A Unified Approach to Discrepancy Minimization

Nikhil Bansal
University of Michigan, Ann Arbor, MI, USA

Aditi Laddha
Georgia Tech, Atlanta, GA, USA

Santosh Vempala
Georgia Tech, Atlanta, GA, USA

Abstract

We study a unified approach and algorithm for constructive discrepancy minimization based on a stochastic process. By varying the parameters of the process, one can recover various state-of-the-art results. We demonstrate the flexibility of the method by deriving a discrepancy bound for smoothed instances, which interpolates between known bounds for worst-case and random instances.

2012 ACM Subject Classification Theory of computation → Design and analysis of algorithms

Keywords and phrases Discrepancy theory, smoothed analysis

Digital Object Identifier 10.4230/LIPIcs.APPROX/RANDOM.2022.1

Category RANDOM

Related Version *Full Version*: https://arxiv.org/abs/2205.01023

Funding *Nikhil Bansal*: Supported in part by the NWO VICI grant 639.023.812.
Aditi Laddha: Supported in part by NSF awards CCF-2007443 and CCF-2134105.
Santosh Vempala: Supported in part by NSF awards CCF-2007443 and CCF-2134105.

Acknowledgements We are grateful to Yin Tat Lee and Mohit Singh for helpful discussions.

1 Introduction

Given a universe of elements $U = \{1, \ldots, n\}$ and a collection $\mathcal{S} = \{S_1, \ldots, S_m\}$ of subsets $S_i \subseteq U$, the discrepancy of the set system $\mathcal{S}$ is defined as

$$\mathrm{disc}(\mathcal{S}) = \min_{x:U\to\{-1,1\}} \max_{i\in[m]} \Big| \sum_{j\in S_i} x(j) \Big| .$$

That is, the discrepancy is the minimum imbalance that must occur in at least one of the sets in $\mathcal{S}$ over all bipartitions of U. More generally for an $m \times n$ matrix A, the discrepancy of A is defined as $\mathrm{disc}(A) = \min_{x\in\{-1,1\}^n} \|Ax\|_\infty$. Note that the definition for set systems corresponds to choosing A as the incidence matrix of $\mathcal{S}$, i.e., $A_{ij} = 1$ if $j \in S_i$ and 0 otherwise. Discrepancy is a well-studied area with several applications in both mathematics and theoretical computer science (see [14, 17, 28]).

Spencer's problem. In a celebrated result, Spencer [34] showed that the discrepancy of any set system with $m = n$ sets is $O(\sqrt{n})$, and more generally $O(\sqrt{n \log(2m/n)})$ for $m \geq n$. To show this, he developed a general partial-coloring method (a.k.a. the entropy method), building on a counting argument of Beck [13], that has since been used widely for various other problems. A similar approach was developed independently by Gluskin [20]. Roughly, here the elements are colored in $O(\log n)$ phases. In each phase, an $\Omega(1)$ fraction of the elements get colored while incurring a small discrepancy for each row.

Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2022).
Editors: Amit Chakrabarti and Chaitanya Swamy; Article No. 1; pp. 1:1–1:22
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Beck-Fiala and Komlós problems. Another central question is the Beck-Fiala problem where each element appears in at most k sets in $\mathcal{S}$. Equivalently, every column of the incidence matrix is k-sparse. The long-standing Beck-Fiala conjecture [15] states that $\mathrm{disc}(\mathcal{S}) = O(\sqrt{k})$. A further generalization is the Komlós problem, also called the vector balancing problem, about the discrepancy of matrices A with column ℓ_2-norms at most 1. Komlós conjectured that $\mathrm{disc}(A) = O(1)$ for any such matrix. Note that the Komlós conjecture implies the Beck-Fiala conjecture.

Banaszczyk showed an $O(\sqrt{\log n})$ bound for the Komlós problem based on a deep geometric result [3]. Here, the full coloring is constructed directly (in a single phase), and this result has also found several applications. The resulting $O(\sqrt{k \log n})$ bound for the Beck-Fiala problem is also the best known bound for general k.[1]

In contrast, the partial coloring method only gives weaker bounds of $O(\log n)$ and $O(k^{1/2} \log n)$ for these problems – the $O(\log n)$ loss is incurred due to the $O(\log n)$ phases of partial coloring.

Limitations of Banaszczyk's result. Even though Banaszczyk's method gives better bounds for the Komlós problem, it is not necessarily stronger, and is incomparable to the partial coloring method. E.g., it is not known how to obtain Spencer's $O(\sqrt{n})$ result (or anything better than the trivial $O(\sqrt{n \log n})$ random-coloring bound) using Banaszczyk's result. A very interesting question is whether there is a common generalization that unifies both these results and techniques.

Algorithmic approaches. Both the partial coloring method and Banaszczyk's result were originally non-algorithmic, and a lot of recent progress has resulted in their algorithmic versions. Starting with the work of [4], several different algorithmic approaches are now known for the partial coloring method [27, 33, 21, 18], based on various elegant ideas from linear algebra, random walks, optimization and convex geometry.

In further progress, an algorithmic version of the $O(\sqrt{\log n})$ bound for the Komlós problem was obtained by [5], see also [7], and [6] for the more general algorithmic version of Banaszczyk's result. In related work, Levy et al. [26] gave deterministic polynomial time constructive algorithms for the Spencer and Komlós settings matching $O(\sqrt{n \log(2m/n)})$ and $O(\sqrt{\log n})$ respectively.

A key underlying idea behind many of these results is to perform a discrete Brownian motion (random walk with small steps) in the $\{-1, 1\}^n$ cube, where the update steps are correlated and chosen to lie in some suitable subspace. However, the way in which these subspaces are chosen for the partial coloring method and the Komlós problem are quite different. We give a high level description of these approaches as this will be crucial later on.

In the partial coloring approach, the walk is performed in a subspace orthogonal to the *tight discrepancy constraints*. If the discrepancy for some row A_i reaches its target discrepancy bound, the update Δx to the coloring satisfies $A_i \cdot \Delta x = 0$. As the walk continues over time, the subspace dimension gets smaller and smaller until the walk is stuck. At this point, the subspace is reset and the *next phase* resumes.

On the other hand, the algorithm for the Komlós problem does not consider the discrepancy constraints at all, and chooses a different subspace with a certain sub-isotropic property which ensures the discrepancy incurred for a row is roughly proportional to its ℓ_2 norm,

[1] For $k = o(\log n)$ an improved bound follows from the $2k - 1$ bound by [15].

while ensuring that the rows with large ℓ_2-norm incur zero-discrepancy. In particular, in contrast to the partial coloring method, all the elements are colored in a *single phase*, and the discrepancy constraints are ignored.

The need for a combined approach. Even though the $O(\sqrt{k \log n})$ bound for the general Beck-Fiala problem is based on Banaszczyk's method, all the important special cases where the conjectured $O(\sqrt{k})$ bound holds are based on the partial coloring method. For example, Spencer's problem with $m = O(n)$ sets corresponds to special case of the Beck-Fiala problem with $k = O(n)$. So Spencer's six-deviations result resolves the Beck-Fiala conjecture for this case, which we do not know how to obtain from Banaszczyk's result.

The Beck-Fiala conjecture also holds for the case of random set systems with $m \geq n$. In particular, Potukuchi [32] considers the model where each column has 1's in k randomly chosen rows and shows that the discrepancy is $O(\sqrt{k})$ with high probability. See also [19, 9, 22, 1] for related results. Potukuchi's result crucially relies on the partial coloring approach, and it is not clear at all how to exploit the properties of random instances in Banaszcyck's approach.

Thus a natural question and a first step towards resolving the Beck-Fiala and Komlós conjecture, and making progress on other discrepancy problems, is whether there exist more general techniques to obtain both Spencer's and Potukuchi's result and the $O(\sqrt{k \log n})$ bound for the Beck-Fiala problem in a unified way.

1.1 Our results

We present a new unified framework that recovers all the results mentioned above, and various other state-of-the-art results as special cases. Our algorithm is based on a derandomization of a stochastic process that is guided by a barrier-based potential function. We were inspired by an elegant idea of Lee and Singh [23] who showed how the barrier function approach can be used to give a proof of Spencer's result without any partial coloring phases. A related idea was also explored in [21]. The barrier function approach itself has been used extensively in various settings such as graph and matrix sparsification [12, 24], covariance estimation [35], isoperimetric inequalities [25], bandit algorithms [2] and also in the context of discrepancy minimization [10, 21, 11].

Given a matrix A, the algorithm starts with the all-zero coloring x_0. Let $x_t \in [-1,1]^n$ be the coloring at time. The algorithm maintains a barrier $b_t > 0$ over time and defines the slack of row i at time t as

$$s_i(t) = b_t - \underbrace{\sum_{j=1}^{n} a_i(j) x_t(j)}_{\text{current discrepancy}} \underbrace{-\lambda \sum_{j=1}^{n} a_i(j)^2 (1 - x_t(j)^2)}_{\text{remaining variance}}. \tag{1}$$

Notice that when all $x_t(j)$ eventually reach ± 1, the *remaining variance* term is zero and the slack measures the gap between the discrepancy and the barrier.

