On the discrepancy of random low degree set systems

Nikhil Bansal* Raghu Meka †

March 4, 2020

Abstract

Motivated by the celebrated Beck–Fiala conjecture, we consider the random setting where there are n elements and m sets and each element lies in t randomly chosen sets. In this setting, Ezra and Lovett showed an $O((t \log t)^{1/2})$ discrepancy bound when $n \leq m$ and an $O(1)$ bound when $n \gg m^t$.

In this paper, we give a tight $O(\sqrt{t})$ bound for the entire range of n and m, under a mild assumption that $t = \Omega((\log \log m)^2)$. The result is based on two steps. First, applying the partial coloring method to the case when $n = m \log^{O(1)} m$ and using the properties of the random set system we show that the overall discrepancy incurred is at most $O(\sqrt{t})$. Second, we reduce the general case to that of $n \leq m \log^{O(1)} m$ using LP duality and a careful counting argument.

1 Introduction

Let $(V, \mathcal{S})$ be a set system with $V = [n]$ and $\mathcal{S} = \{S_1, \ldots, S_m\}$ a collection of subsets of V. For a two-coloring $x \colon V \to \{-1, 1\}$, the discrepancy of a set S is defined as $x(S) = |\sum_{i \in S} x(i)|$, and measures the imbalance from an even-split of S. The discrepancy of the system $(V, \mathcal{S})$ is defined as

$$\operatorname{disc}(\mathcal{S}) = \min_{x \colon V \to \{-1,1\}} \max_{S \in \mathcal{S}} x(S).$$

That is, it is the minimum imbalance of all sets in $\mathcal{S}$, over all possible two-colorings of V. More generally, we define the discrepancy for a $m \times n$ matrix A, as $\operatorname{disc}(A) = \min_{x \in \{-1,1\}^n} \|Ax\|_\infty$. Note that $\operatorname{disc}(\mathcal{S}) = \operatorname{disc}(A)$ if A is the incidence matrix of the system $(V, \mathcal{S})$.

Discrepancy is a widely studied topic and has applications to many areas in mathematics and computer science. For more background we refer the reader to the books [6, 15, 7]. In particular, discrepancy is closely related to the problem of rounding fractional solutions to a linear system of equations [13], and has found several applications in approximation algorithms and optimization.

An important problem, motivated by the rounding fractional solutions to column-sparse linear systems, is to understand the discrepancy of sparse systems where each element $i \in [n]$ lies in at most t sets. In a classic result, Beck and Fiala [4] showed that the discrepancy of such systems is at most $2t - 1$. This bound was recently improved by Bukh to $2t - \log^* t$ [5]. Improved bounds with a better dependence on t, but at the expense of dependence on n, are also known and after a long line of work, the best such bound is

*CWI and TU Eindhoven, Netherlands. bansal@gmail.com. Supported by a NWO VICI grant 639.023.812, and an ERC consolidator grant 617951.

†UCLA. raghum@cs.ucla.edu. Supported by NSF grant CCF-1553605.

$O(t^{1/2}(\log n)^{1/2})$ due to Banaszczyk [1]. These results have also been made algorithmic in recent years [2, 3, 12].

It is a long-standing conjecture that the discrepancy of such set systems is $O(t^{1/2})$ [4]. Despite much work, the problem is open even for very special cases such as when the hypergraph corresponding to the set system is simple, i.e. any two sets intersect in at most one element. Another interesting question to get the tight $O(t^{1/2})$ bound in the case when we have the additional property that the sets are also of size at most t. Here the best known bound is $O((t \log t)^{1/2})$ based on a direct application of the Lovász Local Lemma.

Random set system model. Recently, Ezra and Lovett [8] consider the problem in a natural random model, where there are n elements and m sets and each element $i \in [n]$ lies in exactly t random sets. That is, the t-tuple of sets containing i is chosen uniformly at random among the $\binom{m}{t}$ possibilities. In the following, by a random set system we refer to this model.

Ezra and Lovett [8] proved the following two results in the random model. (i) For $n \leq m$, the expected discrepancy is $O((t \log t)^{1/2})$, and (ii) for $n \gg m^t$, the expected discrepancy is $O(1)$. We remark that an $\Omega(t^{1/2})$ lower bound on the expected discrepancy also holds in the random model (e.g. when $n = m = 2t$, as can be seen easily using the spectral lower bound method [6]).

There are two natural questions left open from their work. First, whether these results can be extended to the entire range of n and m, i.e. when $n \in [m, m^t]$. This is particularly interesting, as the result of [8] in the regime when $n \leq m$ is based on Lovász Local Lemma, which fails for inherent reasons[1] when $n \gg m$. A second natural question is whether their bound can be improved to the optimum bound of $O(t^{1/2})$, especially for the important case of $n = m$. Again, the local lemma inherently loses an additional $(\log t)^{1/2}$ factor when $n = \Theta(m)$.

1.1 Our results and overview

Our main result addresses both these questions, and is the following.

Theorem 1. *Let $(V, \mathcal{S})$ be a random set system on n elements and m sets, where each element lies in t sets. Then, for every n and m, there is an algorithm that runs in time polynomial in n and m, and finds a coloring with expected discrepancy $O(t^{1/2})$, provided that $t = \Omega((\log \log m)^2)$.*

The result is based on two main ideas.

