

Balancing Gaussian vectors in high dimension

Paxton Turner

*Department of Mathematics
Massachusetts Institute of Technology
77 Massachusetts Avenue,
Cambridge, MA 02139-4307, USA*

PAX@MIT.EDU

Raghu Meka

*Department of Computer Science
UCLA and Massachusetts Institute of Technology
3732H Boelter Hall
Los Angeles, CA 90095, USA*

RAGHUVARDHAN@GMAIL.COM

Philippe Rigollet

*Department of Mathematics
Massachusetts Institute of Technology
77 Massachusetts Avenue,
Cambridge, MA 02139-4307, USA*

RIGOLLET@MATH.MIT.EDU

Editors: Jacob Abernethy and Shivani Agarwal

Abstract

Motivated by problems in controlled experiments, we study the discrepancy of random matrices with continuous entries where the number of columns n is much larger than the number of rows m . Our first result shows that if $\omega(1) = m = o(n)$, a matrix with i.i.d. standard Gaussian entries has discrepancy $\Theta(\sqrt{n} 2^{-n/m})$ with high probability. This provides sharp guarantees for Gaussian discrepancy in a regime that had not been considered before in the existing literature. Our results also apply to a more general family of random matrices with continuous i.i.d entries, assuming that $m = O(n/\log n)$. The proof is non-constructive and is an application of the second moment method. Our second result is algorithmic and applies to random matrices whose entries are i.i.d. and have a Lipschitz density. We present a randomized polynomial-time algorithm that achieves discrepancy $e^{-\Omega(\log^2(n)/m)}$ with high probability, provided that $m = O(\sqrt{\log n})$. In the one-dimensional case, this matches the best known algorithmic guarantees due to Karmarkar–Karp. For higher dimensions $2 \leq m = O(\sqrt{\log n})$, this establishes the first efficient algorithm achieving discrepancy smaller than $O(\sqrt{m})$.

Keywords: Controlled experiments, covariate balance, discrepancy, random matrix, second moment method, number partitioning, greedy algorithm

1. Introduction

Randomized controlled experiments are often dubbed the “gold standard” for estimating treatment effects because of their ability to create a treatment and a control group that have the same features on average. Indeed, pure randomization, i.e., assigning each observation uniformly at random between the treatment and control group, leads to two groups with approximately the same size, the same average age, the same average height, etc. Unfortunately, because of random fluctuations, this approach may not lead to the best

balance between the attributes of the control group and those of the treatment group. Yet, near perfect balance is highly desirable since it often leads to a more accurate estimator of the treatment effect. This quest for balance was initiated at the dawn of controlled experiments. Indeed, W.S. Gosset, a.k.a Student (of t -test fame) already questioned the use of pure randomization when it leads to unbalanced covariates (Student, 1938), and R.A. Fisher proposed randomized block designs as a better solution in certain cases (Fisher, 1935). One traditional approach to overcome this limitation is to simply *rerandomize* the allocation until the generated assignment is deemed balanced enough (Morgan and Rubin, 2012; Li et al., 2018). Rerandomization is effectively a primitive form of optimization that consists in keeping the best of several random solutions. However, it was not until recently that covariate balancing was recognized for the combinatorial optimization problem that it really is. With this motivation, Bertsimas et al. (2015); Kallus (2018) proposed algorithms based on mixed integer programming that, while flexible, did not come with theoretical guarantees. More recently, Harshaw et al. (2019) used new algorithms from Bansal et al. (2018) with theoretical guarantees to generate experimental designs with a tunable degree of randomization versus covariate balance and characterized the resulting trade-off between model robustness and efficiency for a specific treatment effect estimator computed on data collected in such experiments.

In this work, we investigate both the theoretical and algorithmic aspects associated to this question by framing it in the broader scope of *vector balancing*. In particular, this question bears strong theoretical footing in discrepancy theory.¹

Let $X_1, \dots, X_n \in \mathbb{R}^m$ denote a collection of vectors and let \mathbf{X} denote the $m \times n$ matrix whose column i is X_i . The *discrepancy* $\mathcal{D}(X_1, \dots, X_n)$ of this collection is defined as follows.²

$$\mathcal{D}_n := \mathcal{D}(X_1, \dots, X_n) = \min_{\sigma \in \{\pm 1\}^n} \left| \sum_{i=1}^n \sigma_i X_i \right|_{\infty} = \min_{\sigma \in \{\pm 1\}^n} |\mathbf{X}\sigma|_{\infty} \quad (1.1)$$

Discrepancy theory is a rich and well-studied area with applications to combinatorics, optimization, geometry, and statistics, among many others (see the comprehensive texts Matoušek, 1999; Chazelle, 2000). A fundamental result in the area due to Spencer (1985) states that if $\max_i |X_i|_{\infty} \leq 1$ and $m = n$, then $\mathcal{D}_n \leq 6\sqrt{n}$. Spencer’s proof is nonconstructive and relies on a technique known as *partial coloring*. In the last decade, starting with the breakthrough work of Bansal (2010), several algorithmic versions of the partial coloring method have been introduced to efficiently find a signing σ that approximately attains the minimum in (1.1). These include approaches based on random walks (Bansal, 2010; Lovett and Meka, 2012), random projections (Rothvoss, 2017), and multiplicative weights (Levy et al., 2017). In the regime where $m \geq n$, these algorithms can be used to compute a signing (or allocation) $\sigma \in \{-1, 1\}^n$ with objective value $O(\sqrt{n \log(2m/n)})$. Moreover, this guarantee is tight in the sense that examples are known with discrepancy matching this bound.

The aforementioned results make minimal structural assumptions on the vectors X_1, \dots, X_n and treat the input as worst-case. However, in the context of controlled experiments, it is natural to assume that X_1, \dots, X_n are, in fact, independent copies of a random vector $X \in \mathbb{R}^m$. While more general results are possible, the reader should keep in mind the canonical example where $X \sim \mathcal{N}_m(0, I_m)$ is a standard Gaussian vector, and in particular where the entries of X are of order 1. We dub the study of \mathcal{D}_n in this context *average-case discrepancy*.

It was first shown in Karmarkar et al. (1986) via a nonconstructive application of the second moment method that when $m = 1$, the average-case discrepancy is $\mathcal{D}_n = \Theta(\sqrt{n} 2^{-n})$ with high probability, assum-

1. The recent work Harshaw et al. (2019) takes a similar point of view, though here our purpose is to focus purely on optimal covariate balance.
 2. In the interest of clarity, we free ourselves from important considerations in the practical design of controlled experiments such as having two groups of exactly the same size.

ing that the underlying distribution has a sufficiently regular density. This result was extended to specific multidimensional regimes. First, Costello (2009) showed that $\mathcal{D}_n = \Theta(\sqrt{n} 2^{-n/m})$ in the constant dimension regime $m = O(1)$. The optimal discrepancy is also known in the super-linear regime $m \geq 2n$ where it was shown that $\mathcal{D}_n = O(\sqrt{n \log(2m/n)})$.³ In particular, there is a striking gap between this benchmark and the discrepancy $|\mathbf{X}\sigma^{\text{rdm}}|_\infty = \Theta(\sqrt{n \log m})$ achieved by a random signing σ^{rdm} , especially in the sub-linear regime. Motivated by applications to controlled experiments, Krieger et al. (2019) studied the average-case discrepancy problem with the aim to improve on this gap. The authors devised a simple and efficient greedy scheme that, in the univariate case, outputs an allocation σ^{gree} satisfying $|\mathbf{X}\sigma^{\text{gree}}| = O(n^{-2})$. In addition, Krieger et al. (2019) argue that $|\mathbf{X}\sigma^{\text{gree}}| = O(n^{-2/m})$ for any *constant* dimension m .

This state of the art leaves three important questions open:

1. Can a sub-polynomial discrepancy be achieved in polynomial time even in dimension 1?
2. What is the optimal discrepancy in the intermediate regime where $\omega(1) = m = o(n)$?
3. Do there exist efficient allocations that perform better than the random allocation in super-constant dimension?

The answer to the first question is well known. Indeed, the best known algorithm for number partitioning is due to Karmarkar and Karp (1982) and yields $\sigma \in \{-1, 1\}^n$ such that $|\mathbf{X}\sigma|_\infty = e^{-\Omega(\log^2 n)}$ with high probability (see also Boettcher and Mertens, 2008). While this result provides a super-polynomial improvement over algorithms built for the worst case, a significant gap remains between the information-theoretic bounds and the algorithmic ones despite extensive work on the subject (Boettcher and Mertens, 2008; Borgs et al., 2001; Hoberg et al., 2017). This suggests the possibility of a statistical-to-computational gap similar to those that have been observed starting with sparse PCA (Berthet and Rigollet, 2013a,b) and more recently in other planted problems (Brennan et al., 2018; Bandeira et al., 2018). Moreover, while the greedy algorithm of Krieger et al. (2019) is loosely based on ideas from this algorithm, no multivariate extension of this algorithm is known even for the case $m = 2$. Note that in the super-linear regime $m \geq 2n$, the work of Chandrasekaran and Vempala (2014) also proposes a polynomial-time algorithm based on Lovett and Meka (2012) showing an absence of substantial statistical-to-computational gaps.

In this paper, we provide answers to the remaining two questions raised above. First, we show that the discrepancy of standard Gaussian vectors is $\Theta(\sqrt{n} 2^{-n/m})$ with high probability for the remaining regime $\omega(1) = m = o(n)$. Moreover, we complement this existential result by giving the first randomized polynomial-time algorithm that achieves discrepancy $e^{-\Omega(\log^2(n)/m)}$ when $2 \leq m = O(\sqrt{\log n})$. Note that while this remains an intrinsically low-dimensional result, it covers already super-constant dimension. This first algorithmic result paves the way for potential algorithmic advances in a wider range of high-dimensional problems. In particular, our existential result sets an information-theoretic benchmark against which future algorithmic results can be compared as well as a baseline to establish potential statistical-to-computational gaps in high dimensions. These improved discrepancy bounds also have direct applications to randomized control trials. For example, in the case of an additive linear response with all covariates observed, the discrepancy attained by the allocation controls the fluctuations of the difference-in-means treatment effect estimator (Krieger et al., 2019).

Besides discrepancy, another point of view on balancing covariates in randomized trials is that of pairwise matching. In this setup, the experimenter first divides the sample into two equal-sized groups and then pairs up individuals who have similar covariates. The quality of the optimal matching is naturally measured by the Wasserstein-1 distance between the two groups of covariates. For the unidimensional case, Greevy

3. The upper bound established in Chandrasekaran and Vempala (2014) presents additional polylogarithmic terms that are negligible for most of the range $m \geq 2n$. This is also the regime considered by Harshaw et al. (2019).

et al. (2004) proposed a scheme that consists of performing a minimum cost bipartite matching, and this can be implemented in near-linear time using modern tools from computational optimal transport (Altschuler et al., 2017, 2019). In addition, recent results from optimal transport (Ledoux and Zhu, 2019) study the case of i.i.d Gaussian covariates in constant dimension m and derive sharp asymptotic bounds on the Wasserstein-1 distance of the form $n^{-1/m}$. This perspective deserves further study and provides another promising approach for improving the quality of inference in randomized control trials.

2. Main results

In this section, we give an overview of our main results. Detailed computations and proofs are postponed to subsequent sections.

2.1. Existential result

Our first main result shows that when $X_1, \dots, X_n \stackrel{iid}{\sim} \mathcal{N}(0, I_m)$ and $\omega(1) = m = o(n)$, then the discrepancy is asymptotically $\sqrt{\frac{\pi n}{2}} 2^{-n/m}$ with high probability. As in the one-dimensional case (Karmarkar et al., 1986), this result highlights that drastic cancellations are possible, with high probability, when the number of vectors grows asymptotically faster than the dimension.

Theorem 1 *Fix an absolute constant $\gamma > 1$ and suppose that $\omega(1) = m = o(n)$. Let $X_1, \dots, X_n \stackrel{iid}{\sim} \mathcal{N}(0, I_m)$ be independent standard Gaussian random vectors. Then*

$$\lim_{n \rightarrow \infty} \mathbb{P} \left[\mathcal{D}(X_1, \dots, X_n) \leq \gamma \sqrt{\frac{\pi n}{2}} 2^{-n/m} \right] = 1. \quad (2.2)$$

If $\gamma' < 1$, then

$$\lim_{n \rightarrow \infty} \mathbb{P} \left[\mathcal{D}(X_1, \dots, X_n) \geq \gamma' \sqrt{\frac{\pi n}{2}} 2^{-n/m} \right] = 1. \quad (2.3)$$

The work of Costello (2009) handles the case $m = O(1)$, and shows that the limiting probability in (2.2) is exactly $1 - \exp(-2\gamma^m)$. We also note that the series of papers by Borgs et al. (2001, 2008a,b) provides an even more complete description of the unidimensional case.

Our results are not limited specifically to Gaussian distributions. A mild extension of our techniques allows us to derive a similar result for a more general family of distributions, assuming that $m = O(n/\log n)$.

Remark 2 *Let $C > 0$ denote a sufficiently small absolute constant, and suppose that $m \leq Cn/\log n$. Let \mathbf{X} denote an $m \times n$ random matrix whose entries are i.i.d random variables having a common density $f : \mathbb{R} \rightarrow \mathbb{R}$ such that*

$$\int f(x)^2 dx < \infty, \quad \int x^4 f(x) dx < \infty, \quad \text{and} \quad f(x) = f(-x), \forall x \in \mathbb{R}.$$

Then there exist absolute positive constants $c \leq c'$ such that

$$\lim_{n \rightarrow \infty} \mathbb{P} \left[c\sqrt{n}2^{-n/m} \leq \mathcal{D}(X_1, \dots, X_n) \leq c'\sqrt{n}2^{-n/m} \right] = 1.$$

We omit the proof of the above remark and focus on the Gaussian case for simplicity and because for Gaussian vectors, our analysis covers the whole range $m = o(n)$.

The proof of the upper bound in Theorem 1 is a nonconstructive application of the second moment method, in a similar spirit to the analysis of [Karmarkar et al. \(1986\)](#) on the one-dimensional case as well as Achlioptas–Moore’s analysis of the threshold for random k -SAT ([Achlioptas and Moore, 2002](#)). Recall that the second moment method states that for a nonnegative random variable S , we have

$$\mathbb{P}[S > 0] \geq \frac{\mathbb{E}[S]^2}{\mathbb{E}[S^2]}. \quad (2.4)$$

As described in more detail in Section 3, our strategy is to let S count the number of signings with discrepancy at most $\gamma 2^{-n/m} \sqrt{\pi n/2}$ and show that the right-hand-side of (2.4) tends to 1 asymptotically. We also note that the lower bound in Theorem 1 is a straightforward consequence of the Markov inequality (first moment method) applied to S (see Proposition 7).