We define the potential

$$\Phi(t) = \sum_i s_i(t)^{-p} \tag{2}$$

for some fixed $p > 1$, that penalizes the rows with small slacks and blows up to infinity if some slack approaches zero. If we can ensure that the slacks are always positive and the potential is bounded, then the discrepancy is upper bounded by value of the barrier when the algorithm terminates.

At each time step, the algorithm picks a random direction v_t that is orthogonal to some of the rows with the least slack, and satisfies some additional properties, and updates the coloring by a small amount in the direction v_t. The barrier b_t is also updated. These updates are chosen to ensure that the potential does not increase in expectation, and hence all the slacks stay bounded away from 0. We give a more detailed overview in Section 2.

By changing the parameters p, λ depending on the problem at hand, we obtain several results using a unified approach.

1. Set coloring [34]. For any set system on n elements and $m \geq n$ sets, $\mathrm{disc}(\mathcal{S}) = O(\sqrt{n \log(2m/n)})$.
2. Komlós problem [7]. For any $A \in \mathbb{R}^{m \times n}$ with columns norms $\left\|A^j\right\|_2 \leq 1$, $\mathrm{disc}(A) = O(\sqrt{\log n})$.
3. Random/Spectral Hypergraphs [32]. Let $A \in \{0,1\}^{m \times n}$ be the incidence matrix of a set system with n elements and m sets, where element lies in at most k sets and let $\gamma = \max_{v \perp \mathbf{1}, \|v\|=1} \|Av\|$. Then for $m \geq n$, $\mathrm{disc}(\mathcal{S}) = O(\sqrt{k} + \gamma)$.
4. Gaussian Matrix [16]. For a random matrix $A \in \mathbb{R}^{m \times n}$ with each entry $A_{ij} \sim \mathcal{N}(0, \sigma^2)$ independently, with probability at least $1 - (1/m^3)$, $\mathrm{disc}(A) = O\left(\sigma\left(\sqrt{n} + \sqrt{\log m}\right) \cdot \sqrt{\log \frac{2m}{n}}\right)$.

More generally, given a matrix A, we state the following result based on optimizing the various parameters of the algorithm, depending on the properties of A. This allows our framework to be applied in a black-box manner to a given problem at hand.

▶ **Theorem 1.** *For a $A \in \mathbb{R}^{m \times n}$ with $\left\|A^j\right\|_2 \leq L$ and $|a_i(j)| \leq M$ for all $i \in [m], j \in [n]$, let $h : \mathbb{R}^+ \to \mathbb{R}^+$ be a non-increasing function such that for every subset $S \subseteq [n]$ and $i \in [m]$,*

$$\sum_{j \in S} a_i(j)^2 \leq |S| \cdot h(|S|). \tag{3}$$

Then, for any $p > 1$, there exists a vector $x \in \{-1,1\}^n$ such that $\|Ax\|_\infty \leq 5b_0 + 2M$, where

$$b_0 = \min\left(\sqrt{8(p+1)(48m)^{1/p} \cdot \beta},\ 250L\sqrt{\log(2m)}\right). \tag{4}$$

where $\beta = \int_{t=0}^{n-2} h(n-t)(n-t)^{-1/p} dt$.

Let us see how Theorem 1 directly leads to the results stated above.

Set coloring. As $\|A^j\|_2 \leq \sqrt{m}$, we have $L = \sqrt{m}$, and as $\sum_{j \in S} a_i(j)^2 \leq |S|$, we can set $h(t) = 1$ for all $t \in [n]$. Consider (4) and suppose $p \geq 1.1$ so that $p/(p-1) = O(1)$. Then

$$\beta = \int_{t=0}^{n-2} h(n-t) \cdot (n-t)^{-1/p} dt = O(n^{1-1/p}),$$

and the first bound in (4) gives $b_0 = O(pn^{1/2}(m/n)^{1/p})$. Setting $p = \log(2m/n)$ gives Spencer's $O(\sqrt{n \log(2m/n)})$ bound.

Interestingly, the above result gives a new proof of Spencer's six-deviations result based on a direct single-phase coloring. In contrast, all the previously known proofs of this result [4, 27, 33, 18] required multiple partial coloring phases.

Komlós problem. Here $L = 1$ and the second term in (4) directly gives a $O(\sqrt{\log m})$ bound[2]. This also implies an $O(\sqrt{\log n})$ bound as at most n^2 rows can have ℓ_1-norm more than 1, and we can assume that $m \leq n^2$.

Similarly, bounding $h(t)$ using standard concentration bounds, directly gives the following results for various models of random matrices.

▶ **Theorem 2** (Sub-Gaussian Matrix). *Let $A \in \mathbb{R}^{m \times n}$ with each column drawn independently from a distribution $\mathcal{D}$, where the marginal of each coordinate is sub-Gaussian with mean 0 and variance σ^2. Then, for $n \leq m \leq 2^{O(\sqrt{n})}$, $\mathrm{disc}(A) = O(\sigma\sqrt{n \log(2m/n)})$, with probability at least $1 - (1/m^2)$.*

▶ **Theorem 3** (Random Matrix). *Let $A \in \mathbb{R}^{m \times n}$, $m \geq n$ such that every column of A is drawn independently from the uniform distribution on $\{x \in \mathbb{R}^m : \|x\|_2 \leq 1\}$. Then $\mathrm{disc}(A) = O(1)$ with probability at least $1 - (1/m^2)$.*

1.1.1 Flexibility of the method

An important advantage of the method is it flexibility, which can be used to obtain several additional results.

Subadditivity. Given $A, B \in \mathbb{R}^{m \times n}$, can we bound $\mathrm{disc}(A + B)$ given bounds on $\mathrm{disc}(A)$ and $\mathrm{disc}(B)$? Such questions can be directly handled by this framework by considering a weighted combination of two different potential functions – one for A and another for B.

More precisely, let us define $\mathrm{sdisc}(A)$, the *Stochastic Discrepancy* of a matrix A, to be the upper bound on discrepancy obtained by the Potential Walk described in Algorithm 1. For this notion, we have the following approximate subadditivity for arbitrary matrices.

▶ **Theorem 4** (Subadditivity of Stochastic Discrepancy). *For any two arbitrary matrices $A, B \in \mathbb{R}^{m \times n}$, there exists $x \in \{-1, 1\}^n$ such that*

$$|\langle a_i, x\rangle| \lesssim \mathrm{sdisc}(A) \quad \textit{for every row } a_i \textit{ of } A, \textit{ and}$$
$$|\langle b_i, x\rangle| \lesssim \mathrm{sdisc}(B) \quad \textit{for every row } b_i \textit{ of } B.$$

In particular, this implies that $\mathrm{sdisc}(A + B) \lesssim \mathrm{sdisc}(A) + \mathrm{sdisc}(B)$.

Here $a \lesssim b$ means that $a = O(1)b$. The theorem is algorithmic if A, B are given. It also implies that for any matrix A, we have $\mathrm{sdisc}(A) \lesssim \min_B(\mathrm{sdisc}(B) + \mathrm{sdisc}(A - B))$.

Similar questions have been studied previously in the context of understanding the discrepancy of unions of systems [30, 31]. For example, other related quantities such as the γ_2-norm and the determinant lower bound are also subadditive [30, 31], We remark that the additive bound cannot hold for the (actual) discrepancy or even hereditary discrepancy[3], and a logarithmic loss is necessary. For this reason, the previous additive bounds based on γ_2-norm and the determinant lower bound lose extra polylogarithmic factors when translated to discrepancy.

A direct application of Theorem 4 is the following.

2 It would be interesting to construct an explicit family of examples where the discrepancy obtained by our approach is $\Omega(\sqrt{\log n})$.

3 A classical example due to Hoffman gives two set systems A and B, each with hereditary discrepancy 1, but their union has discrepancy $\Omega(\log n/ \log\log n)$ [29].

▶ **Theorem 5** (Semi-Random Komlós). *Let $C \in \mathbb{R}^{m \times n}$ be an arbitrary matrix with columns satisfying $\left\|C^j\right\|_2 \le 1$ for all $j \in [n]$, and $R \in \mathbb{R}^{m \times n}$ be a matrix with entries drawn i.i.d. from $\mathcal{N}(0, \sigma^2)$. Then, for $n \le m \le 2^{O(\sqrt{n})}$, with probability at least $1 - (1/m^2)$,*

$$\mathrm{disc}(C+R) = O\left(\sqrt{\log n} + \sigma\sqrt{n \log(2m/n)}\right).$$

For $m = O(n)$, the bound above is $O(\sqrt{\log n} + \sigma\sqrt{n})$, which is better than the bound of $O(\sqrt{\log n}(1 + \sigma\sqrt{n}))$ obtained by directly applying the best-known bound for the Komlós problem to $C + R$.