Reduction of n to k. We show that the problem with arbitrary n, m, t can be *reduced* to the case of $n \leq k$, where $k = Cm \log^2 m$, with high probability, for a fixed constant C. More precisely, let A be the $m \times n$ incidence matrix of the random set system, and let a_i denote the i-th column of A. We first apply the Beck–Fiala theorem [4] to the vectors $a_{k+1}, \ldots, a_n$ to find $x' \colon \{k+1, \ldots, n\} \to \{-1, 1\}$ such that the signed sum $b := \sum_{i>k} x'(i) a_i$ satisfies $\|b\|_\infty \leq 2t - 1$.

We show that with high probability, there is a *fractional* coloring $x'' \colon [k] \to [-1, 1]$ of $a_1, \ldots, a_k$ with discrepancy vector $\sum_{i \in [k]} x''(i) a_i$ exactly equal to $-b$. Together x'' and x' give a fractional coloring of $[n]$ with discrepancy 0, where the elements $1, \ldots, k$ have colors in $[-1, 1]$, while the elements $k+1, \ldots, n$ are

[1] As the average set size is $nt/m \gg t$.

colored $\{-1, 1\}$. As the first k columns are still random, this gives a "reduction" of the random Beck–Fiala problem from general n to k.

The existence of the coloring x'' follows from the following result.

Theorem 2. *For all $c > 0$, there exists a constant $C > 0$ such that the following holds. Let $a_1, \dots, a_k \in \{0, 1\}^m$ be random vectors with t ones. Let $P := \{\sum_{i=1}^k a_i x_i : \ x_i \in [-1, 1], i \in [k]\} \subset \mathbb{R}^m$ be the set of discrepancy vectors achievable by fractional colorings of $a_1, \dots, a_k$. Then for $k \geq Cm \log^2 m$, with probability at least $1 - 1/m^c$, it holds that $2tB_\infty^m \subset P$, where $B_\infty^m = \{y \in \mathbb{R}^m : \|y\|_\infty \leq 1\}$ is the unit ℓ_∞-ball in $\mathbb{R}^m$.*

To prove Theorem 2, we use LP duality to give an equivalent condition for the property $2tB_\infty^m \subset P$, and then use a counting argument to show that this condition is satisfied with high probability for k random vectors. We first prove a weaker bound of $k = O(m^3 \log m)$ using a standard ϵ-net argument. Later, we give a much more careful argument to improve this to $k = O(m \log^2 m)$.

Partial Coloring. It remains to modify the fractional coloring x'' on $[k]$ to an integral $\{-1, +1\}$ coloring, while incurring low discrepancy. To achieve this, we apply the partial coloring procedure of Lovett and Meka [14] over $O(\log k)$ iterations. The main issue here is to ensure that the overall discrepancy stays bounded by $O(t^{1/2})$ after all the iterations. To this end, we use the property that the starting set system on k columns is random to control the parameters used in the partial coloring lemma. As the partial coloring method gives no control on which subset of the original k columns remains after each iteration, we incur a penalty due to a union bound over all the subsets of the original columns. This is where we require that k is not too large relative to m. In particular, we show the following.

Theorem 3. *Let A be a random $m \times k$ matrix where $k \geq 2m$ and each column has t ones. Then, there is an algorithm that runs in time polynomial in k, and for $t = \Omega(\log^2(ek/m))$ and any $x^{(0)} \in [-1, 1]^k$, with probability at least $1 - \exp(-t)$, finds $x \in \{-1, 1\}^k$ such that for all rows v_j of A, $|\langle v_j, x - x^{(0)} \rangle| = O(\sqrt{t})$.*

Combining Theorems 2 and 3 directly gives Theorem 1. In particular, the condition $t = \Omega((\log \log m)^2)$ arises as $t = \Omega(\log^2(ek/m))$ and $k/m = O(\log^2 m)$.

We first prove Theorem 3 in Section 2, and then prove Theorem 2 in Section 3.

Related work. Very recently, two other groups [11, 9] have independently obtained related results. These results consider the regime where $n \gg m$, and use Fourier-analytic methods to show that an $O(1)$ discrepancy can be achieved for random low degree systems for $n = \tilde{\Omega}(m^2)$ [11] and $n = \tilde{\Omega}(m^3)$ [9], where $\tilde{\Omega}(\cdot)$ ignores polylogarithmic factors. Their results are non-algorithmic.

1.2 Preliminaries

We recall the standard Hoeffding/Chernoff bounds (see e.g., [10] for reference).

Lemma 4. *Let $X_1, \dots, X_n$ be independent random variables. Assume that $0 \leq X_i \leq M$ and $\mathbb{E}[X_i] = \mu_i$ for all i, and let $X = X_1 + \cdots + X_n$ and $\mu = \mathbb{E}[X]$. Then for any $\delta > 0$,*

$$\Pr[X \geq (1 + \delta)\mu] \leq \left(\frac{e^\delta}{(1 + \delta)^{1+\delta}} \right)^{\mu/M}.$$

For $\delta > 2e - 1$ this gives,

$$\Pr[X \geq (1+\delta)\mu] \leq \exp(-(1+\delta)\ln(1+\delta)\mu/2M).$$

Similarly, for $0 \leq \delta < 1$,

$$\Pr[X \leq (1-\delta)\mu] \leq \exp(-\delta^2\mu/2M).$$

Stochastic Domination. For random variables X and Y, we say that X stochastically dominates Y if $\Pr[X \geq a] \geq \Pr[Y \geq a]$ for all $a \in \mathbb{R}$. It is well-known that X dominates Y iff there is a coupling $(\hat{X}, \hat{Y})$ of X and Y such that $\Pr[\hat{X} \geq \hat{Y}] = 1$. It directly follows from this coupling that if $X_1, \ldots, X_n$ are independent copies of X and $Y_1, \ldots, Y_n$ are independent copies of Y, then $\Pr[X_1 + \cdots + X_n \geq t] \geq \Pr[Y_1 + \cdots + Y_n \geq t]$ for any $t \in \mathbb{R}$.