In addition to our result for $m = o(n)$, using similar techniques we also provide a precise characterization of Gaussian discrepancy in the linear regime $m \leq \delta n$, where δ is a sufficiently small absolute constant. In Appendix B, we show that the discrepancy is $\Theta(\sqrt{n} 2^{-1/\delta})$ with probability at least 99%, asymptotically as $n \rightarrow \infty$. This provides further evidence of a conjecture of [Aubin et al. \(2019\)](#) that the discrepancy when $m = \delta n$ is asymptotically $c(\delta)\sqrt{n}$ with high probability for an explicit function $c(\delta)$.⁴ In particular, our result combined with those of [Chandrasekaran and Vempala \(2014\)](#) confirms that the discrepancy is $\Theta(c(\delta)\sqrt{n})$ with asymptotic probability at least 99% when $m = \delta n$ for all $\delta > 0$.

Complementary to our work, we discuss recent existential results on average-case discrepancy in the discrete case when X_1, \dots, X_n are i.i.d vectors in $\{0, 1\}^m$. Extending prior work of [Ezra and Lovett \(2016\)](#), [Franks and Saks \(2018\)](#) and [Hoberg and Rothvoss \(2018\)](#) use a nonconstructive Fourier-analytic argument to show, for two different models of random sparse binary vectors, that the discrepancy is $O(1)$ if $n = \tilde{\Omega}(m^3)$ ([Franks and Saks, 2018](#)) and $n = \tilde{\Omega}(m^2)$ ([Hoberg and Rothvoss, 2018](#)). In addition, for the continuous case, [Franks and Saks \(2018\)](#) show that the discrepancy of random unit vectors is $O(\exp(-\sqrt{n/m^3}))$. [Potukuchi \(2018\)](#) uses the second moment method to show the discrepancy is $O(1)$ if $n = \Omega(m \log m)$ in the specific case where the entries of X_1 are uniform on $\{0, 1\}$. In other recent work, [Bansal and Meka \(2019\)](#) establish an average-case version of the Beck–Fiala conjecture, giving an algorithmic proof that the discrepancy of uniformly random t -sparse binary vectors is at most $O(\sqrt{t})$ for the entire range of parameters m, n if $t = \Omega(\log \log m)$. It is an open question as to whether there exists a polynomial-time algorithm achieving $O(1)$ discrepancy for random $\{-1, +1\}$ vectors or sparse $\{0, 1\}$ vectors with $n = \text{poly}(m)$ ([Hoberg and Rothvoss, 2018](#); [Franks and Saks, 2018](#)).

2.2. Algorithmic result

Our second main result is algorithmic and applies to a large family of continuous distributions. We construct a randomized polynomial-time algorithm called Generalized Karmarkar–Karp (**GKK**) that achieves discrepancy $\exp(-\Omega(\log^2(n)/m))$ with high probability, assuming $m = O(\sqrt{\log n})$. This establishes the first such efficient algorithm achieving quasi-polynomially-small discrepancy for this regime. Our algorithm and analysis extend those of [Karmarkar and Karp \(1982\)](#) in the one-dimensional case to higher dimensions.⁵

4. See Appendix B for a more precise description of their results.

5. [Karmarkar and Karp \(1982\)](#) give two algorithms for number partitioning. The first one is a simple greedy heuristic, but its analysis was only performed for the uniform distribution over a decade later by [Yakir \(1996\)](#). Our algorithm presented here generalizes the second one which was rigorously analyzed in the original paper of [Karmarkar and Karp \(1982\)](#).

Theorem 3 *Let \mathbf{X} denote a random $m \times n$ matrix with iid entries having a common density $\rho : [-\Delta, \Delta] \rightarrow \mathbb{R}$ which is L -Lipschitz and bounded above by some constant $D > 0$. Suppose that*

$$m \leq C \sqrt{\frac{\log n}{\max(1, \log \Delta)}},$$

*for some sufficiently small absolute constant $C = C(D, L) > 0$. Then the algorithm **GKK** outputs, in polynomial time, a signing $\sigma \in \{-1, +1\}^n$ such that*

$$|\mathbf{X}\sigma|_\infty \leq \exp\left(-\frac{c \log^2 n}{m}\right),$$

with probability at least $1 - \exp(-cn^{1/4})$ for some absolute constant $c > 0$.

This result easily extends to distributions with unbounded support. For example, if \mathbf{X} has i.i.d standard Gaussian entries, then setting $\Delta = O(\sqrt{\log n})$ and conditioning on the (high probability) event $\{|\mathbf{X}_{ij}| \leq \Delta \forall i, j\}$, we can apply Theorem 3 to show that **GKK** yields discrepancy $\exp(-c \log^2(n)/m)$ for the Gaussian matrix \mathbf{X} .

It is an open question as to whether or not the guarantee of Theorem 3 can be improved to achieve sub-quasi-polynomial discrepancy efficiently, even in dimension one. Note that for $m = 1$, [Hoberg et al. \(2017\)](#) provide evidence of hardness of a $O(2^{\sqrt{n}})$ -approximation to the optimal discrepancy in worst case via a reduction from the Minkowski problem and the shortest vector problem. We leave the following question.

Question 1 *Suppose that $m = n^\gamma$ for some $\gamma \in (0, 1)$. Let \mathbf{X} denote a random $m \times n$ matrix with independent standard Gaussian entries. What is the smallest possible value of $|\mathbf{X}\sigma|_\infty$ that can be achieved algorithmically in polynomial time?*

In particular, it is an open problem as to whether the partial coloring method can be used to guarantee subconstant discrepancy for standard Gaussians when $m = n^\gamma$. We suspect that the answer is negative. It seems that even attaining discrepancy $o(\sqrt{m})$ serves as a natural bottleneck for such an approach.

3. Gaussian discrepancy in sub-linear dimension

The main goal of this section is to prove the following proposition. Throughout, we adopt the shorthand notation $u_n \lesssim_n v_n$ for $u_n \leq v_n(1 + o(1))$ and $u_n \simeq_n v_n$ for $u_n = v_n(1 + o(1))$.

Proposition 4 *Fix $\gamma > 1$, $\omega(1) = m = o(n)$, and let $X_1, \dots, X_n \stackrel{iid}{\sim} \mathcal{N}(0, I_m)$ be independent standard Gaussian random vectors. Then*

$$\lim_{n \rightarrow \infty} \mathbb{P}\left[\mathcal{D}(X_1, \dots, X_n) \leq \gamma \sqrt{\frac{\pi n}{2}} 2^{-n/m}\right] = 1.$$

We first outline our proof strategy based on the second moment method. Set $\varepsilon = \varepsilon(n) = \gamma 2^{-n/m} \sqrt{\pi n/2}$ and define S , the number of low discrepancy solutions, to be

$$S = \sum_{\sigma \in \{\pm 1\}^n} \mathbb{I}\left(\left|\sum_{i=1}^n \sigma_i X_i\right|_\infty \leq \varepsilon\right). \quad (3.5)$$

Our goal is to show that $\mathbb{E}[S^2]/\mathbb{E}[S]^2 = 1 + o(1)$. By the second moment method (2.4), this implies the desired result.

The next lemma gives a useful form for the first and second moments of S and follows from a straightforward calculation. Its proof is postponed to Appendix A.

Lemma 5 *The random variable S defined as in (3.5) has its first two moments given by*

$$\mathbb{E}[S] = 2^n \mathbb{P}\left(|Z| \leq \frac{\varepsilon}{\sqrt{n}}\right)^m \quad (3.6)$$

where $Z \sim \mathcal{N}(0, 1)$, and

$$\mathbb{E}[S^2] = 2^n \sum_{k=0}^n \binom{n}{k} \mathbb{P}_{\rho_k} (|\sqrt{n}X| \leq \varepsilon, |\sqrt{n}Y| \leq \varepsilon)^m. \quad (3.7)$$

Here $\rho_k = 1 - 2k/n$ and \mathbb{P}_{ρ_k} denotes the joint distribution of (X, Y) with $X, Y \sim \mathcal{N}(0, 1)$ having correlation ρ_k .




Figure 1: $\alpha \mapsto \phi_n(\alpha)$ for $n = 1000$, $m = \lfloor \sqrt{1000} \rfloor$, and $\varepsilon = 1/n$.

Given this representation, we proceed in two steps to prove an upper bound on the second moment $\mathbb{E}[S^2]$:

- (i) We first apply a truncation argument to show that the contribution from the $k \leq n/4$ and $k \geq 3n/4$ terms in the summand of (3.7) is negligible. See Lemma 14 and its proof in Appendix A for details.
- (ii) Then we show that the dominant contribution in the summation (3.7) is asymptotically bounded by $\mathbb{E}[S]^2$ and comes from an interval of length $\Theta(\sqrt{n})$ around $k \simeq n/2$. This part is somewhat delicate and we apply the *Laplace method* to obtain sharp bounds.

By step (i), it suffices to control the leading term

$$L := 2^n \sum_{k=n/4}^{3n/4} \binom{n}{k} \mathbb{P}_{\rho_k} (|\sqrt{n}X| \leq \varepsilon, |\sqrt{n}Y| \leq \varepsilon)^m. \quad (3.8)$$

To that end, approximate the above binomial coefficient using Lemma C.2 in Berthet et al. (2018): For any $l \in (0, 1/2]$, $\alpha \in (l, 1-l)$ such that $n\alpha$ is an integer, it holds

$$\exp\left(-\frac{1}{12l^2n}\right) \leq \sqrt{2\pi n\alpha(1-\alpha)} \exp(-nh(\alpha)) \binom{n}{\alpha n} \leq \exp\left(\frac{1}{12n}\right),$$

where $h(\alpha) = -\alpha \log \alpha - (1-\alpha) \log(1-\alpha)$ denotes the binary entropy with $h(0) = h(1) = 0$. Therefore, it holds that

$$L \lesssim_n \frac{2^n}{\sqrt{2\pi n}} \sum_{k=n/4}^{3n/4} \exp(\phi_n(\alpha_k)) \quad (3.9)$$

where $\alpha_k = k/n$ and

$$\phi_n(\alpha) = nh(\alpha) + m \log(\mathbb{P}_{1-2\alpha} [|\sqrt{n}X| \leq \varepsilon, |\sqrt{n}Y| \leq \varepsilon]) - \frac{1}{2} \log \alpha(1-\alpha). \quad (3.10)$$

Moreover, as justified in Lemma 15 (see Appendix A), for n sufficiently large, $\phi_n(\alpha)$ is a strictly concave function on $[0.25, 0.75]$ with a unique maximum at $\alpha = 0.5$. See Figure 1 for the graph of $\phi_n(\alpha)$ for a specific setting of the parameters. Thus we can make the Riemann sum approximation

$$L \lesssim_n \frac{2^n}{\sqrt{2\pi n}} \sum_{k=n/4}^{3n/4} \exp(\phi_n(\alpha_k)) \lesssim_n \frac{\sqrt{n}2^n}{\sqrt{2\pi}} \int_{1/4}^{3/4} \exp(\phi_n(\alpha)) d\alpha. \quad (3.11)$$

Our goal now is to employ the Laplace method (see, e.g., Murray, 1984), a well-known technique from asymptotic analysis, to compute explicitly the asymptotic growth of the right-hand-side above. It consists in performing a second-order Taylor expansion of ϕ_n in order to reduce the problem to the computation of a Gaussian integral.

Lemma 6 *Suppose that $m = o(n)$ and set $\varepsilon = \gamma 2^{-n/m} \sqrt{n\pi/2}$. Recall the definition of S from (3.5). Then*

$$L \lesssim_n \mathbb{E}[S]^2. \quad (3.12)$$

Proof We apply the Laplace method to the integral in (3.11). Let $\eta \in (0, 1)$ be arbitrary, and define $g_n(\alpha) = \phi_n(\alpha)/n$. Since $h''(\alpha)$ is continuous, Lemma 15 implies that there exists $\delta = \delta(\eta)$ and $N = N(\eta)$ such that

$$\frac{1}{n} |\phi_n''(\alpha) - \phi_n''(1/2)| \leq \eta, \quad \forall \alpha \in (1/2 - \delta, 1/2 + \delta), n \geq N. \quad (3.13)$$

The above inequality follows by writing $g_n''(\alpha) = h''(\alpha) + r_n(\alpha)$, where $r_n(\alpha)$ is a remainder term that goes to 0 uniformly in $\alpha \in (0.25, 0.75)$ as $n \rightarrow \infty$, using Lemma 15. Using that the remainder term is small and $h''(\alpha)$ is continuous at $\alpha = 1/2$, we arrive at (3.13).

By (3.13) and Taylor's theorem,

$$\phi_n(\alpha) - \phi_n(1/2) \leq \frac{1}{2}(\phi_n''(1/2) + \eta n)(\alpha - 1/2)^2, \quad \forall \alpha \in (1/2 - \delta, 1/2 + \delta), n \geq N. \quad (3.14)$$

Moreover,

$$\phi_n''(1/2) + \eta n < 0 \quad (3.15)$$

for n sufficiently large because $\eta \in (0, 1)$ and $\phi_n''(1/2) \simeq_n -4n$ by Lemma 15. Therefore, since ϕ_n is increasing on $(0.25, 0.75)$ for n sufficiently large,

$$\begin{aligned} \frac{\sqrt{n}}{\exp(\phi_n(1/2))} \int_{1/4}^{1/2-\delta} \exp(\phi_n(\alpha)) d\alpha &\lesssim_n 10\sqrt{n} \exp(\phi_n(1/2 - \delta) - \phi_n(1/2)) \\ &\lesssim_n 10\sqrt{n} \exp\left(\frac{1}{2}(\phi_n''(1/2) + \eta n)\delta^2\right) = o(1), \end{aligned} \quad (3.16)$$

where we applied (3.14) and (3.15). By symmetry of $\phi_n(\alpha)$ about $\alpha = 1/2$, the integral as in (3.16) from $1/2 + \delta$ to $3/4$ is negligible. Moreover, by (3.14),

$$\begin{aligned} \int_{1/2-\delta}^{1/2+\delta} \exp(\phi_n(\alpha)) d\alpha &\lesssim_n \int_{1/2-\delta}^{1/2+\delta} \exp\left(\phi_n(1/2) + \frac{1}{2}(\phi_n''(1/2) + \eta n)(\alpha - 1/2)^2\right) d\alpha \\ &\lesssim_n \exp(\phi_n(1/2)) \sqrt{\frac{2\pi}{|\phi_n''(1/2) + \eta n|}} = 2^n f_n^m \sqrt{\frac{2\pi}{n(1 - \eta/4)}}, \end{aligned} \quad (3.17)$$

where

$$f_n = \mathbb{P}_0(|\sqrt{n}X| \leq \varepsilon, |\sqrt{n}Y| \leq \varepsilon).$$

Since $\eta \in (0, 1)$ was arbitrary, we conclude by (3.6), (3.9), (3.11), (3.16), (3.16), and the definition of f_n that

$$L \lesssim_n \frac{2^n}{\sqrt{2\pi n}} \cdot n \cdot \int_{1/4}^{3/4} \exp(\phi_n(\alpha)) d\alpha \lesssim_n 2^{2n} f_n^m = \mathbb{E}[S]^2.$$

Proof [Proof of Proposition 4] We see that $\mathbb{E}[S^2]/\mathbb{E}[S]^2 \lesssim_n 1$ as $n \rightarrow \infty$ applying Lemma 5, Lemma 14, (3.8), (3.9), and Lemma 6. Proposition 4 follows by the second moment method. ■

We complement Proposition 4 with a near-matching lower bound.