As another application, consider a matrix C with n columns and two sets of rows, A and B, where each row in A has entries in $\{0, 1\}$, and the column norm of every column restricted to rows in B is at most 1. Suppose that A has $O(n)$ rows. Applying the framework gives a coloring with $O(\sqrt{n})$ discrepancy for rows in A and $O(\sqrt{\log n})$ for rows in B.[4] Notice that using previous techniques, if we apply the partial coloring method to get $O(\sqrt{n})$ discrepancy for A, this would give $O(\log n)$ for rows of B. On the other hand, if we apply try to obtain $O(\sqrt{\log n})$ discrepancy for B, all the known methods would incur $O(\sqrt{n \log n})$ discrepancy for A.

Relaxing the function $h(\cdot)$. Recall that the function h in Theorem 1, that controls how the ℓ_2 norms of rows decrease when restricted to subsets S of columns, and plays an important role in the bounds. In many random or pseudo-random instances however, a worst case bound on h can be quite pessimistic. For example, here even though most rows decrease significantly when restricted to S, h can remain relatively high due to a few outlier rows. The following result gives improved bound for such settings where for any subset S of columns, most row sizes restricted to S do not deviate much from their expectation if S is chosen at random.

▶ **Theorem 6** (Pseudo-Random Bounded Degree Hypergraphs). *Let $A \in \{0,1\}^{m \times n}$ such that $\left\|A^j\right\|_1 \le k$. Suppose there exists $\beta \le k$ s.t. for any $S \subseteq [n]$ and any $c > 0$, the number of rows of A with*

$$\left|\sum_{j \in S} a_i(j) - \|a_i\|_1 \cdot (|S|/n)\right| \ge c\beta \tag{5}$$

is at most $c^{-2}|S|$. Then $\mathrm{disc}(A) = O(\sqrt{k} + \beta)$.

As discussed in [32], one can set $\beta \le \max_{v \perp \mathbf{1}, \|v\|=1} \|Av\|$ in (5), which in particular gives Potukuchi's result [32] for random k-regular hypergraphs as $\beta = O(k^{1/2})$ in this case.

Combining with Theorem 4, this extends to the following semi-random setting. Consider a random k-regular hypergraph A with n vertices and n edges. Suppose an adversary can arbitrarily modify A by adding or deleting vertices from edges such that degree of any vertex changes by at most t. How much can this affect the discrepancy of the hypergraph?

▶ **Theorem 7** (Semi-Random Hypergraphs). *Consider a random k-regular hypergraph with incidence matrix $A \in \mathbb{R}^{m \times n}$ with $m \ge n$, and let $C \in \{-1, 0, 1\}^{m \times n}$ be an arbitrary matrix with at most t non-zero entries per column. Then $\mathrm{disc}(A + C) = O\left(\sqrt{k} + \sqrt{t \log n}\right)$ with probability $1 - n^{-\Omega(1)}$.*

---

[4] This answers a question of Haotian Jiang.

2 The Framework

Given a matrix $A \in \mathbb{R}^{m\times n}$, we start at some x_0 and our goal is to reach an x_T in $\{-1,1\}^n$ with small discrepancy. The basic idea will be to apply a small random update (of size δ) to x_t at step t for T steps, where the update will be chosen with care. We use the slack function and the potential function defined in (1) and (2) to implement this approach. The figure below gives a high level description of the algorithm.

Algorithm 1 PotentialWalk.

Input: A matrix $A \in R^{m\times n}$, a potential function $\Phi : \mathbb{R} \times \mathbb{R}^n \to \mathbb{R}^+$.
Let $x_0 = 0, t = 0$. Let $T = (n-2)/\delta^2$.
for $t \in [T]$ **do**
  Select v_t such that: (i) $\mathbb{E}_\varepsilon[\Phi(t+1, x_t + \varepsilon\delta v_t)] \leq \Phi(t, x_t)$, (ii) $x_t \pm \delta v_t \in [-1,1]^n$, and (iii) $\langle x_t, v_t\rangle = 0$, where ε is a Rademacher random variable (± 1 with probability $1/2$).
  Let $x_{t+1} = x_t + \varepsilon\delta v_t$.
Output: x_T

2.1 Example: Komlós setting

We first give an overview of the ideas by describing how the framework above works for the Komlós setting. Recall that here $A \in \mathbb{R}^{m\times n}$ has columns satisfying $\left\|A^j\right\|_2 \leq 1$. To minimize notation, let us assume here that $m = n$ (this is also the hardest case for the problem).

At time t, let $\mathcal{V}_t = \{j \in [n] : |x_t(j)| < 1 - 1/2n\}$ and let $n_t = |\mathcal{V}_t|$. These are the variables that are "alive", and not yet "frozen". To ensure that $x_t \in [-1,1]^n$, the update v_t will only change the variables in $\mathcal{V}_t$. We also set $\langle v_t, x_t\rangle = 0$, which ensures that $\|x_t\|^2 = \delta^2 t$ for any $t \in [0,T]$. So v_t satisfies

$$v_t(j) = 0 \text{ for all } j \notin \mathcal{V}_t \text{ and } \langle v_t, x_t\rangle = 0. \tag{6}$$

As $|x_t(j)| \geq (1-1/2n)$ for all $j \notin \mathcal{V}_t$, we have $(n-n_t)(1-1/2n)^2 \leq \sum_{j\notin\mathcal{V}_t} x_t(j)^2 \leq \sum_{j\in[n]} x_t(j)^2 = \delta^2 t$. So the number of alive variables at time t satisfies $n_t \geq n - (\delta^2 t)/(1-(1/(2n)))^2 > n - \delta^2 t - 1$.

Blocking large rows. To ensure the two-sided bound $|\sum_j a_i(j)x(j)| < b_0$, we create a new row $-a_i$ for each row a_i at the beginning. Now, as the squared 2-norm of every column of A is at most 2, at any time t, the number of rows with $\sum_{j\in\mathcal{V}_t} a_i(j)^2 > 12$ is at most $|\mathcal{V}_t|/6 = n_t/6$. Let us call such rows *large* (at time t). Otherwise, the row is *small.* We additionally constrain v_t so that

$$\langle a_i, v_t\rangle = 0 \text{ for all rows } \{i : \sum_{j\in\mathcal{V}_t} a_i(j)^2 > 12\}. \tag{7}$$

This ensures that a row only starts to incur any discrepancy once it becomes small. So at step t, we will define the slacks only for small rows and only such rows will contribute to the potential $\Phi(t)$. Let $\mathcal{I}_t$ denote the set of small rows at time t. In the slack function (1), we will set $b_t = b_0$ for all t and $\lambda = 2^{-5}b_0$. So, at the beginning of the algorithm, when $x_0(j) = 0$ for all j, we have $\Phi(0) = \sum_{i\in\mathcal{I}_0}(b_0 - \lambda\cdot\sum_{j\in[n]} a_i(j)^2)^{-p} \leq \frac{|\mathcal{I}_0|}{(b_0-12\lambda)^p} \leq n\left(\frac{2}{b_0}\right)^p$.

At any time t, the change in potential $\Phi(t+1)-\Phi(t)$ is due to (i) new rows becoming small and entering $\mathcal{I}_{t+1}$ and (ii) and the change slack of rows in $\mathcal{I}_t$. As each row has discrepancy 0 until it becomes small, the total contribution of step (i) over the entire algorithm is at most $n(2/b_0)^p$. So the main goal will be to show that Φ does not rise due to step (ii). This will ensure that the potential throughout the algorithm is at most $2n(2/b_0)^p$, which gives the $\sum_j a_i(j)x(j) < b_0$ for all i.

Bounding the increase in Φ. We now describe the main ideas of the algorithm and computations for the change in Φ in step (ii). The desired $O(\sqrt{\log n})$ will then follow directly by optimizing the parameters b_0 and p in (1).

Let $e_{t,i}$ denote a vector in $\mathbb{R}^n$ with j-th entry $a_i(j)^2 x_t(j)$. At step t, x_t changes as $x_{t+1} - x_t = \varepsilon\delta \cdot v_t$ and, by a simple calculation, the approximate change in $s_i(t)$ is:

$$s_i(t+1) - s_i(t) \simeq \left(2\lambda\langle e_{t,i}, v_t\rangle - \langle a_i, v_t\rangle\right)\varepsilon\delta + \lambda\langle a_i^{(2)}, v_t^{(2)}\rangle\delta^2\,,$$

where ε is a Rademacher random variable and $a^{(2)}$ denotes the vector with j-th entry $a(j)^2$. The error terms not included above are all higher powers of δ, and can be ignored for small enough δ as long as all coefficients are bounded. We formalize this in Section 2.2.