Partial Coloring Lemma. The algorithmic partial coloring lemma due to Lovett and Meka [14], takes as input some fractional coloring and target discrepancy bounds for each row, and finds another partial coloring satisfying these row-wise discrepancy bounds and where at least half the variables are set to -1 or $+1$.

Lemma 5. *(Partial Coloring Lemma [14]). Let $v_1, \ldots, v_m \in \mathbb{R}^n$, and $x_0 \in [-1,1]^n$ be a starting point. There is an efficient randomized algorithm that given parameters $c_1, \ldots, c_m$ such that $\sum_{j=1}^m \exp(-c_j^2/16) \leq n/16$ and $\delta > 0$, runs in time $O((m+n)^3 \cdot \delta^{-2} \cdot \log(nm/\delta))$ and with probability at least 0.1 finds a point $x \in [-1,1]^n$ such that*

1. *$|\langle x - x_0, v_j \rangle| \leq c_j \|v_j\|_2$ for each $j \in [m]$.*
2. *$|x_i| \geq 1 - \delta$ for at least $n/2$ indices $i \in [n]$.*

Note that the probability of success can be boosted by running the algorithm multiple times, and we will assume that the probability of failure of the algorithm is exponentially small in n and m. When $|x_i| \geq 1 - \delta$, we say that variable i is *frozen* and otherwise it is *alive*. Setting $\delta = 1/n$, rounding the frozen variables to the nearest -1 or $+1$ at the end of the algorithm can lead to an additional discrepancy of at most 1. So we will assume henceforth that $\delta = 0$.

2 Coloring the k columns

We now prove Theorem 3. We first describe the algorithm, and then give the analysis.

The algorithm. Our input consists of a random matrix A with $n_0 \leq k$ columns, and some fractional coloring $x^{(0)} \in [-1,1]^{n_0}$.

The algorithm will proceed in several iterations $i = 0, 1, \ldots, O(\log n_0)$. At the beginning of iteration i, let S_i denote some subset of columns corresponding to the alive variables, and let $A^{(i)}$ and $x^{(i)}$ denote the matrix A and the fractional coloring restricted to these columns. We initialize $A^{(0)} = A$. We use j to index the rows.

The iterations will consist of three different phases: (i) $i = 0$, (ii) $1 \leq i \leq \log t$ and (iii) $i > \log t$.

- *Phase 0.* Given $A^{(0)}$ and the starting coloring $x^{(0)}$, we reduce the number of fractional variables to $n_1 = m$, by picking any basic feasible solution to the linear program

$$A^{(0)}x = A^{(0)}x^{(0)} \qquad \text{subject to } -1 \le x_i \le 1 \quad \text{for } i \in [n_0]$$

  As $A^{(0)}x = A^{(0)}x^{(0)}$ consists of at most m linearly independent constraints, the solution will have at most m variables that are not set to -1 or 1. Note that the resulting matrix $A^{(1)}$ is no longer random.

  Below, we will assume that $n_i = m2^{1-i}$ for $i = 1, 2, \ldots$, by (possibly) adding some columns already colored $\{-1, 1\}$.

- *Phase 1.* For each iteration $i = 1, \ldots, \log_2 t$, apply the algorithm in Lemma 5 to $A^{(i)}$ and $x^{(i)}$ with the discrepancy bound

$$c_j \|v_j\|_2 = ct^{1/2}/i^2, \tag{1}$$

  where c is some fixed constant that will be specified later. If the condition $\sum_{j=1}^{m} \exp(-c_j^2/16) \le n_i/16$ in Lemma 5 is not satisfied, declare a failure and abort the algorithm.

- *Phase 2.* For $i > \log_2 t$, we apply the algorithm in Lemma 5 with $c_j = 0$ for sets larger than $ct^{1/2}$ and $c_j = \infty$ otherwise. Again, the algorithm aborts if $\sum_{j=1}^{m} \exp(-c_j^2/16) \le n_i/16$ does not hold during any iteration. If all the sets have size at most $ct^{1/2}$, the algorithm terminates and the remaining uncolored elements are set arbitrarily to $\{-1, 1\}$.

2.1 Analysis

If the algorithm does not abort in any iteration, clearly the total discrepancy for any set in Phase 1 is at most $ct^{1/2} \sum_{i \ge 1} i^{-2} = O(t^{1/2})$. Moreover, in Phase 2 a set incurs zero discrepancy as long as it is larger than $ct^{1/2}$, hence the total discrepancy added is also at most $ct^{1/2}$.

We will show that the probability that the algorithm aborts is at most $\exp(-t)$. If the algorithm aborts, then we simply output the $O(t)$ discrepancy coloring given by the Beck–Fiala Theorem [4]. Clearly, the expected discrepancy of the resulting algorithm is $O(t \exp(-t)) + O(t^{1/2}) = O(t^{1/2})$, hence proving Theorem 3.

We begin with a simple lemma that we will use repeatedly.