Proposition 7 *Let $\omega(1) = m = o(n)$, fix $\gamma < 1$, and let $X_1, \dots, X_n \stackrel{iid}{\sim} \mathcal{N}(0, I_m)$ be independent standard Gaussian random vectors. Then*

$$\lim_{n \rightarrow \infty} \mathbb{P} \left[\mathcal{D}(X_1, \dots, X_n) \leq \gamma \sqrt{\frac{\pi n}{2}} 2^{-n/m} \right] = 0.$$

Proof Recall the definition of S as in (3.5), which counts the number of signings with discrepancy $\varepsilon = \gamma 2^{-n/m} \sqrt{\pi n/2}$. By the Markov inequality, (A.20), and (3.6),

$$\mathbb{P}[S > 1] \leq \mathbb{E}[S] = 2^n \mathbb{P} \left[|Z| < \gamma \sqrt{\frac{\pi n}{2}} 2^{-n/m} \right]^m \lesssim_n \gamma^m \rightarrow 0$$

because $\omega(1) = m = o(n)$ and $\gamma < 1$. This completes the proof. ■

Our first main result, Theorem 1, is a direct consequence of Propositions 4 and 7.

4. Algorithmic discrepancy minimization in low dimension

Now we describe our approach for proving Theorem 3. In this section we introduce the generalized Karmarkar–Karp algorithm **GKK**. Recall that the goal is to find algorithmically $\sigma \in \{\pm 1\}^n$ such that $|\mathbf{X}\sigma|_\infty$ is small. As in Karmarkar and Karp (1982), our algorithm is a *differencing method*, which means that throughout the algorithm, we maintain a set of vectors S , and our basic operations consist of removing two vectors, say x and y , from S and then adding the difference to S : $S \leftarrow S \cup \{x - y\} \setminus \{x, y\}$. We perform a sequence of these differencing operations in a judicious way until there is a single vector v remaining in S . Note that at any given time, the elements of S correspond to (disjoint) partial signed sums of the original vectors X_1, \dots, X_n . Hence, the final vector $v \in S$ is indeed a signed sum of the original vectors. It is possible to keep track of the final signing by tracking the differences, though we do not do so explicitly.

Next, we informally describe the **GKK** differencing method in detail. For simplicity, we assume that $\Delta = 1$ in this description. The algorithm **GKK** is a recursive procedure that consists of $\Theta(\log n)$ phases. For the first phase of the recursion, given a collection of n vectors lying in $[-1, 1]^m$, we partition this cube into sub-cubes of side length $\alpha = n^{-\Omega(1/m)}$. The idea is that with sub-cubes of this size, we are likely to have multiple points in each sub-cube, and these points would be very close to each other. We then randomly difference the vectors in each sub-cube until there is at most one point left in each sub-cube. Next, we enter a *clean-up* step to deal with the leftover vectors. First we combine the leftover vectors (at most one per each

sub-cube) via a standard differencing algorithm that we call **REDUCE** into a single ‘bad’ vector $v^{(0)}$ and let $G' \subset [-\alpha, \alpha]^m$ denote the vectors formed from random differencing. Next we make the entries of the bad vector small by adding signed combinations of a few vectors from G' . Namely, we draw at random points from G' and greedily difference them against $v^{(0)}$ until the resulting vector is sufficiently small in the Euclidean norm. Specially, our update procedure for this clean-up step is

$$\begin{aligned} v^{(k)} &= v^{(k-1)} + a^* \mathbf{u}_k \\ a^* &= \operatorname{argmin}_{a \in \{\pm 1\}} \left| v^{(k-1)} + a \mathbf{u}_k \right|_2. \end{aligned} \quad (4.18)$$

where \mathbf{u}_k is drawn at random from the remaining vectors in G' .

Once we have $v^{(k)} \in [-O_m(\alpha), O_m(\alpha)]^m$, we stop drawing random vectors from G' , and this ends the first phase of recursion. The remaining vectors form the input to the second phase, which applies the same procedure as above on the smaller cube $[-\alpha, \alpha]^m$. Moreover, subsequent phases follow the same pattern: **partition**, **difference**, and **clean-up**. After each phase, the input cube shrinks by a factor of $n^{-\Omega(1/m)}$. Hence, after a logarithmic number of phases, the remaining vectors lie in a cube of side length $n^{-\Omega(n/m)} = e^{-\Omega(\log^2 n/m)}$. We then apply **REDUCE** to combine the remaining vectors into a single vector with discrepancy as in Theorem 3.

We remark that our algorithm also features a resampling step that happens immediately after partitioning. In each phase, this resampling procedure labels points as ‘good’ or ‘bad’ so that the good points are independent and have independent coordinates that have a nice distribution. This same resampling trick was also used in [Karmarkar and Karp \(1982\)](#) and is essential for (most of) the remaining random vectors at the end of each phase to have a nice distribution facilitating a recursive analysis. Moreover, the **partition** and **difference** steps of our algorithm are also similar to those used in [Karmarkar and Karp \(1982\)](#) for the one-dimensional case.

In summary, the algorithm **GKK** consists of several phases of a subroutine **PRDC**, which stands for partition, resample, difference, clean-up, that we now define explicitly. In the first part of the clean-up phase, we remark that the aforementioned algorithm **REDUCE** is applied. However, we defer the explicit description of this algorithm, which uses standard techniques, to Appendix C, instead stating its key property of use.

Lemma 8 *Given $X_1, \dots, X_N \in \mathbb{R}^m$, the algorithm **REDUCE** is polynomial-time and outputs $\sigma \in \{\pm 1\}^N$ such that*

$$\left| \sum_{i=1}^N \sigma_i X_i \right|_{\infty} \leq \max_{S \subset [N]: |S|=m} \sum_{j \in S} |X_j|_{\infty}. \quad (4.19)$$

In the explicit description of **PRDC** below, $\gamma > 0$ denotes a fixed absolute constant to be set later (see Appendix E).

PRDC:

Input: A number $\alpha_t > 0$. A set of vectors $S_t \subset [-\alpha_t, \alpha_t]^m$. A single vector $v_t \subset \gamma m [-\alpha_t, \alpha_t]^m$. A pdf $g_t : [-\alpha_t, \alpha_t]^m \rightarrow \mathbb{R}$. Define $N_t = 2^m \lceil |S_t|^{1/(4m)} \rceil^m$.

1. **Partition:** Define $\alpha_{t+1} = \alpha_t / \lceil |S_t|^{1/(4m)} \rceil$. Divide the cube $[-\alpha_t, \alpha_t]^m$ into N_t disjoint sub-cubes C_1, \dots, C_{N_t} that are of the form $\alpha_{t+1} z + [0, \alpha_{t+1}]^m$ for some integer vector $z \in \mathbb{Z}^m$.
2. **Resample:** Independently for every vector x in S_t , if $x \in C_j$, then label x as ‘good’ with probability $(\min_{y \in C_j} g_t(y)) / g_t(x)$. Otherwise, label x to be ‘bad.’ Let G_t denote the set of good points and B_t denote the set of bad points.

3. **Difference:** For every sub-cube C_j , pick uniformly at random two points in $G_t \cap C_j$, include their difference in G'_t , and remove them from G_t . Continue this until $G_t \cap C_j$ has at most 1 good point for every j . Let B'_t be the union of B_t, v_t , and the leftover good points.

4. **Clean-up:**

(a) Apply **REDUCE** to the vectors in B'_t to obtain σ . Define $v_t^{(0)} = \sum_{b_i \in B'_t} \sigma_i b_i$.

(b) For $k = 0, 1, 2, \dots$

If $\left|v_t^{(k)}\right|_2 \geq \gamma m \alpha_{t+1}$: remove uniformly at random a point $x \in G'_t$. Define $v_t^{(k+1)} = v_t^{(k)} + a^* x$ where $a^* = \operatorname{argmin}_{a \in \{\pm 1\}} |v_t^{(k)} + ax|_2$. Define $G'_t \leftarrow G'_t \setminus \{x\}$.

Else: $v_{t+1} := v_t^{(k)}$. **BREAK**

Output: $S_{t+1} := G'_t, v_{t+1}, \alpha_{t+1} := \alpha_t / \lceil |S_t|^{1/(4m)} \rceil$

Now we explicitly describe our main algorithm **GKK** in terms of the subroutine **PRDC**. Recall that ρ is the density corresponding to a particular entry of \mathbf{X} . First we need the following definition.

Definition 9 (Triangular distribution) *A random vector $\mathbf{y} \in \mathbb{R}^m$ follows a triangular distribution on the cube $[-R, R]^m$ if the distribution of \mathbf{y} is given by $\mathbf{u} - \mathbf{v}$, where \mathbf{u} and \mathbf{v} are independent and uniformly distributed on $[0, R]^m$. Notationally, we write $\mathbf{y} \sim \operatorname{Tri}[-R, R]^m$.*

GKK:

Input: An $m \times n$ matrix \mathbf{X} . A probability density function $\rho : [-\Delta, \Delta] \rightarrow \mathbb{R}$. Let $T = \lceil C^* \log n \rceil$ where $C^* := (2 \log(10/3))^{-1}$.

1. Set $S_1 = \operatorname{col}(\mathbf{X}), \alpha_1 = \Delta, v_1 = \mathbf{0}$, and $g_1 = \rho^{\otimes m}$.

2. For $t = 1, 2, \dots, T$:

(a) Run **PRDC** on the input data S_t, v_t, α_t, g_t to output S_{t+1}, v_{t+1} , and α_{t+1} .

(b) Set $g_{t+1}(x) = \frac{1}{\alpha_{t+1}} f(x/\alpha_{t+1})$, where $f(x)$ is the triangular density on $[-1, 1]^m$.

3. Apply **REDUCE** to the vectors in $S_T \cup \{v_T\}$ to obtain σ . Let $v = \sum_{s_i \in S_T \cup \{v_T\}} \sigma_i s_i$.

Output: $|v|_\infty$

We remark that the first three steps of **PRDC** are similar to those in the corresponding subroutine in [Karmarkar and Karp \(1982\)](#) for the one-dimensional case. The clean-up step and its analysis on the other hand are quite different. In particular, we use **REDUCE** to combine the ‘bad’ vectors left over from resampling into a single bad vector $v^{(0)}$. This subroutine is quite similar to the algorithm used by Beck–Fiala to show that t -sparse vectors have discrepancy at most $2t - 1$ ([Beck and Fiala, 1981](#)). In contrast, [Karmarkar and Karp \(1982\)](#) use a greedy iterative algorithm for dealing with bad points in dimension 1, but it is not clear how to generalize their algorithm to also work in higher dimensions. In the next part of the clean-up step, we must bring the bad vector $v^{(0)}$ into a smaller range. [Karmarkar and Karp \(1982\)](#) do this by randomly sampling points from G' and greedily differencing them against $v^{(0)}$ until the resulting number is small. Here we use the same approach, but since we are working in higher dimensions, we measure the resulting vector in the Euclidean norm. In this part of the clean-up step, the key difference between our work and [Karmarkar and Karp \(1982\)](#) lies in our analysis, which includes elements of the analysis of stochastic gradient descent, as well as martingale concentration and the Khintchine inequality (see [Appendix E](#)).

We also comment on the reason for the bound $m = O(\sqrt{\log n})$ in Theorem 3. First observe that by our choice of $\alpha = n^{-\Omega(1/m)}$ for the side-lengths of the sub-cubes at the first phase, it is necessary that $m = O(\log n)$; otherwise the sub-cubes are not smaller than the original cube. The reason we require the stronger condition $m = O(\sqrt{\log n})$ is so that not too many points are labeled ‘bad’ in the resampling step of our algorithm. We direct the reader to Appendix D for the analysis and further discussion.

4.1. Analysis of GKK

The proof of Theorem 3 follows from a sequence of inductive assumptions. Recall that S_t denotes the points input to the t^{th} phase of **PRDC**, excluding the single ‘bad’ vector $v_t \in \gamma m[-\alpha_t, \alpha_t]^m$, where γ is a fixed absolute constant to be determined. Recall that $C^* = (2 \log(10/3))^{-1}$, as set in the definition of **GKK**, and that $\Delta > 0$ is the side length of the cube containing the initial set of vectors S_1 .

Proposition 10 *Let X_1, \dots, X_n be iid random vectors, each having a joint density $g : [-\Delta, \Delta]^m \rightarrow \mathbb{R}$. Consider the output S_t, v_t, α_t that results after the $(t - 1)$ -th phase of **PRDC** in step 2 of **GKK**. Then conditioned on $|S_j| = n_j$ for $1 \leq j \leq t$, we have*

- the n_t points in S_t are iid and follow a triangular distribution on $[-\alpha_t, \alpha_t]^m$, and
- the random vector v_t is independent of the vectors in S_t .

Proposition 10 ensures that the distribution of the output of each phase of recursion is preserved, allowing us to apply induction. At the heart of this result is the following marginal calculation which implies that the good points have a uniform distribution on their respective sub-cubes. Conditioning on $X_1 \in C_1$, if L denotes the label of X_1 as ‘good’ or ‘bad’, then (X_1, L) has a mixed joint density $p(x, \ell)$ where $x \in C_1$ and $\ell \in \{\text{‘good’}, \text{‘bad’}\}$, which by Bayes’ rule satisfies

$$p(x|L = \text{‘good’}) = \frac{p(x, \text{‘good’})}{\mathbb{P}[L = \text{‘good’}]} = \frac{g(x) \cdot \frac{\min_{y \in C_1} g(y)}{g(x)}}{\int_{C_1} p(y, \text{‘good’}) dy} = \frac{1}{\text{Vol}(C_1)},$$

for all $x \in C_1$.

The proofs of Propositions 11 and 12 below are postponed to Appendices D and E, respectively. The former relies on showing that a large fraction of the points input to the t^{th} phase are labeled ‘good’ in the **resample** step, and the latter requires us to show that few of the random differences created in step 3 of **PRDC** are lost in the **clean-up** step.

Proposition 11 *Suppose that $1 \leq t \leq C^* \log n$ and $m \leq C \sqrt{(\log n) / \max(1, \log \Delta)}$, where C is a sufficiently small absolute constant. Then for some fixed θ , conditioned on the events $|S_j| \geq \theta^{j-1} n$ for all $1 \leq j \leq t$, it holds that the set G'_t of random differences created in step 2 of the t^{th} phase of **PRDC** satisfies $|G'_t| \geq \beta |S_t|$ for some fixed β with probability at least $1 - \exp(-c_1 \sqrt{n})$, where $c_1 > 0$ is an absolute constant. In particular, we may set $\theta = 0.3$ and $\beta = 0.4$.*

Proposition 12 *Suppose that $1 \leq t \leq C^* \log n$ and $m \leq C \sqrt{\log n}$, where C is a sufficiently small absolute constant. Then conditioned on the events $|G'_t| \geq \beta |S_t|$ and $|S_j| \geq \theta^{j-1} n$ for $1 \leq j \leq t$, it holds that the set S_{t+1} (the input to the $(t + 1)$ -th iteration of **PRDC**) satisfies $|S_{t+1}| \geq \theta |S_t|$ with probability at least $1 - \exp(-c_2 n^{1/4})$, where $c_2 > 0$ is an absolute constant. In particular, we may choose $\beta = 0.4$ and $\theta = 0.3$.*

The proof of Theorem 3 follows easily from the previous two propositions and is found in Appendix F.