Then, up to second order terms in δ, $\Phi(t+1) - \Phi(t) \simeq f(t)\delta^2 + g(t)\varepsilon\delta$ where,

$$f(t) = -p\lambda \sum_{i\in\mathcal{I}} \frac{\langle a_i^{(2)}, v_t^{(2)}\rangle}{s_i(t)^{p+1}} + \frac{p(p+1)}{2}\sum_{i\in\mathcal{I}} \frac{\left(2\lambda\langle e_{t,i}, v_t\rangle - \langle a_i, v_t\rangle\right)^2}{s_i(t)^{p+2}},$$

$$g(t) = p\sum_{i\in\mathcal{I}} \frac{\left(2\lambda\langle e_{t,i}, v_t\rangle - \langle a_i, v_t\rangle\right)}{s_i(t)^{p+1}}.$$

Note that the expectation of the second term $g(t)\varepsilon\delta$ is zero. So it suffices to prove that there is a choice of v_t such that $f(t) \leq 0$. This will ensure the expected change of Φ is at most *zero*, and there will be a choice of ϵ that ensures Φ is non-increasing. The difficulty in making $f(t)$ at most zero is that the positive part (the second term of $f(t)$) has an extra factor of $s_i(t)$ in the denominator. So if some $s_i(t)$ becomes very small, the positive term could dominate. To ensure this doesn't happen, we choose v_t to be in a subspace that makes this positive term zero for the smallest slack indices.

Blocking small slacks. Let $\mathcal{J}_t$ be the subset of $\mathcal{I}$ corresponding to all but the $\lfloor n_t/12\rfloor$ smallest values of $s_i(t)$ at time t. Select v_t such that

$$\left(2\lambda\langle e_{t,i}, v_t\rangle - \langle a_i, v_t\rangle\right) = 0 \text{ for all } i \in \mathcal{I}\backslash\mathcal{J}_t, \tag{8}$$

Then as $\sum_i s_i(t)^{-p}) \leq \Phi(t)$, and the smallest $n_t/12$ slacks are "blocked", we have

$$\max_{j\in\mathcal{J}_t} \frac{1}{s_j(t)} \leq \left(\frac{\Phi(t)}{n_t/12}\right)^{1/p},$$

and so,

$$\begin{aligned} f(t) &\leq p\left(\frac{p+1}{2}\sum_{i\in\mathcal{J}_t} \frac{\left(2\lambda\langle e_{t,i}, v_t\rangle - \langle a_i, v_t\rangle\right)^2}{s_i(t)^{p+1}} \max_{j\in\mathcal{J}_t} s_j(t)^{-1} - \lambda\sum_{i\in\mathcal{I}} \frac{\langle a_i^{(2)}, v_t^{(2)}\rangle}{s_i(t)^{p+1}}\right) \\ &\leq p\left(\frac{p+1}{2}\sum_{i\in\mathcal{J}_t} \frac{\left(2\lambda\langle e_{t,i}, v_t\rangle - \langle a_i, v_t\rangle\right)^2}{s_i(t)^{p+1}} \left(\frac{12\Phi(t)}{n_t}\right)^{1/p} - \lambda\sum_{i\in\mathcal{I}} \frac{\langle a_i^{(2)}, v_t^{(2)}\rangle}{s_i(t)^{p+1}}\right)\end{aligned}$$

In addition to (6) and (8), suppose v_t also satisfies

$$\sum_{i\in\mathcal{J}_t} \frac{\langle 2\lambda e_{t,i} - a_i, v_t\rangle^2}{s_i(t)^{p+1}} \leq 12 \cdot \sum_{i\in\mathcal{J}_t} \frac{\langle a_i^{(2)}, v_t{}^{(2)}\rangle}{s_i(t)^{p+1}}. \tag{9}$$

Choosing the update v_t. Later in Section 2.2, we will see how to find a vector v_t satisfying (6), (8), (7), and (9). Then,

$$f(t) \leq p \sum_{i\in\mathcal{J}_t} \frac{\langle a_i^{(2)}, v_t^{(2)}\rangle}{s_i(t)^{p+1}} \left(6(p+1)\left(\frac{12\Phi(t)}{n_t}\right)^{1/p} - \lambda\right).$$

To show that $f(t) \leq 0$, it thus suffices to have $6(p+1)\left(12\Phi(t)/n_t\right)^{1/p} - \lambda \leq 0$.

As $\Phi(t)^{\frac{1}{p}} \leq 2(2n)^{1/p}/b_0$ by the inductive hypothesis, and $n_t \geq 1$, it suffices to have $12(p+1)\left(24n\right)^{1/p} - \lambda \cdot b_0 \leq 0$. Choosing $p = \log n$ so that $n^{1/p} = O(1)$, and as $\lambda = 2^{-5}b_0$, we can pick $b_0 = O(\sqrt{\log n})$ to satisfy the above. This gives the desired discrepancy bound.

2.2 The General Framework

We now describe the algorithm more formally. Given a matrix $A \in \mathbb{R}^{m\times n}$ with $\left\|A^j\right\|_2 \leq 1$ for all $j \in [n]$, extend A such that for each original row a_i of A, there are two rows a_i and $-a_i$ in A. Additionally, partition every row a_i into 2 rows, a_i^S and a_i^L, with small and large entries, as follows:

$$a_i^S(j) = \begin{cases} 0 & \text{if } |a_i(j)| > 1/2\lambda \\ a_i(j) & \text{otherwise} \end{cases}, \quad a_i^L(j) = \begin{cases} a_i(j) & \text{if } |a_i(j)| > 1/2\lambda \\ 0 & \text{otherwise,} \end{cases}$$

where λ is a parameter to be determined later. After this transformation, for any $x \in \mathbb{R}^n$, $\|Ax\|_\infty = \max_i \langle a_i^S + a_i^L, x\rangle$, and the squared 2-norm of any column of A is at most 2.

Let $\mathcal{I}$ denote the index set of all rows of A, and $\mathcal{I}^S$ denote the index set of rows of the first type above.

The step-size of the algorithm is δ and the algorithm will run for $T = \frac{n-2}{\delta^2}$ steps. Starting with $x_0 = 0$, let $v_t \in \mathbb{R}^n$ with $\langle x_t, v_t\rangle = 0$. For $t \in [T]$,

$$x_t = \begin{cases} x_{t-1} + \delta v_{t-1} & \text{w.p. } 1/2, \\ x_{t-1} - \delta v_{t-1} & \text{w.p. } 1/2. \end{cases}$$

As t increases, some variables will start approaching 1 in magnitude. To ensure that $x_t \in [-1,1]^n$, we restrict v_t to be in the space of *alive* variables, defined as $\mathcal{V}_t = \{i \in [n] : |x_t(i)| < 1 - 1/(2n)$.

For any $t \in [T]$, $\|x_t\|^2 = \delta^2 t$ as

$$\|x_t\|^2 = \|x_{t-1} + \delta v_t\|^2 = \|x_{t-1}\|^2 + \delta^2\|v_t\|^2 = \delta^2(t-1) + \delta^2 = \delta^2 t. \tag{10}$$

Let $n_t = |\mathcal{V}_t|$ denote the number of alive variables at t. By (10), $(n-n_t)(1-\epsilon)^2 \leq \delta^2 t$, which gives $n_t \geq n - \frac{\delta^2 t}{(1-1/(2n))^2} > n - \delta^2 t - 1$.

To select a v_t such that for all $t \in [T]$, $x_t \in [-1,1]^n$ and $\langle a_i, x_t\rangle$ is bounded for all rows, we classify the rows according to how many variables are still "uncolored" in a row.

Let the set of *s-Alive* rows at time t be defined as $\mathcal{I}_t = \{i \in \mathcal{I}^S : \sum_{j\in\mathcal{V}_t} a_i(j)^2 \leq 20\}$. The choice of 20 here is arbitrary, and large enough constant works.

We can now define the slack and the potential function.

Slack. For any $i \in \mathcal{I}$, the slack function is defined as

$$s_i(t) = b_t - \langle a_i, x_t \rangle - \lambda \cdot \sum_{j=1}^{n} a_i(j)^2(1 - x_t(j)^2).$$

We call b_t the barrier, and for $t \in [T]$, we also move it as $b_t = b_{t-1} + \delta^2 d_{t-1}$, for some function d_t. We set $\lambda = cb_0$ where $c = 1/42$ and b_0 is the initial barrier.

Potential function. The potential function has a parameter $p > 1$ and is defined as

$$\Phi(t) = \sum_{i \in \mathcal{I}_t} s_i(t)^{-p}.$$

We will only consider slacks for *alive* rows and ensure that they are always positive. Moreover, we will consider only the small *s-Alive* rows as the rows in $\mathcal{I}^L$ will be easily handled. To ensure that $s_i(t)$ does not become too "small" for any *s-Alive* row, the choice of v_t should not decrease the smallest slacks. This motivates the following definitions.