Lemma 6. *Let A be a $m \times \ell$ random matrix where each column has t ones, and let M be some fixed subset of r rows of A. For $s \ge 10t\ell/m$, let $B(s)$ denote the event that each row of M contains at least s 1's. Then over the random choice of A,*

$$\Pr[B(s)] \le \exp(-rs \log(sm/t\ell)/2).$$

Proof. For $i \in [\ell]$, let X_i denote the number of 1's in column i of M. Each X_i is independent and has the hypergeometric distribution $H(m, t, r)$ with mean $\mathbb{E}[X_i] = tr/m$. Let Y_i denote the sum of r independent Bernoulli random variables each with mean t/m, and note that Y_i is distributed as the binomial $\text{Bin}(r, t/m)$. It is well known that $H(m, t, r)$ is more sharply concentrated around its mean than the corresponding binomial distribution $\text{Bin}(r, t/m)$, and in particular $\exp(\lambda X_i) \le \exp(\lambda Y_i)$ for all $\lambda \in \mathbb{R}$, (see e.g. [10], page 395). This implies that the tail bounds in Lemma 4 for $\sum_{i=1}^{\ell} Y_i$ also hold for $\sum_{i=1}^{\ell} X_i$.

As $B(s)$ implies that $\sum_{i=1}^{\ell} X_i \ge rs$, we have that $\Pr[B(s)] \le \Pr[\sum_{i=1}^{\ell} X_i \ge rs]$. By the discussion above, to bound $\Pr[B(s)]$, it suffices to use Lemma 4 to upper bound $\Pr[\sum_{i=1}^{\ell} Y_i \ge rs]$. As $\sum_{i=1}^{\ell} Y_i$ is distributed as $\text{Bin}(r\ell, p)$, and hence as a sum of $r\ell$ independent Bernoulli random variables with mean p,

using Lemma 4 with $\mu = pr\ell$ and $(1+\delta) = rs/\mu = s/(p\ell) = sm/t\ell$, which is at least 10 (and hence $> 2e-1$) by our assumption, we have that

$$\begin{aligned}\Pr[B(s)] &\leq \Pr\left[\sum_{i=1}^{\ell} Y_i \geq rs\right] = \Pr\left[\sum_{i=1}^{\ell} Y_i \geq (1+\delta)\mu\right] \\ &\leq \exp(-\mu(1+\delta)\log(1+\delta)/2) = \exp(-rs\log(sm/t\ell)/2). \quad \square\end{aligned}$$

We next bound the failure probability in Phase 2, and then in Phase 1.

Lemma 7. *The probability that the algorithm fails during Phase 2 is at most* $\exp(-t)$.

Proof. Consider some iteration i for $i > \log_2 t$. Let $\ell = n_i = m2^{1-i}$. The iteration i aborts if the number of rows with size at least $s := ct^{1/2} + 1$ exceeds $\ell/16$. Call such rows *big* and let $r = \ell/16$. Note that $sm/t\ell \geq c2^{i-1}/t^{1/2} \geq 10$ as $i \geq \log_2 t$.

By Lemma 6 and taking a union bound over all such $r \times \ell$ submatrices of $A^{(0)}$, this probability is at most

$$\begin{aligned}\binom{n_0}{\ell}\binom{m}{r} \cdot \Pr[B(s)] &\leq \binom{k}{\ell}\binom{m}{r} \cdot \exp(-rs\log(sm/t\ell)/2) \\ &\leq \left(\frac{ek}{\ell}\right)^{2\ell} \cdot \exp\left(-\frac{\ell s \log(sm/t\ell)}{32}\right). \qquad (2)\end{aligned}$$

where the second step uses that $\binom{m}{r} \leq \binom{k}{\ell}$ as $k \geq 2m$ and $r \leq \ell \leq m$. Writing $\gamma = ek/m$, we have $ek/\ell = \gamma m/\ell$ and the expression in (2) becomes

$$\exp(2\ell(\log\gamma + \log m/\ell - s/64 - \log sm/t\ell)) = \exp(2\ell(\log\gamma - s/64 - \log s/t))$$

Choosing $c \geq 128$ so that $s/64 \geq 2t^{1/2}$, this is at most $\exp(-2\ell(2t^{1/2} - \log\gamma - \log t))$. The lower bound on t in Theorem 3 implies that $t^{1/2} = \Omega(\log\gamma)$ and hence this is at most $\exp(-\ell t^{1/2})$.

As $\ell = m2^{1-i}$ in iteration i and as $\ell \geq 2t^{1/2}$ in each iteration, the overall probability of failure over all iterations $i > \log_2 t$ is at most $\sum_i \exp(-m2^{1-i}t^{1/2}) \leq e^{-t}$. $\square$

Lemma 8. *The probability that the algorithm fails during Phase 1 is at most* $\exp(-\Omega(m/\log^4 t))$.

Proof. Let us fix an iteration $i \leq \log_2 t$ in phase 1, and let $d = ct^{1/2}/i^2$ be the discrepancy bound in (1) and $\ell = m2^{1-i}$ be the number of variables. Call a row *small* if it has size $s \leq s_0$, where $s_0 = d^2/(80i) = c^2t/(80i^5)$. For a row j of size s, note that by (1), $c_j = d/\sqrt{s}$, and hence the contribution of small rows to $\sum_{j=1}^{m} \exp(-c_j^2/16)$ is at most

$$m\exp(-d^2/16s_0) \leq m\exp(-5i) \leq \ell/32.$$

It remains to show that, with high probability, the contribution of the remaining rows to $\sum_{j=1}^{m} \exp(-c_j^2/16)$ is also at most $\ell/32$. To this end, we conservatively assume that $c_j = 0$ for big rows and hence we need to bound only the probability that there are more than $\ell/32$ rows.

If we pick ℓ columns from the random matrix $A^{(0)}$, the expected row size is $\mu = t\ell/m = t2^{1-i}$. As $s_0 = c^2t/80i^5$ and $\mu = t2^{1-i}$, we can pick c large enough so that $s_0 \geq 10\mu$, and hence $s_0 \geq 10t\ell/m$.