Acknowledgments

We thank Tselil Schramm for useful conversations on discrepancy and the anonymous reviewers for their helpful suggestions.

References

- D. Achiloptas and C. Moore. The asymptotic order of the random k -sat threshold. In *The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings*, Vancouver, BC, Canada, November 2002. IEEE.
- N. Alon and J. Spencer. *The Probabilistic Method*. John Wiley and Sons, Inc., New Jersey, 3 edition, 2008.
- J. Altschuler, J. Weed, and P. Rigollet. Near-linear time approximation algorithms for optimal transport via Sinkhorn iteration. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 1961–1971, 2017.
- J. Altschuler, F. Bach, A. Rudi, and J. Weed. Massively scalable Sinkhorn distances via the Nyström method. In *Advances in Neural Information Processing Systems 32 (NeurIPS 2019)*, 12 2019. To appear.
- B. Aubin, W. Perkins, and L. Zdeborova. Storage capacity in symmetric binary perceptrons. *Journal of Physics A: Mathematical and Theoretical*, 2019.
- A. S. Bandeira, A. Perry, and A. S. Wein. Notes on computational-to-statistical gaps: predictions using statistical physics. *arXiv:1803.11132*, 2018.
- N. Bansal. Constructive algorithms for discrepancy minimization. In *Proceedings of the 2010 IEEE 51st Annual Symposium on Foundations of Computer Science, FOCS '10*, pages 3–10, Washington, DC, USA, 2010. IEEE Computer Society. ISBN 978-0-7695-4244-7. doi: 10.1109/FOCS.2010.7. URL <http://dx.doi.org/10.1109/FOCS.2010.7>.
- N. Bansal and R. Meka. On the discrepancy of random low degree set systems. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 2557–2564, 2019. doi: 10.1137/1.9781611975482.157. URL <https://doi.org/10.1137/1.9781611975482.157>.
- N. Bansal, D. Dadush, S. Garg, and S. Lovett. The gram-schmidt walk: a cure for the banaszczyk blues. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 587–597, 2018. doi: 10.1145/3188745.3188850. URL <https://doi.org/10.1145/3188745.3188850>.
- J. Beck and T. Fiala. “integer-making” theorems. *Discrete Applied Mathematics*, 3(1):1 – 8, 1981. ISSN 0166-218X. doi: [https://doi.org/10.1016/0166-218X\(81\)90022-6](https://doi.org/10.1016/0166-218X(81)90022-6). URL <http://www.sciencedirect.com/science/article/pii/0166218X81900226>.
- Q. Berthet and P. Rigollet. Optimal detection of sparse principal components in high dimension. *Ann. Statist.*, 41(1):1780–1815, 2013a.
- Q. Berthet and P. Rigollet. Complexity theoretic lower bounds for sparse principal component detection. In S. Shalev-Shwartz and I. Steinwart, editors, *COLT 2013 - The 26th Conference on Learning Theory, Princeton, NJ, June 12-14, 2013*, volume 30 of *JMLR W&CP*, pages 1046–1066, 2013b.

- Q. Berthet, P. Rigollet, and P. Srivastava. Exact recovery in the Ising blockmodel. *Annals of Statistics (to appear)*, 2018.
- D. Bertsimas, M. Johnson, and N. Kallus. The power of optimization over randomization in designing experiments involving small samples. *Operations Research*, 63(4):868–876, 2015.
- S. Boettcher and S. Mertens. Analysis of the karmarkar-karp differencing algorithm. *CoRR*, abs/0802.4040, 2008. URL <http://arxiv.org/abs/0802.4040>.
- C. Borgs, J. Chayes, and B. Pittel. Phase Transition and Finite-Size Scaling for the Integer Partitioning Problem. *Random Structures and Algorithms*, 19:247–288, 2001.
- C. Borgs, J. Chayes, S. Mertens, and C. Nair. Proof of the local REM conjecture for number partitioning I: Constant energy scales. *Random Structures and Algorithms*, 34:217–240, December 2008a.
- C. Borgs, J. Chayes, S. Mertens, and C. Nair. Proof of the local REM conjecture for number partitioning II: Growing energy scales. *Random Structures and Algorithms*, 34:241–284, December 2008b.
- M. Brennan, G. Bresler, and W. Huleihel. Reducibility and computational lower bounds for problems with planted sparse structure. In S. Bubeck, V. Perchet, and P. Rigollet, editors, *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pages 48–166. PMLR, 06–09 Jul 2018.
- K. Chandrasekaran and S. S. Vempala. Integer feasibility of random polytopes: Random integer programs. In *Proceedings of the 5th Conference on Innovations in Theoretical Computer Science, ITCS '14*, pages 449–458, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-2698-8. doi: 10.1145/2554797.2554838. URL <http://doi.acm.org/10.1145/2554797.2554838>.
- B. Chazelle. *The Discrepancy Method: Randomness and Complexity*. Cambridge University Press, Cambridge, 2000.
- K. Costello. Balancing Gaussian Vectors. *Israeli Journal of Math*, 172:145–156, 2009.
- E. Ezra and S. Lovett. On the Beck-Fiala Conjecture for Random Set Systems. In K. Jansen, C. Mathieu, J. Rolim, and C. Umans, editors, *APPROX/RANDOM*, volume 60 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 29:1–29:10, Dagstuhl, Germany, 2016. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. doi: 10.4230/LIPIcs.APPROX-RANDOM.2016.29.
- R. Fisher. *The design of experiments*. Oliver and Boyd, Edinburgh, 1935.
- C. Franks and M. Saks. On the discrepancy of random matrices with many columns. *arXiv*, pages 1–24, July 2018.
- R. Greevy, B. Lu, J. H. Silber, and P. Rosenbaum. Optimal multivariate matching before randomization. *Biostatistics*, 5(2):263–275, 2004.
- C. Harshaw, F. Sävje, D. Spielman, and P. Zhang. Balancing covariates in randomized experiments using the gramschmidt walk. *CoRR*, abs/1911.03071, 2019. URL <https://arxiv.org/abs/1911.03071>.
- R. Hoberg and T. Rothvoss. A Fourier-analytic approach for the discrepancy of random set systems. *arXiv*, pages 1–19, July 2018.

- R. Hoberg, H. Ramadas, T. Rothvoss, and X. Yang. Number balancing is as hard as minkowski’s theorem and shortest vector. In *Integer Programming and Combinatorial Optimization - 19th International Conference, IPCO 2017, Waterloo, ON, Canada, June 26-28, 2017, Proceedings*, pages 254–266, 2017. doi: 10.1007/978-3-319-59250-3_21. URL https://doi.org/10.1007/978-3-319-59250-3_21.
- N. Kallus. Optimal a priori balance in the design of controlled experiments. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 80(1):85–112, 2018.
- N. Karmarkar, R. Karp, G. Lueker, and A. Odlyzko. Probabilistic Analysis of Optimum Partitioning. *Journal of Applied Probability*, 23:626–645, September 1986.
- N. Karmarkar and R. Karp. The differencing method of set partitioning. Technical report, University of California, Berkeley, 12 1982.
- A. M. Krieger, D. Azriel, and A. Kapelner. Nearly random designs with greatly improved balance. *Biometrika*, 106(3):695–701, 05 2019. ISSN 0006-3444. doi: 10.1093/biomet/asz026. URL <https://doi.org/10.1093/biomet/asz026>.
- M. Ledoux and J.-X. Zhu. On optimal matching of Gaussian samples III. Available on the first author’s webpage, 2019.
- A. Levy, H. Ramadas, and T. Rothvoss. Deterministic discrepancy minimization via the multiplicative weight update method. In *Integer Programming and Combinatorial Optimization - 19th International Conference, IPCO 2017, Waterloo, ON, Canada, June 26-28, 2017, Proceedings*, pages 380–391, 2017. doi: 10.1007/978-3-319-59250-3_31. URL https://doi.org/10.1007/978-3-319-59250-3_31.
- X. Li, P. Ding, and D. B. Rubin. Asymptotic theory of rerandomization in treatment–control experiments. *Proceedings of the National Academy of Sciences*, 115(37):9157–9162, 2018.
- S. Lovett and R. Meka. Constructive discrepancy minimization by walking on the edges. In *Proceedings of the 2012 IEEE 53rd Annual Symposium on Foundations of Computer Science, FOCS ’12*, pages 61–67, Washington, DC, USA, 2012. IEEE Computer Society. ISBN 978-0-7695-4874-6. doi: 10.1109/FOCS.2012.23. URL <https://doi.org/10.1109/FOCS.2012.23>.
- J. Matoušek. *Geometric Discrepancy: an Illustrated Guide*. Springer, New York, 1999.
- K. L. Morgan and D. B. Rubin. Rerandomization to improve covariate balance in experiments. *Ann. Statist.*, 40(2):1263–1282, 04 2012.
- J. Murray. *Asymptotic Analysis*, volume 48. Springer, New York, 1984.
- A. Potukuchi. Discrepancy in random hypergraph models. *CoRR*, abs/1811.01491, 2018. URL <http://arxiv.org/abs/1811.01491>.
- T. Rothvoss. Constructive discrepancy minimization for convex sets. *SIAM J. Comput.*, 46(1):224–234, 2017. doi: 10.1137/141000282. URL <https://doi.org/10.1137/141000282>.
- J. Spencer. Six Standard Deviations Suffice. *Transactions of the American Mathematical Society*, 289: 679–706, 1985.

Student. Comparison between balanced and random arrangements of field plots. *Biometrika*, 29(3-4):363–379, 1938.

R. Vershynin. *High-Dimensional Probability: An Introduction with Applications in Data Science*. Cambridge University Press, New York, 2018.

B. Yakir. The differencing algorithm ldm for partitioning: A proof of a conjecture of karmarkar and karp. *Math. Oper. Res.*, 21(1):85–99, February 1996. ISSN 0364-765X. doi: 10.1287/moor.21.1.85. URL <http://dx.doi.org/10.1287/moor.21.1.85>.

Appendix A. Proofs from Section 3

First, we calculate the first and second moments of S as defined in (3.5).

Proof [Proof of Lemma 5] Let $X_i^{(j)}$ denote the j th element of the vector X_i . Since these elements are independent, we get

$$\mathbb{E}[S] = \sum_{\sigma \in \{\pm 1\}^n} \prod_{j=1}^m \mathbb{P}\left(\left|\sum_{i=1}^n \sigma_i X_i^{(j)}\right| \leq \varepsilon\right) = 2^n \mathbb{P}\left(|Z| \leq \frac{\varepsilon}{\sqrt{n}}\right)^m$$

where $Z \sim \mathcal{N}(0, 1)$. This completes the proof of (3.6).

To prove (3.7), let $d(\tau, \sigma)$ denotes the Hamming distance between σ and τ . Observe that if τ and σ satisfy $d(\tau, \sigma) = k$, then $X := \frac{1}{\sqrt{n}} \sum_{i=1}^n \sigma_i X_i^{(j)}$ and $Y := \frac{1}{\sqrt{n}} \sum_{i=1}^n \tau_i X_i^{(j)}$ are ρ_k -correlated standard Gaussians random variables. Thus

$$\begin{aligned} \mathbb{E}[S^2] &= \sum_{\sigma, \tau \in \{\pm 1\}^n} \mathbb{P}\left(\left|\sum_{i=1}^n \sigma_i X_i\right|_{\infty} \leq \varepsilon, \left|\sum_{i=1}^n \tau_i X_i\right|_{\infty} \leq \varepsilon\right) \\ &= \sum_{\sigma} \sum_{k=0}^n \sum_{\tau: d(\tau, \sigma)=k} \mathbb{P}_{\rho_k}\left(|\sqrt{n}X| \leq \varepsilon, |\sqrt{n}Y| \leq \varepsilon\right)^m \\ &= 2^n \sum_{k=0}^n \binom{n}{k} \mathbb{P}_{\rho_k}\left(|\sqrt{n}X| \leq \varepsilon, |\sqrt{n}Y| \leq \varepsilon\right)^m, \end{aligned}$$

which proves the lemma. ■

The following small-ball probability estimates are required for the proof of the truncation argument, Lemma 14.

Lemma 13 *Let Z denote a standard Gaussian random variable, and let X, Y denote ρ -correlated standard Gaussian random variables with $\rho \in (-0.5, 0.5)$. Then for $0 < z < 1$, we have for some absolute constant $c > 0$ that*

$$-cz^3 \leq \mathbb{P}[|Z| \leq z] - \sqrt{\frac{2}{\pi}} z \leq 0, \tag{A.20}$$

and for all $z \in (0, \infty)$, we have

$$\mathbb{P}_{\rho}[|X| \leq z, |Y| \leq z] \leq \frac{2}{\pi \sqrt{1 - \rho^2}} z^2. \tag{A.21}$$

Proof Observe that $z \mapsto \mathbb{P}[|Z| \leq z]$ is a concave function for $z \geq 0$. Hence, it lies below the tangent line to this curve at $z = 0$, which is precisely the function $z \mapsto \sqrt{2/\pi}z$. This proves the right-hand-side of (A.20). To prove the left-hand-side, we apply Taylor expansion and observe that for $|z| \leq 1$, it holds that

$$\mathbb{P}[|Z| \leq z] = \sqrt{\frac{2}{\pi}}z - \frac{1}{6}\sqrt{\frac{2}{\pi}}z^3 \pm O(z^5) \geq \sqrt{\frac{2}{\pi}}z - cz^3$$

for some absolute constant $c > 0$.