- *Blocked* rows: Let $\mathcal{C}_t$ be the subset of $\mathcal{I}_t$ corresponding to the $\lfloor n_t/12 \rfloor$ smallest values of $s_i(t)$.
- Let $\mathcal{J}_t = \mathcal{I}_t \backslash \mathcal{C}_t$. These are the "large slack" rows.

To prove that all the slacks are positive, we will upper bound the potential throughout by bounding the change in $\Phi(t)$ at each step. Note that $\Phi(t)$ will experience jumps whenever a new index gets added to $\mathcal{I}_t$, however the total contribution of jumps is easily shown to be bounded (see Lemma 19) and can essentially be ignored. To bound the one-step change in Φ, we use the second order Taylor expansion of $\Phi(t+1)$ centered at $\Phi(t)$. Details of this can be found in the arXiv version of this paper [8].

2.3 Algorithm and Analysis

Recall that $e_{t,i}$ denotes the vector in $\mathbb{R}^n$ with j-th entry $a_i(j)^2 x_t(j)$. We can now state the algorithm for selecting v_t.

Algorithm 2 Algorithm for Selecting v_t.

Initialize $x_0 \leftarrow 0$

for $t = 1, \ldots, T = \frac{n-2}{\delta^2}$ **do**

  Let $\mathcal{W}_t = \{w \in \mathbb{R}^n : w(i) = 0,\ \forall i \notin \mathcal{V}_t\}$ `// restrict to alive variables`

  Let $\mathcal{U}_t = \{w \in \mathcal{W}_t : \langle w, 2\lambda e_{t,i} - a_i \rangle = 0, \forall i \in \mathcal{C}_t$ and $\langle w, x_t \rangle = 0\}$
`// restrict to large slack rows`

  Let $\mathcal{Y}_t = \{w \in \mathcal{W}_t : \langle w, a_i \rangle = 0, \forall i \in \mathcal{I} \backslash \mathcal{I}_t\}$ `// restricted to s-Alive rows`

  Let $\mathcal{G}_t$ denote the subspace

$$\mathcal{G}_t = \left\{ w \in \mathcal{W}_t : \sum_{i \in \mathcal{J}_t} \frac{\langle (2\lambda e_{t,i} - a_i), w \rangle^2}{s_i(t)^{p+1}} \leq 40 \sum_{i \in \mathcal{J}_t} \frac{\langle a_i^{(2)}, w^{(2)} \rangle}{s_i(t)^{p+1}} \right\} \tag{11}$$

  Consider the subspace $\mathcal{Z}_t = \mathcal{U}_t \cap \mathcal{Y}_t \cap \mathcal{G}_t$ and let $W = \{w_1, w_2, \ldots, w_k\}$ be an orthonormal basis for $\mathcal{Z}_t$. Choose

$$v_t = \arg\min_{w \in W} \sum_{i \in \mathcal{J}_t} \langle 2\lambda e_{t,i} - a_i, w \rangle^2 s_i(t)^{-(p+1)}. \tag{12}$$

We now re-state our main theorem. In words, the assumption of the theorem is that there is a non-decreasing function $h(.)$ such that for any row, the squared norm in any subset of coordinates S is proportional to $h(|S|)$ times the size of the subset S. Under this condition, we can bound the discrepancy as a function of h.

▶ **Theorem 1.** *For a $A \in \mathbb{R}^{m\times n}$ with $\left\|A^j\right\|_2 \le L$ and $|a_i(j)| \le M$ for all $i \in [m], j \in [n]$, let $h : \mathbb{R}^+ \to \mathbb{R}^+$ be a non-increasing function such that for every subset $S \subseteq [n]$ and $i \in [m]$,*

$$\sum_{j\in S} a_i(j)^2 \le |S| \cdot h(|S|). \tag{3}$$

Then, for any $p > 1$, there exists a vector $x \in \{-1, 1\}^n$ such that $\|Ax\|_\infty \le 5b_0 + 2M$, where

$$b_0 = \min\left(\sqrt{8(p+1)(48m)^{1/p}\cdot\beta},\ 250L\sqrt{\log{(2m)}}\right). \tag{4}$$

where $\beta = \int_{t=0}^{n-2} h(n-t)(n-t)^{-1/p}dt$.

The case when $h(t) = h$ is often useful, for which case we have following corollary.

▶ **Corollary 8.** *For a matrix $A \in \mathbb{R}^{m\times n}$ with $\|A^j\| \le L$ and $|a_i(j)| \le M$ for all $i \in [n], j \in [m]$, let h be such that for every subset $S \subseteq [n]$ and every $i \in [m]$, $\sum_{j\in S} a_i(j)^2 \le |S| \cdot h$. Then, $\mathrm{disc}(A) \le 5b_0 + 2M$, where $b_0 = \min(26\sqrt{hn\log(2m/n)},\ 250L\sqrt{\log{(2m)}})$.*

Roadmap of the proof. The first main lemma below (Lemma 10) establishes that there is a large feasible subspace from which v_t as defined above can be chosen. Using this we prove Lemma 11, which bounds the change in potential. This will allow us to bound the discrepancy of each row and hence prove Theorem 1.

A key fact used for proving Lemma 10 is the following lemma in [7].

▶ **Lemma 9** ([7])**.** *Let $G, H \in \mathbb{R}^{m\times n}$ be matrices such that $|G_{ij}| \le \alpha|H_{ij}|$ for all $i \in [m]$ and $j \in [n]$. Let $K = diag(H^\top H)$. Then for any $\beta \in (0, 1]$, there exists a subspace $W \subseteq \mathbb{R}^n$ such that $\dim(W) \ge (1-\beta)n$, and $\forall w \in W$, $w^\top G^\top Gw \le \frac{\alpha^2}{\beta} \cdot w^\top Kw$.*

We now arrive at the main Lemma.

▶ **Lemma 10** (Subspace Dimension)**.** *For all $t \in T$, $\dim(\mathcal{Z}_t) \ge \lceil 2n_t/3\rceil$.*

Setting the parameters. To show the two bounds in (4), we will set the parameters b_t, d_t (the change in b_t) and p in two ways:

$$\textit{Case 1: } d_t = 4(p+1)\cdot h(n_t)\cdot \max_{i\in\mathcal{J}_t} s_i(t)^{-1} \text{ for all } t \in [T], \text{ and } p, b_0 \text{ arbitrary} \tag{13}$$

$$\textit{Case 2: } p = 2\log(2m),\ b_0 = 840(p+1)\cdot \max_{j\in\mathcal{J}_t} s_j(t)^{-1} \text{ and } d_t = 0 \text{ for all } t \in [T]. \tag{14}$$

Bounding the potential. The next lemma shows that in both these cases, the potential function remains bounded.

▶ **Lemma 11** (Bounded Potential)**.** *In either of the cases given by (13) and(14), we have that $\Phi(t) \le 4m(2/b_0)^p$, for all $t = 0, \ldots, T$.*

The next lemma gives a bound on the minimum value of slack for any active row, given the bound on potential function.

▶ **Lemma 12.** *For any $t \in \{0, \ldots, T\}$, if $\Phi(t) \le 4m(2/b_0)^p$, then* $\max_{i \in \mathcal{J}_t} s_i(t)^{-1} \le \frac{2}{b_0}\left(\frac{48m}{n_t}\right)^{\frac{1}{p}}$.

▶ **Lemma 13.** *For any $t \in [T]$, the choice of v_t satisfies*

$$\sum_{i \in \mathcal{J}_t} \frac{\langle 2\lambda e_{t,i} - a_i, v_t\rangle^2}{s_i(t)^{p+1}} \le \sum_{i \in \mathcal{J}_t} \frac{8h(n_t)}{s_i(t)^{p+1}}. \tag{15}$$

These lemmas will allow us to prove the main theorem (see Appendix).

3 Applications

3.1 Set Coloring

We bound the discrepancy of a set system $(U, \mathcal{S})$ with $|U| = n$, $|\mathcal{S}| = m$, and $m \ge n$. As $\|A^j\|_2 \le \sqrt{m}$, we have $L = \sqrt{m}$, and as $\sum_{j \in S} a_i(j)^2 \le |S|$, we can set $h(t) = 1$ for all $t \in [n]$. Consider (4) and suppose $p \ge 1.1$ so that $p/(p-1) = O(1)$. Then

$$\beta = \int_{t=0}^{n-2} h(n-t) \cdot (n-t)^{-1/p} dt = O(n^{1-1/p}),$$

and the first bound in (4) gives $b_0 = O(pn^{1/2}(m/n)^{1/p})$. Setting $p = \log(2m/n)$ gives Spencer's $O(\sqrt{n\log(2m/n)})$ bound.

3.2 Vector Balancing

We now consider the discrepancy a matrix $A \in \mathbb{R}^{m \times n}$ with column ℓ_2-norms at most 1.