Let $r = \ell/32$. By Lemma 6, and a union bound over all $r \times \ell$ submatrices of $A^{(0)}$, the probability of having more than $\ell/32$ rows of size at least s_0 is bounded by

$$\begin{aligned}\binom{k}{\ell}\binom{m}{\ell/32}\Pr[B(s_0)] &\leq \binom{k}{\ell}^2 \cdot \exp\left(-\frac{\ell s_0 \log(s_0 m/t\ell)}{64}\right) \\ &\leq \exp(2\ell(\log\gamma + \log m/\ell) - \ell s_0 \log(s_0 m/t\ell)/64), \qquad \text{(where } \gamma = ek/m\text{)}.\end{aligned}$$

Noting that $s_0 \log(s_0 m/t\ell) = \Omega(t/i^4)$, the probability above is at most

$$\exp(2\ell(\log\gamma + i) - \ell \cdot \Omega(t/i^4)).$$

As $\log\gamma = O(t^{1/2})$, this is $\exp(-\Omega(-\ell t/i^4)$. As $i \leq \log_2 t$ we have $\ell = m2^{1-i} \geq m/t$, and so the overall failure probability over the iterations $i \leq \log_2 t$ is at most $\exp(-\Omega(m/\log^4 t))$. □

3 Reducing the number of columns

We now prove Theorem 2.

The fractional discrepancy polytope. Let $a_1, \ldots, a_k$ be arbitrary vectors in $\mathbb{R}^m$. Consider the polytope $P := \left\{\sum_{i=1}^k a_i x_i : \ x_i \in [-1,1]\right\}$ of discrepancy vectors obtained by all possible fractional colorings. The convex hull of P is given by its 2^k extreme points

$$p_x := \sum_i x_i a_i \text{ for } x \in \{-1,+1\}^k.$$

For $p \geq 1$, let $B_p^m = \{y \in \mathbb{R}^m : \|y\|_p \leq 1\}$ denote the ℓ_p ball in $\mathbb{R}^m$. For brevity, let $Q := 2tB_\infty^m$. The following lemma characterizes exactly when $Q \subset P$.

Lemma 9. *$Q \subset P$ iff $\|y^T A\|_1 > 2t$, for all $y \in B_1^m$, where A is the matrix with columns $a_1, \ldots, a_k$.*

Proof. Suppose $Q \not\subset P$. As P and Q are convex, by Farkas' lemma, there exists a hyperplane given by normal y, that separates some point $q \in Q \setminus P$ from P. As $0 \in P$, we can assume that there is some $s > 0$ such that $y^T q > s$ and $y^T p_x < s$ for each extreme point p_x of P. As

$$\max_{x\in\{-1,1\}^k} y^T p_x = \max_{x\in\{-1,1\}^k} \sum_{i=1}^k (y^T a_i) x_i = \sum_{i=1}^k |y^T a_i| = \|y^T A\|_1, \tag{3}$$

this is the same as saying that there is some s such that $y^T q > s > \|y^T A\|_1$.

By scaling, we can assume that $\|y\|_1 = 1$ and as $Q = 2tB_\infty$ we have $\max_{q\in Q} y^T q = 2t$. This gives that $s < 2t$, and thus there is some y with $\|y\|_1 = 1$ and $\|y^T A\|_1 < 2t$; a contradiction.

Conversely if $Q \subset P$, then no direction exists that separates some point $q \in Q$ from P, which implies for each y with $\|y\|_1 = 1$, there is some $p_x \in P$ such that $y^T p_x > 2t$, which by (3) gives that $\|y^T A\|_1 > 2t$. □

So Theorem 2 will follow by showing that with high probability, $\|y^T A\|_1 > 2t$ for every $y \in B_1^m$.

3.1 A weaker bound on k

We first sketch a simpler bound of $k = \Omega(m^3 \log m)$. Together with Theorem 3, this implies the $O(t^{1/2})$ discrepancy bound in Theorem 1 for $t = \Omega(\log^2 m)$.

Theorem 10. *Let A be a $m \times k$ random matrix where each column has t ones and $k = O(m^3 \log m)$. Then with probability at least $1 - \exp(-\Omega(m \log m))$, it holds that $\|y^T A\|_1 > 2t$ for all $y \in B_1^m$.*

Let $\delta > 0$, and let N_δ be the set of points $y' \in \mathbb{R}^m$ such that each coordinate y'_i of y' is an integral multiple of δ, and $\|y'\|_1 \leq 1$. Clearly $|N| \leq (3/\delta)^m$ and for any point $y \in B_1^m$ there is some point y' in N_δ with $\|y - y'\|_1 \leq m\delta$.

We fix $\delta = 3/km$ and note that $|N_\delta| \leq \exp(m \log(km))$. As $|(y - y')^T a| \leq \|y - y'\|_1 \|a\|_\infty \leq m\delta$ for any $a \in [-1, 1]^m$, to show Theorem 10 it suffices to show that $\|y^T A\|_1 > 2t + km\delta = 2t + 3$ for all $y \in N_\delta$ with $\|y\|_1 \geq 1 - m\delta \geq 9/10$.

Fix a vector y in the net N_δ with $\|y\|_1 \geq 9/10$. Let X denote the random variable $|y \cdot a|$, where $a \in \{0, 1\}^m$ is chosen randomly with exactly t ones. Henceforth, we assume that $t \leq m/10$, as for $t = \Theta(m)$, an $O(m^{1/2})$ discrepancy follows from the result of Spencer [16].