To prove (A.21), note that the joint density $\psi_\rho(x, y)$ of a pair of standard normal ρ -correlated Gaussians satisfies

$$\psi_\rho(x, y) = \frac{1}{2\pi\sqrt{1-\rho^2}} \exp\left(-\frac{x^2 - 2\rho xy + y^2}{2 - 2\rho^2}\right) \leq \frac{1}{2\pi\sqrt{1-\rho^2}}.$$

The upper bound follows by positive-semidefiniteness of the covariance matrix. Hence, integrating over the rectangle $|x| \leq z, |y| \leq z$ and applying the above upper bound yields the desired result. \blacksquare

Lemma 14 *Suppose that $\omega(1) = m = o(n)$ and let $\varepsilon = \varepsilon(n) = \gamma 2^{-n/m} \sqrt{\pi n/2}$ for some $\gamma > 1$. Then*

$$2^n \sum_{k=0}^{n/4} \binom{n}{k} \mathbb{P}_{\rho_k} (|\sqrt{n}X| \leq \varepsilon, |\sqrt{n}Y| \leq \varepsilon)^m = o(\mathbb{E}[S]^2). \quad (\text{A.22})$$

$$2^n \sum_{k=3n/4}^n \binom{n}{k} \mathbb{P}_{\rho_k} (|\sqrt{n}X| \leq \varepsilon, |\sqrt{n}Y| \leq \varepsilon)^m = o(\mathbb{E}[S]^2). \quad (\text{A.23})$$

Proof Note that (A.23) follows from (A.22) by symmetry, so it suffices to prove (A.22). We may write $m = n/g_n$ for some sequence g_n such that $\omega(1) = g_n = o(n)$. For notational convenience, define

$$f_n(\rho) = \mathbb{P}_\rho(|\sqrt{n}X| \leq \varepsilon, |\sqrt{n}Y| \leq \varepsilon).$$

By Lemma 5, we have

$$\begin{aligned} & \frac{2^n \sum_{k=0}^{n/4} \binom{n}{k} \mathbb{P}_{\rho_k} (|\sqrt{n}X| \leq \varepsilon, |\sqrt{n}Y| \leq \varepsilon)^m}{\mathbb{E}[S]^2} \\ &= \underbrace{\sum_{k=0}^{n/(g_n)^2} \frac{\binom{n}{k}}{2^n} \left(\frac{f_n(\rho_k)}{f_n(0)}\right)^m}_{=:A} + \underbrace{\sum_{k=n/(g_n)^2}^{n/4} \frac{\binom{n}{k}}{2^n} \left(\frac{f_n(\rho_k)}{f_n(0)}\right)^m}_{=:B}. \end{aligned} \quad (\text{A.24})$$

For ε as above and $Z \sim N(0, 1)$, we have by applying (A.20) that

$$2^n \mathbb{P}(|Z| < \varepsilon/\sqrt{n})^m \geq 2^n \left(\sqrt{\frac{2}{\pi n}}\varepsilon\right)^m (1 - c\varepsilon^2/n)^m \gtrsim_n \left(\frac{\gamma+1}{2}\right)^m, \quad (\text{A.25})$$

where c is an absolute constant. To obtain the right-hand-side, note that $\varepsilon/\sqrt{n} \xrightarrow{n \rightarrow \infty} 0$ since $m = o(n)$. Thus, for n sufficiently large it holds that

$$1 - c\varepsilon^2/n \geq \frac{1}{2} \left(1 + \frac{1}{\gamma}\right),$$

which yields the right-hand-side of (A.25). Now using the crude bound $f_n(\rho_k) \leq \mathbb{P}(|\sqrt{n}Z| \leq \varepsilon)$, (A.25), the fact that $f_n(0) = \mathbb{P}(|\sqrt{n}Z| \leq \varepsilon)^2$, and the inequality

$$\sum_{k=0}^j \binom{n}{k} \leq \left(\frac{ne}{j}\right)^j,$$

we have

$$\begin{aligned} A &= \sum_{k=0}^{n/(g_n)^2} \frac{\binom{n}{k}}{2^n} \left(\frac{f_n(\rho_k)}{f_n(0)}\right)^m \\ &\lesssim_n \left(\frac{\gamma+1}{2}\right)^{-m} (e g_n^2)^{n/g_n^2} \\ &= \exp\left(-\frac{n \log \frac{1}{2}(1+\gamma)}{g_n} + \frac{n}{g_n^2} + \frac{2n \log g_n}{g_n^2}\right) = o(1) \end{aligned} \quad (\text{A.26})$$

because $(1/2)(1+\gamma) > 1$, $g_n \rightarrow \infty$, and $n/g_n \rightarrow \infty$ as $n \rightarrow \infty$.

By (A.20) and (A.21) (noting again that $f_n(0) = \mathbb{P}(|\sqrt{n}Z| \leq \varepsilon)^2$), we have

$$B = \sum_{k=n/(g_n)^2}^{n/4} \frac{\binom{n}{k}}{2^n} \left(\frac{f_n(\rho_k)}{f_n(0)}\right)^m \lesssim_n (c')^m \sum_{k=n/(g_n)^2}^{n/4} \frac{\binom{n}{k}}{2^n} \left(\frac{n^2}{k(n-k)}\right)^{m/2} \quad (\text{A.27})$$

where c' is an absolute constant. By the Hoeffding bound, letting c'' denote another absolute constant, we have

$$(\text{A.27}) \lesssim_n (c'')^m g_n^m e^{-n/8} = \exp\left(\frac{n \log c''}{g_n} + \frac{n \log g_n}{g_n} - \frac{n}{8}\right) = o(1)$$

since $g_n \rightarrow \infty$. Since $A, B = o(1)$, we conclude by (A.24) that (A.22) holds, as desired. ■

Lemma 15 *Suppose that $m = o(n)$ and set $\varepsilon = \gamma 2^{-n/m} \sqrt{n\pi/2}$. Then the function $\alpha \mapsto \phi_n(\alpha)$ defined in (3.10) is asymptotically strictly concave on $(0.25, 0.75)$. More precisely,*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \frac{\partial^2}{\partial \alpha^2} \phi_n(\alpha) = -\frac{1}{\alpha(1-\alpha)} < -4, \quad \forall \alpha \in (0.25, 0.75),$$

and the convergence is uniform over $\alpha \in (0.25, 0.75)$. Moreover, for n large enough, $\phi_n(\alpha)$ has a unique maximum over $(0.25, 0.75)$ located at $\alpha = 0.5$.

Proof Because $|\partial_\alpha^2 \log \alpha(1-\alpha)| = O(1)$ for $\alpha \in (0.25, 0.75)$, $m = o(n)$, and

$$h''(\alpha) = -\frac{1}{\alpha(1-\alpha)},$$

to verify the strict concavity of $\phi_n(\alpha)$, it suffices to show that

$$\left| \frac{\partial^2}{\partial \alpha^2} \log \mathbb{P}_{1-2\alpha} [|\sqrt{n}X| \leq \varepsilon, |\sqrt{n}Y| \leq \varepsilon] \right| = O(1), \quad \alpha \in (0.25, 0.75). \quad (\text{A.28})$$

For notational convenience, we write $f_n(\rho) = \mathbb{P}_\rho(|\sqrt{n}X| \leq \varepsilon, |\sqrt{n}Y| \leq \varepsilon)$. We study the logarithmic second derivative

$$J_n(\rho) := \frac{f_n''(\rho)}{f_n(\rho)} - \left(\frac{f_n'(\rho)}{f_n(\rho)}\right)^2 \quad (\text{A.29})$$

by controlling each term individually.

First, recall that for any $\rho \in (-1, 1)$, the distribution \mathbb{P}_ρ admits a density with respect to the Lebesgue measure over \mathbb{R}^2 given by

$$\psi_\rho(x, y) = \frac{1}{2\pi\sqrt{1-\rho^2}} \exp\left(-\frac{x^2 - 2\rho xy + y^2}{2 - 2\rho^2}\right).$$

It holds that

$$f_n'(\rho) = \iint_{[-\frac{\varepsilon}{\sqrt{n}}, \frac{\varepsilon}{\sqrt{n}}]^2} \partial_\rho \psi_\rho(x, y) dx dy.$$

Thus since $\varepsilon = o(\sqrt{n})$ we get,

$$\lim_{n \rightarrow \infty} \frac{f_n'(\rho)}{f_n(\rho)} = \lim_{n \rightarrow \infty} \frac{\frac{\varepsilon^2}{n} \iint_{[-\frac{\varepsilon}{\sqrt{n}}, \frac{\varepsilon}{\sqrt{n}}]^2} \partial_\rho \psi_\rho(x, y) dx dy}{\frac{\varepsilon^2}{n} \iint_{[-\frac{\varepsilon}{\sqrt{n}}, \frac{\varepsilon}{\sqrt{n}}]^2} \psi_\rho(x, y) dx dy} = \frac{\partial_\rho \psi_\rho(0, 0)}{\psi_\rho(0, 0)} = \partial_\rho \log(\psi_\rho)(0, 0). \quad (\text{A.30})$$

Similarly,

$$\lim_{n \rightarrow \infty} \frac{f_n''(\rho)}{f_n(\rho)} = \frac{\partial_\rho^2 \psi_\rho(0, 0)}{\psi_\rho(0, 0)} = \partial_\rho^2 \log(\psi_\rho)(0, 0) + (\partial_\rho \log(\psi_\rho)(0, 0))^2. \quad (\text{A.31})$$

Together with (A.29) and (A.30), the above display yields

$$\lim_{n \rightarrow \infty} J_n(\rho) = \frac{1 + \rho^2}{(1 - \rho^2)^2} = O(1),$$

if $\rho \in (-0.5, 0.5)$. Moreover, the convergence in (A.30) and (A.31) is uniform over $\rho \in (-0.5, 0.5)$. This is because the functions ψ_ρ , $\partial_\rho \psi_\rho$, and $\partial_\rho^2 \psi_\rho$ are all C -Lipschitz on \mathbb{R}^2 for some absolute constant $C > 0$, provided that we restrict $\rho \in (-0.5, 0.5)$. Next, changing variables via $\rho = 1 - 2\alpha$, this verifies (A.28). Thus we have shown that $\phi_n(\alpha)$ is strictly concave on $(0.25, 0.75)$ for n sufficiently large, completing the first part of the proof.

The strict concavity verifies that $\phi_n(\alpha)$ has a unique maximum on $(0.25, 0.75)$. We show that it occurs at $\alpha = 0.5$. It is easy to check that both $h(\alpha)$ and $\alpha \mapsto \log \frac{1}{\sqrt{\alpha(1-\alpha)}}$ have a critical point at $\alpha = 1/2$. So, applying the change of variables $\rho = 1 - 2\alpha$, we just need to verify that $f_n'(0) = 0$. Let $\phi(x) = \frac{1}{\sqrt{2\pi}} e^{-x^2/2}$ denote the density of a standard Gaussian and set $\ell = \varepsilon/\sqrt{n}$. Straightforward calculus shows that

$$\left. \frac{\partial}{\partial \rho} \right|_{\rho=0} \psi_\rho(x, y) = xy\phi(x)\phi(y).$$

Therefore,

$$\left. \frac{\partial}{\partial \rho} \right|_{\rho=0} f_n(\rho) = \left(\int_{-\ell}^{\ell} x\phi(x) \right)^2 = 0.$$

This proves the second part of the lemma, so we're done. ■

Appendix B. Gaussian discrepancy in small linear dimension

The goal of this appendix is to prove the result below, which combined with Theorem 1 and Theorem 2 of Chandrasekaran and Vempala (2014) provides a precise characterization of asymptotic Gaussian discrepancy.

Theorem 16 *Let $X_1, \dots, X_n \stackrel{iid}{\sim} \mathcal{N}(0, I_m)$ be independent standard Gaussian random vectors. Let $\gamma > 1$ denote an arbitrary absolute constant. Then there exists $\Delta = \Delta(\gamma)$ such that for $m \leq \Delta n$,*

$$\liminf_{n \rightarrow \infty} \mathbb{P} \left[\mathcal{D}(X_1, \dots, X_n) \leq \gamma \sqrt{\frac{\pi n}{2}} 2^{-n/m} \right] \geq 0.99. \quad (\text{B.32})$$

In particular, combining Theorem 16 with Theorem 2 of Chandrasekaran and Vempala (2014), we can now estimate the discrepancy up to constant factor, with probability asymptotically larger than 99%, in the entire linear regime $m = \delta n$ where $\delta > 0$. Note that our guarantee on the probability here is weaker than that of the high-probability upper bound from Theorem 1. The constant 0.99 can be boosted to be arbitrarily close to 1 by choosing smaller Δ , though our techniques do not allow us to set the right-hand-side to be 1 for any fixed $\Delta > 0$.

The closely related work of Aubin et al. (2019) also considered Gaussian discrepancy in the linear regime $m = \delta n$ for fixed $\delta > 0$. Subject to a certain numerical hypothesis, the authors showed that

$$\liminf_{n \rightarrow \infty} \mathbb{P} \left[\mathcal{D}(X_1, \dots, X_n) \leq c(\delta) \sqrt{n} \right] > 0, \quad (\text{B.33})$$

where $c(\delta)$, as a function of δ , is the inverse of the function $x \mapsto \log(1/2)/\mathbb{P}[|Z| \leq x]$ and $Z \sim N(0, 1)$. Their proof is an application of the second moment method, similar to ours. They also showed the following high-probability lower bound using the first moment method:

$$\lim_{n \rightarrow \infty} \mathbb{P} \left[\mathcal{D}(X_1, \dots, X_n) \geq (c(\delta) - \varepsilon) \sqrt{n} \right] = 1, \quad (\text{B.34})$$

where $\varepsilon > 0$ is an arbitrary absolute constant. Aubin et al. (2019) conjectures, with strong evidence using heuristics from statistical mechanics, that the event in (B.33) holds with probability tending to 1. We remark that as $\delta \rightarrow 0$, we have $c(\delta) = \Theta(2^{-1/\delta}) = \Theta(2^{-n/m})$. Theorem 16 shows that with a constant factor's worth of ‘extra room’ in the discrepancy threshold, the asymptotic probability in (B.33) can be boosted to be arbitrarily close to 1.

On the algorithmic side, using a mild extension of the techniques of Chandrasekaran and Vempala (2014), in dimension $m = \delta n$ with $\delta \in (0, 1)$, one can show an algorithmic bound of $O(\sqrt{\delta n})$ on the discrepancy, and this is the best known result for this regime. Hence, Theorem 16 suggests the possibility of a statistical-to-computational gap in the small linear regime $m = \delta n$ for $\delta \in (0, 1)$. Note that for $\delta > 1$, the results of Chandrasekaran and Vempala (2014) confirm an absence of statistical-to-computational gaps in the discrepancy.

The proof of Theorem 16 follows closely the steps from Section 3 with some modifications. We begin with a truncation argument as in Lemma 14.