Here $L = 1$ and the second term in (4) directly gives a $O(\sqrt{\log m})$ bound. This also implies an $O(\sqrt{\log n})$ bound as at most n^2 rows can have ℓ_1-norm more than 1, and we can assume that $m \le n^2$. In particular, for a row a_i with $\|a_i\|_2 < 1/n^{1/2}$, we have $|\langle a_i, x\rangle| \le \|a_i\|_1 \le \sqrt{n}\,\|a_i\|_2 < 1$ and it can be ignored. The sum of squares of elements in A is at most n the number of rows with $\|a_i\|_2 > 1/n^{1/2}$ is at most n^2.

3.3 Sub-Gaussian Matrices and Random Matrices

We give the proofs for these application in the appendix.

4 Flexibility of the Method

An advantage of the potential function approach is its flexibility. We describe two illustrative applications. In Section A.2 we show how the bounds for matrices A and B obtained using the framework can be used to directly give bounds for $C = A + B$ by combining the potentials for A and B in a natural way.

In Section 4.1 we consider how the requirement on the function $h(\cdot)$ in Theorem 1 can be relaxed, and use it to bound the discrepancy of sparse hypergraphs (the Beck-Fiala setting) satisfying a certain pseudo-randomness condition.

4.1 Discrepancy of Sparse Pseudo-random Hypergraphs

In this section, we consider 0/1 matrices that satisfy a certain regularity property, namely, for most rows, the sum of their entries in any subset of columns is close to the sum of the full row scaled by the fraction of columns in the subset. This property is satisfied, e.g., by the matrices that correspond to sparse random hypergraphs. In particular, we show the following.

▶ **Theorem 6** (Pseudo-Random Bounded Degree Hypergraphs). *Let $A \in \{0,1\}^{m\times n}$ such that $\|A^j\|_1 \leq k$. Suppose there exists $\beta \leq k$ s.t. for any $S \subseteq [n]$ and any $c > 0$, the number of rows of A with*

$$\Big|\sum_{j\in S} a_i(j) - \|a_i\|_1 \cdot (|S|/n)\Big| \geq c\beta \tag{5}$$

is at most $c^{-2}|S|$. Then $\mathrm{disc}(A) = O(\sqrt{k}+\beta)$.

Proof outline. At a high level the proof is similar to that of Theorem 4, using a weighted potential function. However, rather than just two potentials, we will have to consider a combination of $O(\log n)$ potentials, and it will take some care to make sure this doesn't create an overhead in the discrepancy. We note that the main algorithm remains: at each step choose a vector in a subspace defined by a set of constraints based on the current vector x_t.

We next discuss the details of the algorithm and the proof of Theorem 6. The full proof can be found in the arXiv version of this paper [8].

Partitioning rows according to ℓ_1-norm. First, extend A such that for each original row a_i, there are two rows a_i and $-a_i$ in A. Since our goal is to prove discrepancy $O(\sqrt{k})$, we can ignore all rows will ℓ_1-norm less than $\sqrt{k}$. Then $m \leq n\sqrt{k}$ because the number of rows with ℓ_1-norm greater than $\sqrt{k}$ is at most $2nk/\sqrt{k} = 2n\sqrt{k}$. Let $N = \lceil \log_2 n/k \rceil$ and $\mathcal{Q} = \{0\} \cup [N]$. Partition the rows of A into based on their initial ℓ_1-norm into $|\mathcal{Q}| = N+1$ classes:

- $\mathcal{A}_0 = \{i \in \mathcal{I} : \sqrt{k} \leq \|a_i\|_1 < 2k\}$.
- For each $i \in [N]$, let $\mathcal{A}_i = \{i \in \mathcal{I} : 2^i k \leq \|a_i\|_1 < 2^{i+1}k\}$.

The sum of ℓ_1-norms of rows in A is at most $2nk$, therefore for any i, $2^i k|\mathcal{A}_i| \leq 2nk$ and $|\mathcal{A}_i| \leq 2^{1-i}n$.

We create $N+1$ potential functions $\{\Phi_i(t)\}_{i=0}^N$, one associated with each row partition. The potential functions use the same p, b_0 parameters, and $\lambda = cb_0$ with $c = 1/42$, but have different rate of change of barrier functions $d_q(\cdot)$, based on q. We will run Algorithm 2 on each partition separately but use the same x_t and v_t at each step. In this case, we can select parameters to ensure that each potential function is decreasing in expectation (see Lemma 18). However, there might not exist a vector v_t that ensure that moving in v_t direction decreases all the potential functions simultaneously. To deal with this, we use a weighted combination of Φ_q as the potential function:

$$\Phi(t) = \frac{1}{k}\cdot\Phi_0(t) + \sum_{q\geq 1} 2^{2q}\cdot\Phi_q(t). \tag{16}$$

4.1.1 A suitable subspace

To identify the constrained subspace for the PotentialWalk (Algorithm 6), we use the following definitions. The set of *Active* rows is defined as $\mathcal{I}_t = \{i \in \mathcal{I} : \sum_{j\in\mathcal{V}_t} |a_i(j)| \le 12k\}$. For each class q, let $h_q : \mathbb{R}^+ \to \mathbb{R}$ be a non-increasing function such that for every subset $S \subseteq n$, at most $n_t/16$ rows i from class $\mathcal{A}_q$ violate the condition

$$\sum_{j\in S} |a_i(j)| \le |S| \cdot h_q(|S|). \tag{17}$$

While following the general framework from Section 2.2, we make three crucial changes:

- Move orthogonal to rows with *large deviation*. At step t, the ℓ_1 norm of row a_i will be close to $(n_t/n) \cdot \|a_i\|_1$ for most rows. Let $a_{i,t}$ denote a vector in $\mathbb{R}^n$ with j-th entry $\mathbf{1}_{j\in\mathcal{V}_t} a_i(j)$, i.e., $a_{i,t}$ is row a_i restricted to the alive coordinates at time t. Then the set of *large deviation* rows consists of rows that deviate significantly from this expected value

  $$\mathcal{B}_t = \{i \in \mathcal{I} : |\, \|a_{i,t}\|_1 - \|a_i\|_1 \cdot (n_t/n)| \ge 4\beta\}. \tag{18}$$

  For any $t \in [T]$, (5) implies that $\dim(\mathcal{B}_t) \le \lfloor n_t/16 \rfloor$.
- Ignore *Dead* rows. As soon as the ℓ_1-norm of some row becomes less than 8β, we drop it from the potential function. The set of *dead* rows at step t is defined as

  $$\mathcal{D}_t = \{i \in \mathcal{I} : \|a_{i,t}\|_1 \le 8\beta\}. \tag{19}$$

  For a dead row, rather than keeping track of its discrepancy using a slack function, we uniformly bound the the additional discrepancy gained by a row after it becomes dead.
- *Block* rows based on their initial size. For $q \in \mathcal{Q}$, let $\mathcal{C}_t^q$ be the subset of $\mathcal{A}_q \cap \mathcal{I}_t$ corresponding to the $\lfloor 2^{i-8} n_t^2/n \rfloor$ smallest values of $\{s_i(t) : i \in \mathcal{A}_q \cap \mathcal{I}_t\}$, and let $\mathcal{J}_t^q = \mathcal{A}_i \backslash \{\mathcal{C}_t^q \cup \mathcal{D}_t\}$.

We are ready to state the algorithm for selecting v_t.

Algorithm 3 Algorithm for Selecting v_t.

Let $h_q(n_t) = 2^{q+2}/n$ and $w_q(t) = 2^{5-\frac{q}{4}} \left(\frac{n}{n_t}\right)^{1/4}$
for $t = 1, \ldots, T$ **do**
  Let $\mathcal{W}_t = \{w \in \mathbb{R}^n : w(i) = 0,\ \forall i \in \mathcal{V}_t\}$ `// restrict to alive variables`
  Let $\mathcal{U}_t = \{w \in \mathcal{W}_t : \langle w, 2cb_0 e_{t,i} - a_i\rangle = 0, \forall i \in \mathcal{C}_t$ and $\langle w, x_t\rangle = 0\}$ `// restrict to large slack rows`
  Let $\mathcal{Y}_t = \{w \in \mathcal{W}_t : \langle w, a_i\rangle = 0, \forall i \in \mathcal{I}\backslash\mathcal{I}_t\}$ `// move orthogonal to large norm rows`
  Let $\mathcal{G}_t = \{w \in \mathcal{W}_t : \langle a_i, w\rangle = 0,\ \forall i \in \mathcal{B}_t\}$ `// move orthogonal to large deviation rows`
  Let $\mathcal{Z}_t = \mathcal{U}_t \cap \mathcal{Y}_t \cap \mathcal{G}_t$ and let $W = \{w_1, \ldots, w_k\}$ be an orthonormal basis for $\mathcal{Z}_t$
  Let $v_t \in W$ such that for all $q \in \mathcal{Q}$,

$$\sum_{i\in\mathcal{J}_t^q} \langle 2cb_0 e_{t,i} - a_i, v_t\rangle^2 s_i(t)^{-p-1} \le 8 w_q(t) \cdot h_q(n_t) \sum_{i\in\mathcal{J}_t^q} s_i(t)^{-p-1}. \tag{20}$$

4.2 Proof Outline

The following lemma bounds the number of active classes at step t.