Lemma 11. *For every $y \in B_1^m$, $\mathbb{E}[X] \geq t/2m^2$, assuming $t \leq m/10$.*

Proof. First, as $X = |y \cdot a|$ for $a \in [-1, 1]^n$ and $\|y\|_1 \leq 1$, we have that $0 \leq X \leq 1$, and hence $\mathbb{E}[X] \geq \mathbb{E}[X^2]$. So it suffices to lower bound the second moment as follows.

$$\begin{aligned} \mathbb{E}[X^2] &= \mathbb{E}[(\sum_i a_i y_i)^2] = \sum_i \mathbb{E}[a_i] y_i^2 + \sum_{i \neq j} E[a_i a_j] y_i y_j = \frac{t}{m} \sum_i y_i^2 + \sum_{i \neq j} \frac{t(t-1)}{m(m-1)} y_i y_j \\ &= \frac{t(m-t)}{m(m-1)} \sum_i y_i^2 + \frac{t(t-1)}{m(m-1)} (\sum_i y_i)^2 \geq t/(2m) \sum_i y_i^2 \geq \frac{t}{2m^2}. \quad \square \end{aligned}$$

As $\|y^T A\|_1$ is the sum of k independent random variables $X_1, \ldots, X_k$ distributed as X, using the lower bound on $\mathbb{E}[X]$ and as $0 \leq X \leq 1$, by Lemma 4,

$$\Pr[X_1 + \cdots + X_k < \frac{1}{2} \frac{tk}{2m^2}] \leq \exp(-kt/8m^2) = \exp(-ctm \log m),$$

for $k \geq 8cm^3 \log m$. Taking a union bound over all the points in the net N_δ, and choosing c large enough, Theorem 10 follows as

$$|N_\delta| \exp(-cmt \log m) \leq \exp(m(\log km - tc \log m)) \leq \exp(-\Omega(m \log m)).$$

3.2 A stronger bound on k

Plugging the above bound of $k = O(m^3 \log m)$ in Theorem 3 already implies a $O(t^{1/2})$ discrepancy bound when $t = \Omega((\log k/m)^2) = \Omega(\log^2 m)$. So henceforth we can assume that $t \ll \log^2 m$ (in the argument below we will need that $t = o(m^{1/2})$. We will now prove a refined bound of $k = O(m \log^2 m)$.

We note that one cannot hope for $k = O(m)$, and in particular k must be at least $m(\log m)^{1-\epsilon}$ for any $\epsilon > 0$. This holds even if we require the condition $\|y^T A\|_1 > 2t$ to hold only for the coordinate vectors $y =$

$e_1, \ldots, e_m$. Suppose $k/m < \log^{1-\epsilon} m$, then as the expected number of ones in a row is $kt/m < t \log^{1-\epsilon} m$. For $t = \log^{\epsilon/2} m$ (which satisfies the condition $t = \Omega(\log^2(k/m))$, this is $\ll \log m$, so with probability at least $1 - m^{-\Omega(1)}$ some row j in A will consist of all zeros, and hence violate $\|e_j^T A\|_1 > 2t$.

We first give the main idea before describing the details.

The idea. Consider the net N_δ with $\delta = 1/km$ as before. By Theorem 10, we can assume that $k \leq m^3 \log m$ and hence $\delta = 1/\text{poly}(m)$.

For a point $y \in N_\delta$, let $Y(y)$ be the random variable $|y^T a|$, where a is a random column with t ones. We need to show that with high probability, for each y in the net, the sum of k independent copies of $Y(y)$ is more than $4t$. In the argument in Section 3.1, k had to be large as we were taking a union bound over the exponentially many points of the net N_δ. While we cannot reduce $|N_\delta|$ much, the idea here is to exploit the specific structure of the random vectors a and the event that we care about. For instance, if y is sparse, we get a not too small probability for Y being small, but then there aren't too many such sparse vectors in the net. We exploit such trade-offs below.

More precisely, we consider another random variable $X \geq 0$ that will be stochastically dominated by Y, and the value of X will (essentially) only depend on the values of a and only on the sign pattern of y in certain specific coordinates. This will lead to a much smaller loss in the union bound. We now give the details.

3.2.1 The argument

Fix some $y \in N_\delta$ and recall that $1/2 \leq \|y\|_1 \leq 1$. We say that coordinate i lies in class $j \in \{0, 1, \ldots, h-1\}$, for $h = \log(1/\delta) = O(\log m)$, if $|y_i| \in (2^{-j-1}, 2^{-j}]$. Moreover, i has *positive* sign if $y_i > 0$ and *negative* sign if $y_i < 0$. We will not care about coordinates with value 0. Let us define the *weight* of class j of y as $w_j(y) = \sum_{i \in \text{class } j} |y_i|$.

Define the *class* $c(y)$ of y as a class j with the highest weight. Let $c^-(y)$ and $c^+(y)$ denote the classes $c(y) - 1$ and $c(y) + 1$ respectively (if they exist). As $\|y\|_1 \geq 1/2$, the class $c(y)$ has weight at least $1/2h$. Let $n(y)$ denote the number of coordinates with class $c(y)$, and we thus have

$$2^{c(y)}/2h \leq n(y) \leq 2^{c(y)+1}.$$

As $c(y)$ is the maximum weight class, we also have that the number of coordinates of class $c^+(y)$ and $c^-(y)$ is at most $n(y)$ and $4n(y)$ respectively.

We now define the random variable X with the desired properties.