Lemma 17 *Let $\gamma > 1$ denote an arbitrary absolute constant. Then there exists $\Delta = \Delta(\gamma)$ such that if $m = \delta n$ for $\delta \leq \Delta$ and $\varepsilon = \varepsilon(n) = \gamma 2^{-1/\delta} \sqrt{\pi n/2}$, then*

$$2^n \sum_{k=0}^{n/4} \binom{n}{k} \mathbb{P}_{\rho_k} (|\sqrt{n}X| \leq \varepsilon, |\sqrt{n}Y| \leq \varepsilon)^m = o(\mathbb{E}[S]^2). \quad (\text{B.35})$$

$$2^n \sum_{k=3n/4}^n \binom{n}{k} \mathbb{P}_{\rho_k} (|\sqrt{n}X| \leq \varepsilon, |\sqrt{n}Y| \leq \varepsilon)^m = o(\mathbb{E}[S]^2). \quad (\text{B.36})$$

Proof

The proof follows closely that of Lemma 14, setting $g_n \equiv 1/\delta$. We set

$$f_\delta(\rho) = \mathbb{P}_\rho(|\sqrt{n}X| \leq \varepsilon, |\sqrt{n}Y| \leq \varepsilon) = \mathbb{P}_\rho(|X| \leq \gamma 2^{-1/\delta} \sqrt{\pi/2}, |Y| \leq \gamma 2^{-1/\delta} \sqrt{\pi/2}).$$

Note that the function f_δ is independent of n by our choice of ε . As in (A.24) from Lemma 14, we let

$$A = \sum_{k=0}^{\delta^2 n} \frac{\binom{n}{k}}{2^n} \left(\frac{f_\delta(\rho_k)}{f_\delta(0)} \right)^m, \quad B = \sum_{k=\delta^2 n}^{n/4} \frac{\binom{n}{k}}{2^n} \left(\frac{f_\delta(\rho_k)}{f_\delta(0)} \right)^m.$$

Note that for δ sufficiently small (depending on γ), it holds that $\varepsilon/\sqrt{n} \leq 1$. Therefore, similar to (A.25), we can apply the lower bound from Lemma 13 to conclude that

$$2^n \mathbb{P}[|Z| < \varepsilon/\sqrt{n}]^m \geq 2^n \left(\sqrt{\frac{2}{\pi n}} \varepsilon \right)^m (1 - c\varepsilon^2/n)^m \geq \left(\frac{\gamma+1}{2} \right)^m, \quad (\text{B.37})$$

Hence, as in (A.26) we have

$$A \lesssim_n \left(\frac{\gamma+1}{2} \right)^{-m} (e\delta^{-2})^{\delta^2 n} = \exp \left(-\delta n \log \left(\frac{1}{2}(1+\gamma) \right) + \delta^2 n + 2\delta^2 n \log(1/\delta) \right). \quad (\text{B.38})$$

Hence, if $\delta \leq \Delta(\gamma)$ for $\Delta(\gamma)$ sufficiently small, then we have that $A = o(1)$.

Similar to (A.27), we have by applying (A.20) and (A.21) that

$$B \lesssim_n (c'(\gamma))^m \sum_{k=\delta^2 n}^{n/4} \frac{\binom{n}{k}}{2^n} \left(\frac{n^2}{k(n-k)} \right)^{m/2}. \quad (\text{B.39})$$

By the Hoeffding bound (letting $c''(\gamma)$ denote another constant depending on γ), we have

$$(\text{B.39}) \lesssim_n (c''(\gamma))^m \delta^{-m} e^{-n/8} = \exp(\delta n \log(c''(\gamma)) + \delta n \log(1/\delta) - n/8) = o(1), \quad (\text{B.40})$$

provided that $\delta \leq \Delta(\gamma)$ for $\Delta(\gamma)$ sufficiently small. Since $A = o(1)$ as well for this setting of parameters, the lemma follows. \blacksquare

Our next lemma is a version of Lemma 15 corresponding to the linear regime. We use the log-concavity of the function ϕ_n when we apply the Laplace method to the second moment, as in the sub-linear regime.

Lemma 18 *Let $\eta > 0$ and $\gamma > 1$ be arbitrary constants, and let $\Delta = \Delta(\gamma, \eta)$ denote a sufficiently small absolute constant. Suppose that $m = \delta n$ for $\delta \leq \Delta$, and set $\varepsilon = \gamma 2^{-1/\delta} \sqrt{n\pi/2}$. Then the function $\alpha \mapsto \phi_n(\alpha)$ defined in (3.10) is strictly concave on $(0.25, 0.75)$. More precisely,*

$$\frac{1}{n} \frac{\partial^2}{\partial \alpha^2} \phi_n(\alpha) \leq -\frac{1}{\alpha(1-\alpha)} + \eta < -4 + \eta, \quad \forall \alpha \in (0.25, 0.75). \quad (\text{B.41})$$

Moreover, $\phi_n(\alpha)$ has a unique maximum over $(0.25, 0.75)$ located at $\alpha = 0.5$.

Proof Recall that

$$f_\delta(\rho) = \mathbb{P}_\rho(|X| \leq \gamma 2^{-1/\delta} \sqrt{\pi/2}, |Y| \leq \gamma 2^{-1/\delta} \sqrt{\pi/2}).$$

As in the proof of Lemma 15, it suffices to study the logarithmic second derivative with respect to ρ

$$J_\delta(\rho) := \frac{f_\delta''(\rho)}{f_\delta(\rho)} - \left(\frac{f_\delta'(\rho)}{f_\delta(\rho)} \right)^2 \quad (\text{B.42})$$

and show that $|J_\delta(\rho)| = O(1)$ for $\rho \in (-0.5, 0.5)$. Recall that ψ_ρ denotes the density associated to \mathbb{P}_ρ .

Since $\varepsilon/\sqrt{n} \rightarrow 0$ as $\delta \rightarrow 0$, we have, similar to (A.30), that

$$\lim_{\delta \rightarrow 0} \frac{f_\delta'(\rho)}{f_\delta(\rho)} = \lim_{\delta \rightarrow 0} \frac{\frac{\varepsilon^2}{n} \iint_{[-\frac{\varepsilon}{\sqrt{n}}, \frac{\varepsilon}{\sqrt{n}}]^2} \partial_\rho \psi_\rho(x, y) dx dy}{\frac{\varepsilon^2}{n} \iint_{[-\frac{\varepsilon}{\sqrt{n}}, \frac{\varepsilon}{\sqrt{n}}]^2} \psi_\rho(x, y) dx dy} = \frac{\partial_\rho \psi_\rho(0, 0)}{\psi_\rho(0, 0)} = \partial_\rho \log(\psi_\rho)(0, 0). \quad (\text{B.43})$$

And similar to (A.31), we have

$$\lim_{\delta \rightarrow 0} \frac{f_\delta''(\rho)}{f_\delta(\rho)} = \frac{\partial_\rho^2 \psi_\rho(0, 0)}{\psi_\rho(0, 0)} = \partial_\rho^2 \log(\psi_\rho)(0, 0) + (\partial_\rho \log(\psi_\rho)(0, 0))^2. \quad (\text{B.44})$$

It follows that

$$\lim_{\delta \rightarrow 0} J_\delta(\rho) = \frac{1 + \rho^2}{(1 - \rho^2)^2} = O(1)$$

for $\rho \in (-0.5, 0.5)$. Moreover, similar to the proof of Lemma 15, the convergence in (B.43) and (B.44) is uniform in δ by the Lipschitzness of ψ_ρ , $\partial_\rho \psi_\rho$, and $\partial_\rho^2 \psi_\rho$ over the interval $\rho \in (-0.5, 0.5)$. Therefore, if we take δ sufficiently small with respect to γ, η , then (B.41) holds.

Note that independent of ε , we have that $\rho = 0$ is a critical point of ϕ_n , as shown at the end of the proof of Lemma 15. Applying this and making the change of variables $\rho = 1 - 2\alpha$ verifies the last statement of Lemma 18. \blacksquare

Proof [Proof of Theorem 16] Recall from the definition in (3.8) that

$$L := 2^n \sum_{k=n/4}^{3n/4} \binom{n}{k} \mathbb{P}_{\rho_k}(|\sqrt{n}X| \leq \varepsilon, |\sqrt{n}Y| \leq \varepsilon)^m.$$

Applying Stirling's formula and a Riemann sum approximation as in (3.9) and (3.11), respectively, we have that

$$L \lesssim_n 2^n \sqrt{\frac{n}{2\pi}} \int_{1/4}^{3/4} \exp(\phi_n(\alpha)) d\alpha. \quad (\text{B.45})$$

Since $\phi_n(\alpha)/n$ is independent of n , we can apply the Laplace method directly (see Murray, 1984) along with Lemma 18 to see that

$$\int_{1/4}^{3/4} \exp(\phi_n(\alpha)) d\alpha \lesssim_n \sqrt{\frac{2\pi}{|\phi_n''(1/2)|}} \exp(\phi_n(1/2)) \leq \sqrt{\frac{2\pi}{n(4-\eta)}} 2^{n+1} f_\delta(0)^m. \quad (\text{B.46})$$

assuming $\delta \leq \Delta$ for $\Delta(\gamma, \eta)$ sufficiently small.

Therefore, by Lemma 14, (B.45), (B.46), Lemma 5, the definition of f_δ , and assuming that $\delta \leq \Delta$ for $\Delta(\gamma, \eta)$ sufficiently small, we have

$$\mathbb{E}[S^2] \lesssim_n L \lesssim_n \sqrt{\frac{4}{4-\eta}} (2^n \mathbb{P}[|\sqrt{n}Z| \leq \varepsilon]^m)^2 = \sqrt{\frac{4}{4-\eta}} \mathbb{E}[S]^2.$$

Setting $\eta = 10^{-5}$, we have by the second moment method (2.4) that

$$\mathbb{P}[S > 0] \geq \frac{\mathbb{E}[S]^2}{\mathbb{E}[S^2]} \gtrsim_n \sqrt{1 - \eta/4} \geq 0.99,$$

completing the proof of Theorem 16. ■

Appendix C. The REDUCE algorithm

In this appendix we define the **REDUCE** algorithm, a simple procedure for combining a set of points into a single point whose ℓ_∞ -norm is not too large. This algorithm **REDUCE** is described explicitly below, and its main property of use is described in Lemma 8, whose proof is given below. The analysis of this algorithm uses feasibility as in the classical proof of the Beck-Fiala theorem (Alon and Spencer, 2008).

REDUCE:

Input: $m \times N$ matrix \mathbf{X} with columns X_1, \dots, X_N .

If $N < m$:

 Choose $s \in \{\pm 1\}^N$ arbitrarily.

Else:

1. Let $s^{(0)} = \mathbf{0} \in \mathbb{R}^N$, and let $T_0 = \emptyset$.

2. For $k = 0, 1, 2, \dots$

 If $|T_k| < N - m$

 (a) Find (e.g., using Gaussian elimination) a vector $v \neq \mathbf{0} \in \mathbb{R}^N$ such that $\mathbf{X}v = \mathbf{0}$ and $v_j = 0$ for all $j \in T_k$.

 (b) Define $s^{(k+1)} = s^{(k)} + \lambda v$, where $\lambda > 0$ is the smallest real number such that $|s_j^{(k)} + \lambda v_j| = 1$ for some $j \notin T_k$.

 (c) Define $T_{k+1} = \{j : |s^{(k+1)}| = 1\}$.

 Else: $s := s^{(k)}$. BREAK

Output: $\sigma := \text{sgn}(s)$

Proof [Proof of Lemma 8] We suppose that $N > m$, otherwise, an arbitrary choice of signing gives the desired upper bound. Suppose that we are in the k -th iteration of Step 2 of **REDUCE**. If $|T_k| < N - m$, then there are at most $m + |T_k| < N$ linear constraints on the vector $v \in \mathbb{R}^N$ in step 2(a). So by dimension-counting, there exists a nonempty subspace of feasible v . Next if $s^{(k)} \in [-1, 1]^m$, then λ from step 2(b) exists and furthermore $s^{(k+1)} \in [-1, 1]^m$ by the choice of j in step 2(b). Also, we have that $T_k \subset T_{k+1}$; if $|(s^{(k)})_j| = 1$, then the j -th coordinate remains unchanged for future iterations of step 2. Finally, $|T_k|$ increases at least by 1 in each iteration, so the loop in step 2 is guaranteed to terminate after at most $N - m$ iterations.

It remains to verify that σ satisfies the upper bound from Lemma 8. Observe that $s \in [-1, 1]^m$, $T := \{j : |s_j| = 1\} \geq N - m$, and

$$\sum_{i=1}^N s_i X_i = \mathbf{0}.$$

Therefore,

$$\begin{aligned} \left| \sum_{i=1}^N \sigma_i X_i \right|_{\infty} &\leq \left| \sum_{i=1}^N s_i X_i \right|_{\infty} + \left| \sum_{i \notin T} (\text{sgn}(s_i) - s_i) X_i \right|_{\infty} \\ &\leq \max_{S \subset [N]: |S|=m} \sum_{i \in S} |X_i|_{\infty}. \end{aligned}$$

■

Appendix D. Proof of Proposition 11

We need to show that at each application of resampling in **GKK**, a small number of points are labeled ‘bad’. As discussed in the introduction, the restriction on the dimension $m = O(\sqrt{\log n})$ is needed in our analysis to show that the probability of a point being labeled ‘bad’ is small.

We briefly describe the intuition for this condition by considering the first phase of the algorithm **GKK**. Suppose, for example, that X_1, \dots, X_n are independent triangularly distributed vectors on $[-1, 1]^m$. In step 1 of **PRDC**, the cube $[-1, 1]^m$ is partitioned into sub-cubes of side length $\alpha' = n^{-\Omega(1/m)}$. Next, we enter the resampling step. We show below that the probability of a point being labeled ‘bad’ is at most $O(2^m m \alpha') = O(2^m m n^{-\Omega(1/m)})$. Roughly speaking, the reason for this is that there are $2^m (\alpha')^{-m}$ sub-cubes, and the probability of a point in a particular sub-cube being labeled ‘bad’ is controlled by the product of three terms: 1) the ℓ_1 -Lipschitz constant of the density of X_1 , which is 1, 2) the ℓ_1 -diameter of the sub-cube, which is $m\alpha'$, and 3) the volume of the sub-cube, which is $(\alpha')^m$. Hence, the probability of a point being labeled ‘bad’ is a small constant, assuming that $m = O(\sqrt{\log n})$.

The next two lemmas present the above argument in full detail.

Lemma 19 *Let $\rho : [-\Delta, \Delta] \rightarrow \mathbb{R}$ denote a pdf that is L -Lipschitz and bounded above by some constant $D > 0$. Let $g = \rho^{\otimes m} : [-\Delta, \Delta]^m \rightarrow \mathbb{R}$ denote the density of the distribution of m independent random variables, each individually distributed according to ρ . Then g is L' -Lipschitz in the ℓ_1 norm:*

$$\forall x, y \in [-\Delta, \Delta]^m, \quad |g(x) - g(y)| \leq L' |x - y|_1,$$

where

$$L' = LD^{m-1}.$$

Proof

Define $x^1 = x$, and for $2 \leq k \leq m$, define

$$x^k = x^{k-1} + \mathbf{e}_k (y_k - x_k),$$

where \mathbf{e}_k denotes the k -th elementary basis vector. Then we have

$$\begin{aligned} |g(x) - g(y)| &\leq \sum_{k=1}^m |g(x^k) - g(x^{k-1})| \left(\prod_{i < k} g(y_i) \right) \left(\prod_{i > k} g(x_i) \right) \\ &\leq \sum_{k=1}^m LD^{m-1} |x_k - y_k| \\ &= LD^{m-1} |x - y|_1. \end{aligned}$$

■

Lemma 20 *Let $S = X_1, \dots, X_s \in [-\Delta, \Delta]^m$ denote a sample of iid random vectors, each having a joint density $g = \rho^{\otimes m}$, where ρ is L -Lipschitz and bounded above by $D > 0$. Let B denote the bad points created in step 2 of **PRDC** run on the input $S, v = 0, \alpha = \Delta$, and g . If $m \leq C\sqrt{\log(s)/\max(1, \log \Delta)}$ for a sufficiently small constant $C = C(D, L) > 0$, then*

$$\mathbb{P}[|B| > 0.1s] \leq \exp(-c_1 s),$$

where c_1 is an absolute constant.