▶ **Lemma 14.** *At step t, the following two conditions hold: (i) The number of classes q for which $\mathcal{A}_q \cap \{\mathcal{I}_t \backslash \{\mathcal{B}_t \cup \mathcal{D}_t\}\} \neq \emptyset$ is at most $\log(16n/n_t)$ and (ii) $h_q(t) = 2^{q+2}k/n$ satisfies (17) for all $q \in \mathcal{Q}$.*

So at any step t, the set of active rows is from the first $\log_2(16n/n_t)$ classes of rows. It also helps us define two important parameters associated with a row class q. At step t, consider a $q \in \mathcal{Q}$ with $\mathcal{A}_q \cap \{\mathcal{I}_t \backslash \{\mathcal{B}_t \cup \mathcal{D}_t\}\} \neq \emptyset$.

- Since $n - \delta^2 t - 1 < n_t \leq 16 \cdot 2^{-q} n$, for $q \geq 1$, let $t_q := \max\left\{0,\ n\delta^{-2}\left(1 - 16 \cdot 2^{-q} - 1/n\right)\right\}$. Similarly, let $t_0 := n\delta^{-2}\left(1 - 16k^{-1/2} - 1/n\right)$.
- On the other hand, q must satisfy $2^q \leq \frac{16n}{n_t}$. Let $q_t := \arg\max_{i \geq 0}\left\{2^i \leq 16 \cdot (n/n_t)\right\}$.

The next two lemmas are analogous to Lemma 10 and Lemma 13, respectively.

▶ **Lemma 15.** *For any $t \in [T]$, it holds that $\dim(\mathcal{Z}_t) \geq \lceil n_t/2 \rceil$.*

▶ **Lemma 16.** *For all $t \in [T]$, there exists $v_t \in \mathcal{Z}_t$ such that $\forall q \in \mathcal{Q}$,*

$$\sum_{i \in \mathcal{J}_t^q} \langle 2cb_0 e_{t,i} - a_i, v_t \rangle^2 s_i(t)^{-p-1} \leq 8 w_q(t) \cdot h_q(n_t) \sum_{i \in \mathcal{J}_t^q} s_i(t)^{-p-1}\,. \tag{21}$$

Note that for any row $i \in \mathcal{A}_q$, at $t \leq t_q$, $\langle 2cb_0 e_{i,t} - a_i, v_t \rangle = 0$. So, we can set $d_t^q = 0$ for rows in class q. Lemma 11 and Lemma 16 imply that for all the potential functions to be decreasing, it suffices to have

$$d^q(t) = \begin{cases} 0 & \text{if } t \leq t_q \\ 4(p+1) \cdot w_q(t) \cdot h_q(n_t) \cdot \max_{i \in \mathcal{J}_t^q} s_i(t)^{-1} & \text{otherwise.} \end{cases} \tag{22}$$

The next lemma helps us bound the rate of change of $b_q(t)$, which eventually gives a bound on $b_q(T)$ in Theorem 6.

▶ **Lemma 17.** *For any $t \in \{0, \ldots, T\}$, if $\Phi(t) \leq 8n\left(\frac{2}{b_0}\right)^p \left(\frac{16n}{n_t}\right)$, then*

$$\max_{j \in \mathcal{J}_t^q} s_j(t)^{-1} \leq \begin{cases} k^{1/p} \cdot \frac{2^{1+15/p}}{b_0}\left(\frac{n}{n_t}\right)^{3/p} & \text{if } q = 0 \\ \frac{2^{1+(15-3q)/p}}{b_0}\left(\frac{n}{n_t}\right)^{3/p} & \text{if } q \geq 1. \end{cases} \tag{23}$$

▶ **Lemma 18.** *For $p = 8$ and d_q given by (22), for all $t \in [T]$, we have $\Phi(t) \leq \frac{2^7 n^2}{n_t} \cdot \left(\frac{2}{b_0}\right)^p$.*

Proof of Theorem 6. If row $i \in \mathcal{A}_q$ becomes dead after step $t-1$, then

$$|\langle a_i, x_T \rangle| \leq |\langle a_i, x_t \rangle| + |\langle a_i^S, x_T - x_t \rangle| \leq b_t(q) + 2\sum_{j \in \mathcal{V}_t} |a_i(j)| \leq b_T(q) + 16\beta.$$

Substituting the bound on $\max_{i \in \mathcal{J}_q^t} s_i(t)^{-1}$ from (23), and using $w_q(t) = 2^{5-q/4} \cdot (n/n_t)^{1/4}$ and $h_q(t) = 2^{q+2}/n$, equation (22) gives $d_q(t) = 0$ for $t < t_q$, and

$$d_q(t) = \begin{cases} 9k \cdot \frac{2^{3q/8+14}}{nb_0}\left(\frac{n}{n-\delta^2 t-1}\right)^{5/8} & \text{if } q \geq 1 \text{ and } t \geq t_q \\ 9k^{\frac{9}{8}} \cdot \frac{2^{14}}{nb_0}\left(\frac{n}{n-\delta^2 t-1}\right)^{5/8} & \text{if } q = 0 \text{ and } t \geq t_0. \end{cases}$$

For any $q \geq 1$, summing up $d_q(\cdot)$,

$$b_q(T) = b_0 + \delta^2 \sum_{t=t_q}^{T-1} d_q(0) \leq \delta^2 \int_{t=t_q}^{T} \frac{9k \cdot 2^{3q/4+12+(15-3q)/8}}{nb_0} \left(\frac{n}{n-\delta^2 t - 1}\right)^{5/8} dt$$

$$\leq b_0 + \int_{t=\delta^2 t_q}^{n-2} \frac{9k \cdot 2^{3q/8+14}}{nb_0} \left(\frac{n}{n-t-1}\right)^{5/8} dt = b_0 + \frac{2^{20}k}{b_0}.$$

For $b_0 = 2^{10}\sqrt{k}$, $b_q(T) \leq 2^{11}\sqrt{k}$ for all $q \geq 1$. Similar calculation for $q = 0$ show that $b_0 = 2^{10}\sqrt{k}$ and $b_T(0) = 2^{11}\sqrt{k}$ suffice.

Let $x \in \{-1, 1\}^n$ be obtained from x_T by the rounding $x(j) = \text{sign}(x_T(j))$. Since $T = (n-2)/\delta^2$, $\|x_T\|^2 = n - 2$ with $|x_T(j)| \leq 1$ for all $j \in [n]$. After rounding x_T to x, $\|x\|^2 = n$ and $|\langle a_i, x\rangle| \leq |\langle a_i, x_T\rangle| + |\langle a_i, x - x_T\rangle| \leq 2b_T + 16\beta + \sum_j |x(j) - x_T(j)| \leq b_T + 16\beta + 2.$ ◀

Random and Semi-random Sparse Hypergraphs. This gives an alternate proof of the result [32] of Potukuchi that $\text{disc}(\mathcal{H}) = O(\sqrt{k})$ for regular random k-regular hypergraph $\mathcal{H}$, on n vertices and m edges with $m \geq n$ and $k = o(m^{1/2})$. In particular, Potukuchi showed that such hypergraphs satisfy condition (5) with high probability.

Proof of Theorem 7. By the subadditive property of stochastic discrepancy, $\text{disc}(A + C) \leq O(\sqrt{k}) + O(\sqrt{t \log n})$. However, this bound is not algorithmic because it requires running the algorithm separately on A and $A_c - A$. ◀

References

1 Dylan J Altschuler and Jonathan Niles-Weed. The discrepancy of random rectangular matrices. *arXiv preprint*, 2021. `arXiv:2101.04036`.

2 Jean-Yves Audibert, Sébastien Bubeck, and Gábor Lugosi. Regret in online combinatorial optimization. *Mathematics of Operations Research*, 39(1):31–45, 2014.

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

4 Nikhil Bansal. Constructive algorithms for discrepancy minimization. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 3–10. IEEE, 2010.

5 Nikhil Bansal, Daniel Dadush, and Shashwat Garg. An algorithm for Komlós conjecture matching Banaszczyk's bound. *SIAM Journal on Computing*, 48(2):534–553, 2019.

6 Nikhil Bansal, Daniel Dadush, Shashwat Garg, and Shachar Lovett. The Gram-Schmidt walk: A cure for the Banaszczyk blues. *Theory Comput.*, 15:1–27, 2019.

7 Nikhil Bansal and Shashwat Garg. Algorithmic discrepancy beyond partial coloring. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 914–926, 2017.