The random variable X. Let $i_1, \ldots, i_t$ denote the t locations of 1 in a, that are picked from $[m]$ without replacement. We use the principle of deferred decisions, and assume that the locations $i_1, \ldots, i_{t-1}$ have already been revealed, and that the randomness is only in the t-th choice.

Let $v = y_{i_1} + \cdots + y_{i_{t-1}}$, and note that

$$Y = |y^T a| = |v + y_{i_t}|.$$

Our random variable X will satisfy the following properties.

1. $X \geq 0$ and for each y, a, we have that $X \leq Y$.

2. For every y, X is completely determined by (i) the value of v, (ii) the sign pattern of coordinates in class $c^-(y), c(y), c^+(y)$, (iii) the location of i_t in these three classes (if it falls in these classes), and (iv) on whether any of $\{i_1, \ldots, i_{t-1}\}$ fall in these three classes.

We now define the random variable X, based on a few cases. The above properties are directly verified by inspection.

Let us first assume that $10t < n(y) < m/10$. The remaining side cases are handled easily later.

Balanced Case. We call class $c(y)$ *sign-balanced* if y has at least $n(y)/4$ coordinates in class $c(y)$ with both positive and negative signs.

If $v < 0$, we define $X = 2^{-c(y)-1}$ if i_t lies in class $c(y)$ and $y_{i_t} < 0$. Otherwise, $X = 0$.
Analogously, if $v > 0$, then $X = 2^{-c(y)-1}$ if i_t lies in class $c(y)$ and $y_{i_t} > 0$. Otherwise, $X = 0$.

Note that we always have $X \leq Y$, as in both cases either $X = 0$ or

$$X = 2^{-1-c(y)} \leq |v + y_{i_t}| = Y.$$

Moreover, as $n(y) \geq 10t$, irrespective of the locations of $i_1, \ldots, i_{t-1}$, the probability that i_t lies in class $c(y)$ is at least $(9/10)n(y)/m$. Finally, $X = 2^{-c(y)-1}$ with probability at least $(1/4) \cdot (9/10)n(y)/m \geq n(y)/8m$, irrespective of the value of v.

Unbalanced Case. Without loss of generality, suppose that class $c(y)$ has more than $3n(y)/4$ positive signs (the other case is symmetric). We consider two further cases.

If $v \notin (-2^{-c(y)+1/2}, -2^{-c(y)-3/2})$, we set $X = (1/8)2^{-c(y)}$ if i_t falls in class $c(y)$ and $y_{i_t} > 0$. Otherwise, $X = 0$.

We claim that $X \leq Y$. For, we either have $v \leq -2^{-c(y)+1/2}$, in which case

$$v + y_{i_t} \leq -2^{-c(y)+1/2} + 2^{-c(y)} = -(2^{1/2} - 1)2^{-c(y)} \leq -(1/4)2^{-c(y)},$$

so that $|v + y_{i_t}| \geq (1/4)2^{-c(y)}$. Similarly, if $v \geq -2^{-c(y)-3/2}$, then

$$v + y_{i_t} \geq -2^{-c(y)-3/2} + 2^{-c(y)-1} = 2^{-c(y)-1}(1 - 2^{-1/2}) \geq (1/8)2^{-c(y)},$$

so that $|v + y_{i_t}| \geq (1/8)2^{-c(y)}$. Finally, note that given that i_t falls in class $c(y)$, as the class is unbalanced, we have that $y_{i_t} > 0$ with probability at least $(3/4n(y) - (t-1))/n(y) \geq 1/2$.

On the other hand, if $v \in (-2^{-c(y)+1/2}, -2^{-c(y)-3/2})$, we set $X = 0$ if i_t lies in any of the classes $\{c^-(y), c(y), c^+(y)\}$. Otherwise, we set $X = (1/16)2^{-c(y)}$. We claim that $X \leq Y$ holds, because if i_t does not lie in any of the classes $\{c^-(y), c(y), c^+(y)\}$, then (i) either $|y_{i_t}| < 2^{-c(y)-2}$ in which case

$$|v + y_{i_t}| > |v| - |y_{i_t}| > 2^{-c(y)}(1/2\sqrt{2} - 1/4) \geq (1/16)2^{-c(y)}.$$

Or $|y_{i_t}| > 2^{-c(y)+1}$, in which case

$$|v + y_{i_t}| > |y_{i_t}| - |v| > 2^{-c(y)+1} - 2^{-c(y)+1/2} \geq 2^{-c(y)-1}.$$

In either case, we have $X \leq Y$. Moreover, as $n(y) \leq m/10$, there are at least $m - 6n(y) \geq 2m/5$ coordinates other than these three classes, so that above events happens with probability $\Omega(1)$.

Side cases. We now consider the remaining cases where $n(y) \leq 10t$ or $n(y) \geq m/10$.

If $n(y) < 10t$, we set $X = 0$ if some $i_1, \ldots, i_{t-1}$ already lies in $\{c^-(y), c(y), c^+(y)\}$. Otherwise, we proceed as above depending on whether the class $c(y)$ is balanced or unbalanced; it is easily seen that the previous arguments still hold. Further, the probability that any of $i_1, \ldots, i_{t-1}$ land in these three classes is at most $O(tn(y)/m) = O(t^2/m) \ll 1$.

If $n(y) > m/10$, in the balanced case we proceed as previously. The problem arises in the argument above when $c(y)$ is unbalanced (since we relied on the event that i_t falls outside the three classes happens with decent probability). So instead we do the following: Set $X = 2^{-c(y)}$ if at least $t/20$ $i_1, \ldots, i_{t-1}$ lie inside $c(y)$ and $|v| \geq 10 \cdot 2^{-c(y)}$. Else, set $X = 0$. Clearly, $X \leq Y$ as if $|v| > 10 \cdot 2^{-c(y)}$, then $|v + y_{i_t}| > |v| - |y_{i_t}| \geq 9 \cdot 2^{-c(y)}$.