Proof Let $\alpha' = \Delta/\lceil s^{1/(4m)} \rceil$. Let C_1, \dots, C_N denote the sub-cubes of side length α' formed by partitioning (step 1 of **PRDC**), recalling that $N = (2\Delta)^m(\alpha')^{-m}$. Since X_1, \dots, X_s are independent, we first study the probability that X_1 is bad and then apply a Hoeffding bound.

$$\begin{aligned} \mathbb{P}[X_1 \text{ is bad}] &= \sum_{j=1}^N \int_{C_j} \left(1 - \frac{\min_{y \in C_j} g(y)}{g(x)}\right) g(x) dx \\ &= \sum_{j=1}^N \int_{C_j} \left(g(x) - \min_{y \in C_j} g(y)\right) dx \\ &\leq \sum_{j=1}^N \text{Vol}(C_j) L D^{m-1} \text{diam}_{\ell_1}(C_j) \\ &= (2\Delta)^m L D^{m-1} m \alpha', \end{aligned}$$

where we measure the diameter in the ℓ_1 norm and applied Lemma 19. Since

$$m \leq C\sqrt{\log(s)/\max(1, \log \Delta)},$$

we have

$$p := (2\Delta)^m L D^{m-1} m \alpha' \leq (2\Delta)^m D^{m-1} m \Delta s^{-1/(4m)} \leq 0.05$$

for $C = C(D, L) > 0$ sufficiently small. Since the X_i 's are independent, by Hoeffding's inequality,

$$\mathbb{P}[|B| \geq 0.1s] \leq \mathbb{P}[|B| - ps \geq 0.05s] \leq \exp\left(-\frac{2(0.05)^2 s^2}{s}\right),$$

which completes the proof. ■

Proof [Proof of Proposition 11] The proof is by induction on t . We first handle the base case $t = 1$. By assumption the matrix \mathbf{X} has independent entries, each having a pdf which is L -Lipschitz and bounded above by D . By Lemma 20, with probability at least $1 - \exp(-c_1 n)$, there are at most $0.1n$ points labeled 'bad'. Since $m \leq C\sqrt{\log(n)/\max(1, \log \Delta)}$, for C sufficiently small, there are at most $N_1 \leq (2\Delta)^m \alpha_2^{-m} \leq n^{0.6}$ sub-cubes created by partitioning (step 1 of **PRDC**). Thus, at most that many good points are leftover after random differencing in step 3 of **PRDC**. We conclude that with probability at least $1 - \exp(-c_1 n)$, there are at least

$$\frac{n - 0.01n - n^{0.6}}{2} \geq 0.4n \tag{D.47}$$

points in G'_1 , the set of random differences.

Now we show the inductive step. Let \mathcal{E} denote the event $|S_j| = n_j$ where $n_j \geq (0.3)^{j-1}n$ for all $1 \leq j \leq t$. It suffices to show that

$$\mathbb{P} \left[|G'_{t+1}| \leq 0.4n_t \mid \mathcal{E} \right] \leq \exp(-c_1\sqrt{n}). \quad (\text{D.48})$$

By Proposition 10 in Appendix G, conditionally on \mathcal{E} , the distribution of the points in $S_t = \mathbf{y}_1, \dots, \mathbf{y}_{n_t}$ are iid and follow a triangular distribution on $[-\alpha_t, \alpha_t]^m$. Hence, we have by Lemma 19 that the density of $\alpha_t^{-1}\mathbf{y}_1, \dots, \alpha_t^{-1}\mathbf{y}_{n_t}$ is 1-Lipschitz with respect to ℓ_1 and is bounded above by $D = 1$. Note that, by an application of the chain rule, the probability $\alpha_t^{-1}\mathbf{y}_j$ is labeled ‘good’ using the triangular density on $[-1, 1]^m$ for g in step 2 of **PRDC** is the same as the probability that \mathbf{y}_j is labeled ‘good’ using the triangular density on $[-\alpha_t, \alpha_t]^m$ for g in step 2 of **PRDC**.

Since $t \leq \lceil C^* \log n \rceil$ and $n_j \geq (0.3)^{j-1}n$ for $1 \leq j \leq t$, we have that $n_t \geq \sqrt{n}$. In particular, for $C > 0$ sufficiently small, $s = \sqrt{n}$ satisfies the required lower bound of Lemma 20. Therefore,

$$\mathbb{P} \left[|B_{t+1}| \geq 0.1n_t \mid \mathcal{E} \right] \leq \exp(-c_1n_t) \leq \exp(-c_1\sqrt{n}).$$

For C sufficiently small and $m \leq C\sqrt{\log n}$, there are at most $N_t \leq 2^m n_t^{1/4} \leq n_t^{0.6}$ sub-cubes formed in step 1 of **PRDC**. Hence, at most $n_t^{0.6}$ good points are leftover after the random differencing step of **PRDC**. Halving the number of remaining points as in (D.47) of the base case, we conclude that (D.48) holds with the desired probability in phase t . \blacksquare

Appendix E. Proof of Proposition 12

The goal of this subsection is to prove Proposition 12. The next technical lemma implies that a negligible fraction of points are lost in step 4(b), the clean-up step of **PRDC**.

Lemma 21 *Let $\alpha = \lceil s^{1/(4m)} \rceil^{-1}$, and let $\mathcal{U} = \mathbf{u}_1, \dots, \mathbf{u}_s \stackrel{iid}{\sim} \text{Tri}[-\alpha, \alpha]^m$ denote a sample from a triangular distribution. Let $v^{(0)} \in \mathbb{R}^m$ denote a random vector independent of \mathcal{U} satisfying $|v^{(0)}|_2 \leq Rm^{3/2}$ for some absolute constant $R > 0$. For $k = 1, 2, \dots$, define a sequence of random vectors*

$$v^{(k)} = v^{(k-1)} + a^* \mathbf{u}_k$$

where

$$a^* = \operatorname{argmin}_{a \in \{\pm 1\}} \left| v^{(k-1)} + a \mathbf{u}_k \right|_2.$$

Let c^* denote the absolute constant from Claim E.1. Suppose that $R' \geq 2/c^*$ and

$$K \geq \frac{8R^2 m^2 \sqrt{s}}{R' c^*}.$$

Then with probability at least

$$1 - \exp\left(-\frac{(c^*)^2 K}{8m}\right)$$

there exists $k \leq K$ such that

$$|v^{(k)}|_2 \leq R' m \alpha.$$

Proof By the definition of $v^{(k)}$, we have that

$$0 \leq \left| v^{(K+1)} \right|_2^2 = \left| v^{(0)} \right|_2^2 + \sum_{k=0}^K \left(-2 \left| \langle v^{(k)}, \mathbf{u}_{k+1} \rangle \right| + \left| \mathbf{u}_{k+1} \right|_2^2 \right).$$

Consider the event \mathcal{E} that for all $1 \leq k \leq K$, we have $\left| v^{(k)} \right|_2 \geq R' m \alpha$. Let $\nu^{(k)} = v^{(k)} / \left| v^{(k)} \right|_2$. Observe that $\left| \mathbf{u}_k \right|_2^2 \leq \alpha^2 m$. Applying this and rearranging the inequality above, we have that the event \mathcal{E} implies

$$\sum_{k=0}^K \left| \langle \nu^{(k)}, \mathbf{u}_{k+1} \rangle \right| \leq \frac{R^2 m^3 + \alpha^2 m K}{2R' m \alpha}. \quad (\text{E.49})$$

For $0 \leq j \leq K$, define a sequence of random variables

$$M_j := \sum_{k=0}^j \left(\left| \langle \nu^{(k)}, \mathbf{u}_{k+1} \rangle \right| - c^* \alpha \right).$$

For convenience, we also define $M_{-1} \equiv 0$. Note that M_j is measurable with respect to the sigma-field Ω_j generated by the random variables $v^{(0)}, v^{(1)}, \dots, v^{(j+1)}$. Therefore, $\Omega_{-1} \subset \Omega_0 \subset \dots$ defines a filtration for the sequence of random variables $\{M_j\}_{j \geq -1}$.

Claim E.1 *There exists an absolute constant $c^* > 0$ such that $\{M_j\}_{j \geq -1}$ is a submartingale with respect to the filtration $\{\Omega_j\}_{j \geq -1}$.*

Proof Since $v^{(0)}$ is independent of \mathcal{U} and \mathcal{U} is an independent sample, it follows that \mathbf{u}_{k+1} is independent of $\nu^{(k)}$. Observe that the coordinates of \mathbf{u}_{k+1} are subGaussian. By the Khintchine inequality for the ℓ_1 norm (see Exercises 2.6.5 and 2.6.6 of [Vershynin, 2018](#)), we have

$$\mathbb{E} \left[\left| \langle \nu^{(k)}, \mathbf{u}_{k+1} \rangle \right| \mid v^{(k)} \right] = \mathbb{E} \left[\left| \langle \nu^{(k)}, \mathbf{u}_{k+1} \rangle \right| \mid \nu^{(k)} \right] \geq \alpha c^* \left| \nu^{(k)} \right|_2 = \alpha c^* > 0$$

for an absolute constant $c^* > 0$. ■

Let $c^* > 0$ denote the absolute constant from Claim E.1, and set $R' \geq 2/c^*$. Next, note the equivalence between the following inequalities:

$$\begin{aligned} c^* \alpha K &\geq \frac{c^* \alpha K}{2} + \frac{R^2 m^3 + \alpha^2 m K}{2R' m \alpha} \Leftrightarrow \\ K &\geq \frac{R^2 m^2}{R'(c^* - 1/R')} \alpha^{-2}, \end{aligned} \quad (\text{E.50})$$

assuming that $c^* - 1/R' > 0$. Setting $R' \geq 2/c^*$, it follows that if

$$K \geq \frac{8R^2 m^2 \sqrt{s}}{R' c^*},$$

then (E.50) holds. Next, note by Cauchy-Schwarz that the submartingale M_j has increments bounded by $\alpha \sqrt{m}$. Since (E.50) holds, we may apply the Hoeffding–Azuma inequality to conclude that for such choice of K and R' that

$$\mathbb{P}[\mathcal{E}] \leq \mathbb{P} \left[M_K \leq \frac{R^2 + \alpha^2 m^2 K}{2R' m \alpha} - c^* \alpha K \right] \leq \mathbb{P} \left[M_K \leq -\frac{c^* \alpha K}{2} \right] \leq \exp \left(-\frac{(c^*)^2 K}{8m} \right),$$

as desired. ■

Proof [Proof of Proposition 12]

Let $t \geq 1$ denote the current phase. Let \mathcal{E} denote the event that $|S_j| = n_j$ for all $1 \leq j \leq t$ and $|G'_t| = g'_t$ where $n_j \geq (0.3)^{j-1}n$ for all $1 \leq j \leq t$ and $g'_t \geq (0.4)n_t$. By Proposition 10 and Lemma 27 in Appendix G, conditionally on \mathcal{E} , the points $\mathbf{z}_1, \dots, \mathbf{z}_{g'_t} \in G'_t$ are distributed as $\text{Tri}[-\alpha_{t+1}, \alpha_{t+1}]^m$, and the leftover vector $v_t^{(0)}$ obtained in step 4(a) of **PRDC** is independent of this sample. Moreover, by Lemma 8 and the fact that $|v_t|_\infty \leq |v_t|_2 \leq \gamma m \alpha_t$, it follows that

$$\left| v_t^{(0)} \right|_\infty \leq (\gamma + 1)m\alpha_t.$$

Hence, the Cauchy–Schwarz inequality yields that

$$\left| v_t^{(0)} \right|_2 \leq (\gamma + 1)m^{3/2}\alpha_t.$$

Next, apply Lemma 21 with $\mathcal{U} = \frac{1}{\alpha_t}\mathbf{z}_1, \dots, \frac{1}{\alpha_t}\mathbf{z}_{g'_t}$, $v^{(0)} = \frac{1}{\alpha_t}v_t^{(0)}$, $R = \gamma + 1$, $R' = \gamma$, and $K = (g'_t)^{3/4}$ where $\gamma \geq 2/c^*$. Recall that by assumption $g'_t \geq (0.4)n_t \geq (0.4)(0.3)^{t-1}n$. Since $t \leq \lceil C^* \log n \rceil$, we have that $g'_t \geq \sqrt{n}$. So for C sufficiently small in the bound $m \leq C\sqrt{\log n}$, we have that the lower bound

$$K = (g'_t)^{3/4} \geq \frac{8(\gamma + 1)^2 m^2 \sqrt{g'_t}}{\gamma c^*}$$

holds, and so indeed Lemma 21 applies. Therefore, conditioned on \mathcal{E} , with probability at least

$$1 - \exp\left(-\frac{(c^*)^2 (g'_t)^{3/4}}{8m}\right) \geq 1 - \exp\left(-(c^*)^2 n^{1/4}\right)$$

there exists $k \leq K = (g'_t)^{3/4}$ with

$$\left| v_t^{(k)} \right|_2 \leq \gamma m \alpha_{t+1}.$$

By the lower bounds $n \geq e^{(1/C)m^2}$ and $g'_t \geq \sqrt{n}$, for C sufficiently small, it follows that $(g'_t)^{3/4} \leq (0.01)g'_t$. Hence, conditioned on \mathcal{E} , with probability at least $1 - \exp(-c^* n^{1/4})$ we have $|S_{t+1}| \geq g'_t - (g'_t)^{3/4} \geq (0.3)n_t$, as desired. ■

Appendix F. Proof of Theorem 3

Our main theorem is a direct consequence of Propositions 11 and 12.

Proof [Proof of Theorem 3] Recall that $T = \lceil C^* \log n \rceil$ where $C^* = (2 \log(10/3))^{-1}$, and set $\theta = 0.3$. By the union bound over the T phases of **PRDC** in **GKK**, induction, and Propositions 11 and 12, we have that $|S_t| \geq \theta^{t-1}n$ for all $1 \leq t \leq T$ with probability at least $1 - \exp(-c_3 n^{1/4})$, for some absolute constant $c_3 > 0$. Since $\alpha_{t+1} = \alpha_t / \lceil |S_t|^{1/(4m)} \rceil$, this implies by induction that

$$\alpha_T \leq \max(1, \Delta) \theta^{-T^2/(4m)} n^{-T/(4m)} \leq \max(1, \Delta) \exp\left(-\frac{C^* \log^2 n}{8m}\right)$$

with probability at least $1 - \exp(-c_3 n^{1/4})$.