As $n(y) > m/10$, the probability that fewer than $t/20$ indices $i_1, \ldots, i_{t-1}, i_t$ lie inside the class $c(y)$ is at most $\exp(-\Omega(t))$. Further, if we condition at least $t/20$ of $i_1, \ldots, i_{t-1}$ to lie in $c(y)$, then as $c(y)$ is unbalanced, the probability that $|v| \leq 10 \cdot 2^{-c(y)}$ is $\exp(-\Omega(t))$. So $|X| = 2^{-cy}$ with probability at least $1 - \exp(-\Omega(t)) \geq 1/2$.

3.2.2 The concentration argument

Fix a y and consider the random variable X as defined above. Then, by the arguments above, $\mathbb{E}[X] = \Omega(2^{-c(y)} \cdot n(y)/m)$. As $n(y) \geq 2^{c(y)}/2h$ and $h = O(\log m)$, we have that $\mathbb{E}[X] \geq c/m \log m$ for a fixed constant $c > 0$. Moreover, X is bounded by $M = 2^{-c(y)} \leq 2/n(y)$. Therefore, if we choose k independent copies of $X_1, \ldots, X_k$ of X, then by Lemma 4 with $\mu = k\,\mathbb{E}[X]$, we get

$$\Pr[X_1 + \cdots + X_k < \mu/2] \leq \exp(-\mu/8M) \leq \exp(-\Omega(n(y)k/m \log m)).$$

We now use our definition of the random variable X and take a union bound over the various quantities that the random variable X can depend on. There are at most $2/\delta$ choices for the value v. Now, consider some point y in the net N_δ of class $c(y)$. The behavior of X is completely determined by the sign pattern on $O(n(y))$ coordinates of y (corresponding to the sign-pattern of y restricted to classes $\{c^-(y), c(y), c^+(y)\}$). So in the union bound, we incur loss of $2^{6n(y)}$ (as $|c^-(y)| \leq 4n(y), |c^+(y)| \leq n(y)$). Further, we have at most $\binom{m}{4n(y)} \cdot \binom{m}{n(y)} \cdot \binom{m}{n(y)}$ possibilities for $\{c^-(y), c(y), c^+(y)\}$. Therefore, taking a union bound over all the possible random variables X, we get the failure probability for a fixed $n(y)$ to be at most

$$\exp(-\Omega(n(y)k/m \log m)) \cdot (2em)^{6n(y)} \ll \exp(-\Omega(n(y)C \log m)),$$

if we take $k = Cm \log^2 m$ for a sufficiently big constant C. Adding over all values of $n(y)$ we get that the failure probability in Theorem 2 is at most $m^{-\Omega(C)}$ for $k = Cm \log^2 m$ for C sufficiently big. This finishes the proof of Theorem 2.

Acknowledgements

We would like to thank the Simons Institute at Berkeley for hosting us when this work was done.

References

[1] W. Banaszczyk. Balancing vectors and gaussian measures of n-dimensional convex bodies. *Random Structures & Algorithms*, 12(4):351–360, 1998.

[2] N. Bansal. Constructive algorithms for discrepancy minimization. In *Foundations of Computer Science, FOCS*, pages 3–10, 2010.

[3] N. Bansal, D. Dadush, and S. Garg. An algorithm for Komlós conjecture matching Banaszczyk's bound. In *Foundations of Computer Science, FOCS*, pages 788–799, 2016.

[4] J. Beck and T. Fiala. Integer-making theorems. *Discrete Applied Mathematics*, 3(1):1–8, 1981.

[5] B. Bukh. An improvement of the Beck-Fiala theorem. *Combinatorics, Probability & Computing*, 25(3):380–398, 2016.

[6] B. Chazelle. *The discrepancy method: randomness and complexity*. Cambridge University Press, 2000.

[7] W. Chen, A. Srivastav, G. Travaglini, et al. *A Panorama of Discrepancy Theory*, volume 2107. Springer, 2014.

[8] E. Ezra and S. Lovett. On the Beck–Fiala conjecture for random set systems. In *APPROX/RANDOM*, pages 1–10, 2016.

[9] C. Franks and M. Saks. On the discrepancy of random matrices with many columns. *arXiv preprint arXiv:1807.04318*, 2018.

[10] A. Frieze and M. Karonski. *Introduction to Random Graphs*. Cambridge University Press, 2015.

[11] R. Hoberg and T. Rothvoss. A Fourier-analytic approach for the discrepancy of random set systems. In *Symposium on Discrete Algorithms, SODA*, pages 2547–2556, 2019.

[12] A. Levy, H. Ramadas, and T. Rothvoss. Deterministic discrepancy minimization via the multiplicative weight update method. In *Integer Programming and Combinatorial Optimization, IPCO*, pages 380–391, 2017.

[13] L. Lovász, J. Spencer, and K. Vesztergombi. Discrepancy of set-systems and matrices. *European Journal of Combinatorics*, 7(2):151–160, 1986.

[14] S. Lovett and R. Meka. Constructive discrepancy minimization by walking on the edges. *SIAM J. Comput.*, 44(5):1573–1582, 2015.

[15] J. Matoušek. *Geometric discrepancy: An illustrated guide*. Springer Science, 2009.

[16] J. Spencer. Six standard deviations suffice. *Trans. of Amer. Math. Soc.*, 289(2):679–706, 1985.