Moreover, by the stopping criterion from step 4(b) of **PRDC**, $|v_T|_\infty \leq |v_T|_2 \leq \gamma m \alpha_T$. Applying **REDUCE** to $S_T \cup \{v_T\}$, we see by Lemma 8 that the output $|v|_\infty$ of **GKK** satisfies

$$|v|_\infty \leq \max(1, \Delta)(\gamma m + m - 1) \exp\left(-\frac{C^* \log^2 n}{8m}\right) \leq \exp\left(-\frac{c \log^2 n}{m}\right)$$

for an absolute constant $c > 0$. Note that the right-hand-side follows if we take $C > 0$ sufficiently small in the bound $m \leq C \sqrt{\log(n) / \max(1, \log \Delta)}$. \blacksquare

Appendix G. Distributional properties

Our analysis of **GKK** relies heavily on the fact that the operations in the algorithm preserve important features of the original distribution such as independence. Though not carefully proven in [Karmarkar and Karp \(1982\)](#), these features are crucial to our analysis, so we provide explicit justification of these properties below for completeness.

First we introduce some notation. Given $\alpha > 0$, a fixed collection of vectors $\mathbf{z}_1, \dots, \mathbf{z}_s \in [-\alpha, \alpha]^m$, and a density $g : [-\alpha, \alpha]^m$, divide the cube $[-\alpha, \alpha]^m$ into $N := 2^m (\lceil s^{1/(4m)} \rceil)^m$ sub-cubes C_1, \dots, C_N of side length $\alpha / \lceil s^{1/(4m)} \rceil$ as in step 1 of **PRDC**. Label the points $\mathbf{z}_1, \dots, \mathbf{z}_s$ as in step (2) of **PRDC** using the density g . Define a random collection of ordered pairs $\mathcal{T}_{s,\alpha,g} \subset ([N] \times \{0, 1\})^s$ so that for $1 \leq i \leq s$,

$$(\mathcal{T}_{s,\alpha,g})_i = (j, 1)$$

if and only if $\mathbf{z}_i \in C_j$ and if \mathbf{z}_i is labeled ‘good’, and

$$(\mathcal{T}_{s,\alpha,g})_i = (j, 0)$$

if and only if $\mathbf{z}_i \in C_j$ and \mathbf{z}_i is labeled as ‘bad’.

Usually s, α and g are clear from context, in which case we write \mathcal{T} for $\mathcal{T}_{s,\alpha,g}$. Observe that \mathcal{T} keeps track of which sub-cube v_i lands in and also whether it was labeled good or bad. We refer to \mathcal{T} as the *configuration vector* corresponding to the input of **PRDC**.

We proceed by proving some preliminary lemmas, the first of which states roughly that given random vectors $\mathbf{z}_1, \dots, \mathbf{z}_s$ with a nice conditional distribution, the good points in each sub-cube C_j have a uniform distribution.

Lemma 22 *Suppose that conditioned on an event \mathcal{F} ,*

- *the random vectors $S = \mathbf{z}_1, \dots, \mathbf{z}_s \in \mathbb{R}^m$ are iid, and each vector has the conditional joint density $g : [-\Delta, \Delta]^m \rightarrow \mathbb{R}$.*
- *$S \cup \{v\}$ is a collection of independent random vectors.*

*Run the first two steps of **PRDC** with input $S = \mathbf{z}_1, \dots, \mathbf{z}_s, v$, $\alpha = \Delta$, and density g . Let G denote the good points, and let B denote the bad points. Then conditioned on $\mathcal{T}_{s,\Delta,g}$ and \mathcal{F} ,*

- *the random vectors in $B \cup G$ are mutually independent.*
- *For $1 \leq j \leq N$, a given good point in C_j has a uniform distribution on C_j .*

Proof The first statement follows because (1) $G \cup B = \mathbf{z}_1, \dots, \mathbf{z}_s$ is an independent sample, conditioned on \mathcal{F} , and (2) the ordered pair $(\mathcal{T}_{s,\Delta,g})_i$ is generated independently for each $i \in [s]$. Thus it suffices to show, by symmetry and passing to conditional densities, that

$$g(z|\mathbf{z}_1 \in C_j, \mathbf{z}_1 \text{ good}) = \frac{1}{\text{Vol}(C_j)}$$

for all $z \in C_j$. By Bayes' rule,

$$\begin{aligned} g(z|\mathbf{z}_1 \in C_j, \mathbf{z}_1 \text{ good}) &= \frac{\mathbb{P}[\mathbf{z}_1 \text{ good}|\mathbf{z}_1 = z, \mathbf{z}_1 \in C_j, \mathcal{F}] g(z|\mathbf{z}_1 \in C_j)}{\mathbb{P}[\mathbf{z}_1 \text{ good}|\mathbf{z}_1 \in C_j, \mathcal{F}]} \\ &= \left(\frac{\min_{x \in C_j} g(x)}{g(z)} \cdot \frac{g(z)}{\mathbb{P}[\mathbf{z}_1 \in C_j | \mathcal{F}]} \right) / \left(\frac{\text{Vol}(C_j) \min_{x \in C_j} g(x)}{\mathbb{P}[\mathbf{z}_1 \in C_j | \mathcal{F}]} \right) \\ &= \frac{1}{\text{Vol}(C_j)}, \end{aligned}$$

where the last line follows because

$$\mathbb{P}[\mathbf{z}_1 \text{ good}, \mathbf{z}_1 \in C_j | \mathcal{F}] = \int_{C_j} \mathbb{P}[\mathbf{z}_1 \text{ good}|\mathbf{z}_1 = z, \mathcal{F}] g(z) dz = \text{Vol}(C_j) \min_{x \in C_j} g(x). \quad \blacksquare$$

Lemma 23 Consider the set-up of Lemma 22, and let $\alpha' = \alpha / \lceil s^{1/(4m)} \rceil$. Let G' denote the set of random differences constructed after step 3. of **PRDC** applied to S , v , $\alpha = \Delta$, and g . Then conditioned on the events \mathcal{F} and $\mathcal{T} = \mathbf{T}$, the points in G' are iid and have a triangular distribution on $[-\alpha', \alpha']^m$.

Proof Observe that \mathbf{T} determines the number of points in G' . The points in G' are independent by Lemma 22 and the fact that the points in G are randomly differenced in step 3. of **PRDC**. Since C_j is a translation of the sub-cube $[-\alpha', \alpha']^m$, the difference of two independent, uniformly sampled points from C_j have a triangular distribution on $[-\alpha', \alpha']^m$. \blacksquare

Lemma 24 Consider the set-up of Lemma 23, and let $\ell \in \mathbb{Z}_{\geq 0}$. Let the random variable \mathcal{L} denote the number of points removed from G' in step 4(b) of **PRDC** applied to S , v , $\alpha = \Delta$, and g . Let S' and v' denote the vectors output by **PRDC**. Let $g' = |G'|$. Then conditioned on the events \mathcal{F} , $\mathcal{T} = \mathbf{T}$, and $\mathcal{L} = \ell$,

- The $g' - \ell$ points in S' are iid and follow a triangular distribution on $[-\alpha', \alpha']^m$.
- The random vector v' is independent of the vectors in S' .

Proof Recall that $|G'| = g'$ is determined by \mathbf{T} . Label the points in G' independently at random to be $G' = \mathbf{y}_1, \dots, \mathbf{y}_{g'}$. The points in G' are independent and triangularly distributed on $[-\alpha', \alpha']^m$ by Lemma 23, conditionally on \mathcal{F} and $\mathcal{T} = \mathbf{T}$. Recall the single vector v that was input initially to **PRDC**. In step 4(a), this is combined with vectors in B' to construct a single vector $v^{(0)}$. By Lemma 22, we have that $v^{(0)}$ is independent of G' , conditionally on $\mathcal{T} = \mathbf{T}$ and \mathcal{F} .

Now in step 4(b) of **PRDC**, let us remove points from G' in the order $\mathbf{y}_{g'}, \mathbf{y}_{g'-1}, \dots, \mathbf{y}_{g'-\ell+1}$. By the stopping criterion for step 4(b), we have

$$\{\mathcal{L} = \ell\} = \left\{ \left| v^{(k)} \right|_2 > \gamma m \alpha' \ \forall 1 \leq k \leq \ell - 1, \left| v^{(\ell)} \right|_2 < \gamma m \alpha' \right\}.$$

Since $v^{(k)} = v^{(k-1)} \pm \mathbf{y}_{g'-k+1}$ for $1 \leq k \leq \ell$, the random vector $v^{(k)}$ is independent of $\mathbf{y}_1, \dots, \mathbf{y}_{g'-\ell}$. Therefore, the sample $S' = \mathbf{y}_1, \dots, \mathbf{y}_{g'-\ell}$ is independent of the event $\mathcal{L} = \ell$. Hence, further conditioning on $\mathcal{L} = \ell$ does not affect the distribution of S' , as desired. \blacksquare

Summarizing the content of Lemmas 22, 23, and 24, we have the following proposition.

Proposition 25 *Suppose that conditioned on an event \mathcal{F} ,*

- *the random vectors $S = \mathbf{z}_1, \dots, \mathbf{z}_s \in \mathbb{R}^m$ are iid, and each vector has the conditional joint density $g : [-\Delta, \Delta]^m \rightarrow \mathbb{R}$.*
- *$S \cup \{v\}$ is a collection of independent random vectors.*

Let S', v' denote the vectors output by PRDC applied to $S, v, \alpha = \Delta$, and g . Let $s' \in \mathbb{Z}_{\geq 0}$ and $\alpha' = \alpha / \lceil s^{1/(4m)} \rceil$. Then conditioned on $\mathcal{F}, \mathcal{T} = \mathbf{T}$, and $|S'| = s'$,

- *the s' points in S' are iid and follow a triangular distribution on $[-\alpha', \alpha']^m$.*
- *The random vector v' is independent of the vectors in S' .*

Observe that Proposition 25 and induction imply the next lemma, which guarantees that we have a nice distribution after every phase of PRDC, conditionally on the data $\mathcal{T}^{(j)}$ at each step.

Lemma 26 *Let X_1, \dots, X_n be iid random vectors, each having a joint density $g : [-\Delta, \Delta]^m \rightarrow \mathbb{R}$, conditioned on some event \mathcal{F} . Consider the output S_t, v_t, α_t that results after the $(t-1)$ -th phase of PRDC in step 2 of GKK. For $1 \leq j \leq t-1$, let $\mathcal{T}^{(j)}$ denote the configuration vector resulting from step 2 of the j -th phase of PRDC. Then conditioned on $\mathcal{T}^{(j)} = \mathbf{T}^{(j)}$ for $1 \leq j \leq t-1$ and $|S_j| = n_j$ for $1 \leq j \leq t$, we have*

- *the n_t points in S_t are iid and follow a triangular distribution on $[-\alpha_t, \alpha_t]^m$.*
- *The random vector v_t is independent of the vectors in S_t .*

Next, marginalizing over all possible configuration vectors yields Proposition 10.

Proof [Proof of Proposition 10] We induct on the phase t . Consider the base case $t = 2$. Let $\mathbf{z}_1, \dots, \mathbf{z}_{n_2}$ denote the vectors in S_2 , and let I_i denote a measurable subset of $[-\alpha_2, \alpha_2]^m$ for $1 \leq i \leq n_2$. Recall that $\mathbf{T}^{(1)}$ determines the number of differences in G'_1 , and $|S_2|$ determines the amount of points lost in step 4(b) of PRDC. Then we have, marginalizing over all possible choices of $\mathbf{T}^{(1)}$ compatible with $|S_2| = n_2$,

$$\begin{aligned} & \mathbb{P} \left[\mathbf{z}_i \in I_i \forall 1 \leq i \leq n_2 \mid |S_2| = n_2 \right] \\ &= \sum_{\mathbf{T}^{(1)}} \mathbb{P} \left[\mathbf{z}_i \in I_i \forall 1 \leq i \leq n_2 \mid \mathcal{T}^{(1)} = \mathbf{T}^{(1)}, |S_2| = n_2 \right] \mathbb{P} \left[\mathcal{T}^{(1)} = \mathbf{T}^{(1)} \mid |S_2| = n_2 \right] \end{aligned}$$

By Lemma 26,

$$\mathbb{P} \left[\mathbf{z}_i \in I_i \forall 1 \leq i \leq n_2 \mid \mathcal{T}^{(1)} = \mathbf{T}^{(1)}, |S_2| = n_2 \right] = \mathbb{P} [\mathbf{u}_i \in I_i \forall 1 \leq i \leq n_2]$$

where $\mathbf{u}_1, \dots, \mathbf{u}_{n_2} \stackrel{iid}{\sim} \text{Tri}[-\alpha_2, \alpha_2]^m$. Hence,

$$\mathbb{P} \left[\mathbf{z}_i \in I_i \forall 1 \leq i \leq n_2 \mid |S_2| = n_2 \right] = \mathbb{P}[\mathbf{u}_i \in I_i \forall 1 \leq i \leq n_2],$$

which confirms the first bullet point of Proposition 10 for the base case $t = 2$. Following a similar marginalization procedure, this also implies by Lemma 26 that v_2 , the single vector output by **PRDC**, is independent of S_2 conditionally on $|S_2|$.

Now we handle the inductive step. Let $S_t = \mathbf{y}_1, \dots, \mathbf{y}_{n_t}$ and v_t denote the vectors output by the $(t-1)^{\text{th}}$ phase of **PRDC**. Suppose that conditionally on $\mathcal{F} := \{|S_2| = n_2, \dots, |S_t| = n_t\}$ that S_t is an iid sample of triangularly distributed vectors on $[-\alpha_t, \alpha_t]^m$, and v_t is independent of S_t . By Proposition 25, conditionally on \mathcal{F} , $|S_{t+1}| = n_{t+1}$, and the configuration vector $\mathcal{T}^{(t)} = \mathbf{T}^{(t)}$, the sample S_{t+1} is an iid collection of triangularly distributed vectors on $[-\alpha_{t+1}, \alpha_{t+1}]^m$. Hence, conditioning on $\mathcal{F} \cup \{|S_{t+1}| = n_{t+1}\}$ and applying the same marginalization over the configuration vector $\mathbf{T}^{(t)}$ as in the base case yields the first bullet point of Proposition 10 for the inductive step. The second bullet point follows similarly. ■

The next lemma is used in Appendix E. We omit its proof because it is similar to that of Proposition 10.

Lemma 27 *Let X_1, \dots, X_n be iid random vectors, each having a joint density $g : [-\Delta, \Delta]^m \rightarrow \mathbb{R}$. Apply **GKK** to the matrix \mathbf{X} with columns X_1, \dots, X_n , and consider the good points G'_t created from random differencing in step 3 of the t^{th} phase of **PRDC**. Also consider the random vector $v_t^{(0)}$ formed in step 4(a) of **PRDC**. Then conditioned on $|S_j| = n_j$ for $1 \leq j \leq t$ and $|G'_t| = g'_t$,*

- *the random vectors in G'_t form an independent sample of size g'_t from $\text{Tri}[-\alpha_{t+1}, \alpha_{t+1}]^m$.*
- *The random vector $v_t^{(0)}$ is independent of the vectors in G'_t .*