8 Nikhil Bansal, Aditi Laddha, and Santosh S Vempala. A unified approach to discrepancy minimization. *arXiv preprint*, 2022. `arXiv:2205.01023`.

9 Nikhil Bansal and Raghu Meka. On the discrepancy of random low degree set systems. *Random Structures & Algorithms*, 57(3):695–705, 2020.

10 Nikhil Bansal and Joel Spencer. Deterministic discrepancy minimization. *Algorithmica*, 67(4):451–471, 2013.

11 Nikhil Bansal and Joel H. Spencer. On-line balancing of random inputs. *Random Struct. Algorithms*, 57(4):879–891, 2020.

12 Joshua Batson, Daniel A Spielman, and Nikhil Srivastava. Twice-ramanujan sparsifiers. *SIAM Journal on Computing*, 41(6):1704–1721, 2012.

13 József Beck. Roth's estimate of the discrepancy of integer sequences is nearly sharp. *Combinatorica*, 1(4):319–325, 1981.

14 József Beck. Discrepancy theory. *Handbook of combinatorics*, pages 1405–1446, 1995.

15 József Beck and Tibor Fiala. “Integer-making” theorems. *Discrete Applied Mathematics*, 3(1):1–8, 1981.

16 Karthekeyan Chandrasekaran and Santosh S Vempala. Integer feasibility of random polytopes: random integer programs. In *Proceedings of the 5th conference on Innovations in theoretical computer science*, pages 449–458, 2014.

17 Bernard Chazelle. *The discrepancy method: randomness and complexity*. Cambridge University Press, 2001.

18 Ronen Eldan and Mohit Singh. Efficient algorithms for discrepancy minimization in convex sets. *Random Struct. Algorithms*, 53(2):289–307, 2018.

19 Esther Ezra and Shachar Lovett. On the beck-fiala conjecture for random set systems. *Random Structures & Algorithms*, 54(4):665–675, 2019.

20 Efim Davydovich Gluskin. Extremal properties of orthogonal parallelepipeds and their applications to the geometry of banach spaces. *Mathematics of the USSR-Sbornik*, 64(1):85, 1989.

21 Nicholas Harvey, Roy Schwartz, and Mohit Singh. Discrepancy without partial colorings. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2014)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2014.

22 Rebecca Hoberg and Thomas Rothvoss. A fourier-analytic approach for the discrepancy of random set systems. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2547–2556. SIAM, 2019.

23 Yin Tat Lee and Mohit Singh. Personal communication, 2016.

24 Yin Tat Lee and He Sun. Constructing linear-sized spectral sparsification in almost-linear time. In *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on*, pages 250–269. IEEE, 2015.

25 Yin Tat Lee and Santosh S Vempala. Stochastic localization+ stieltjes barrier= tight bound for log-sobolev. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1122–1129, 2018.

26 Avi Levy, Harishchandra Ramadas, and Thomas Rothvoss. Deterministic discrepancy minimization via the multiplicative weight update method. In *International Conference on Integer Programming and Combinatorial Optimization*, pages 380–391. Springer, 2017.

27 Shachar Lovett and Raghu Meka. Constructive discrepancy minimization by walking on the edges. *SIAM Journal on Computing*, 44(5):1573–1582, 2015.

28 Jiri Matousek. *Geometric discrepancy: An illustrated guide*, volume 18. Springer Science & Business Media, 1999.

29 Jiri Matousek. *Geometric discrepancy: An illustrated guide*, volume 18. Springer, 2009.

30 Jiri Matousek. The determinant bound for discrepancy is almost tight. *Proceedings of the American Mathematical Society*, 141:451–60, 2013.

31 Jiri Matousek and Aleksandar Nikolov. Combinatorial discrepancy for boxes via the γ_2 norm. In *31st International Symposium on Computational Geometry, SoCG*, volume 34, pages 1–15, 2015.

32 Aditya Potukuchi. A spectral bound on hypergraph discrepancy. In *47th International Colloquium on Automata, Languages, and Programming, ICALP*, pages 93:1–93:14, 2020.

33 Thomas Rothvoss. Constructive discrepancy minimization for convex sets. *SIAM Journal on Computing*, 46(1):224–234, 2017.

34 Joel Spencer. Six standard deviations suffice. *Transactions of the American mathematical society*, 289(2):679–706, 1985.

35 Nikhil Srivastava, Roman Vershynin, et al. Covariance estimation for distributions with 2+eps moments. *The Annals of Probability*, 41(5):3081–3111, 2013.

36 Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.

A Appendix: Proof of main theorem

Proof of Lemma 11. We will prove this by induction. Clearly, this holds at $t=0$ as $\Phi(0) \le 2m(2/b_0)^p$. For the inductive step, we will show that for any $j = 0,\dots,T-1$, if $\Phi(j) \le 4m(2/b_0)^p$ then

$$\Phi(j+1) \le \Phi(j) + \frac{1}{Tb_0^p} + |\mathcal{I}_{j+1} \backslash \mathcal{I}_j| \cdot \left(\frac{2}{b_0}\right)^p. \tag{24}$$

Note that $|\mathcal{I}_{j+1} \backslash \mathcal{I}_j|$ is the number of additional rows in $\mathcal{I}^S$ that may become alive at step j. This gives the result by induction as summing (24) over $j = 0,\dots,T-1$ will give

$$\Phi(t+1) \le \Phi(0) + \sum_{j=0}^{T-1} \frac{1}{Tb_0^p} + \left(\frac{2}{b_0}\right)^p \sum_{j=0}^{T-1} |\mathcal{I}_{j+1} \backslash \mathcal{I}_j| \le 2m \cdot \left(\frac{2}{b_0}\right)^p + \frac{1}{b_0^p} \le 4m \cdot \left(\frac{2}{b_0}\right)^p. \tag{25}$$

We now focus on proving (24) for $j = t$.

By the induction hypothesis, $\Phi(t) \le 4m\,(2/b_0)^p$. By Lemma 19, one of the signs for x_{t+1} gives $\mathbb{E}(\Phi(t+1)) - \Phi(t) \le f(t) + \frac{1}{Tnb_0^p} + |\mathcal{I}_{t+1} \backslash \mathcal{I}_t| \cdot \left(\frac{2}{b_0}\right)^p$, where

$$f(t) = -p\delta^2 \sum_{i \in \mathcal{I}_t} \frac{d_t + \lambda\langle a_i^{(2)}, v_t^{(2)}\rangle}{s_i(t)^{p+1}} + \frac{p(p+1)\delta^2}{2} \sum_{i \in \mathcal{I}_t} \frac{(2\lambda\langle e_{t,i}, v_t\rangle - \langle a_i, v_t\rangle)^2}{s_i(t)^{p+2}}.$$

So to prove (24), it suffices to show that $f(t) \le 0$. We first consider the case when b_t, d_t and p are given by (13). As $2\lambda\langle e_{t,i}, v_t\rangle - \langle a_i, v_t\rangle = 0$ for all $i \notin \mathcal{J}_t$, $f(t)$ satisfies

$$f(t) \le -p\delta^2 \sum_{i \in \mathcal{J}_t} \frac{d_t + \lambda\langle a_i^{(2)}, v_t^{(2)}\rangle}{s_i(t)^{p+1}} + \frac{p(p+1)\delta^2}{2} \max_{j \in \mathcal{J}_t} s_j(t)^{-1} \cdot \sum_{i \in \mathcal{J}_t} \frac{(\langle 2\lambda e_{t,i} - a_i, v_t\rangle)^2}{s_i(t)^{p+1}}. \tag{26}$$

By a simple averaging argument described in Lemma 13, we also have that

$$\sum_{i \in \mathcal{I}_t} \frac{(2\lambda\langle e_{t,i}, v_t\rangle - \langle a_i, v_t\rangle)^2}{s_i(t)^{p+1}} \le \sum_{i \in \mathcal{I}_t} \frac{8h(n_t)}{s_i(t)^{p+1}}. \tag{27}$$

Plugging (27) in (26) gives

$$f(t) \le -p\delta^2 \sum_{i \in \mathcal{J}_t} \frac{d_t}{s_i(t)^{p+1}} + \frac{p(p+1)\delta^2}{2} \max_{j \in \mathcal{J}_t} s_j(t)^{-1} \cdot \sum_{i \in \mathcal{J}_t} \frac{8h(n_t)}{s_i(t)^{p+1}}. \tag{28}$$

Therefore, if d_t satisfies equation (13), then $f(t) \le 0$.

We now consider the case in (14). As $v_t \in \mathcal{G}_t$, we have

$$\sum_{i \in \mathcal{J}_t} \frac{(2\lambda\langle e_{t,i}, v_t\rangle - \langle a_i, v_t\rangle)^2}{s_i(t)^{p+1}} \le 40 \cdot \sum_{i \in \mathcal{J}_t} \frac{\langle a_i^{(2)}, {v_t}^{(2)}\rangle}{s_i(t)^{p+1}}. \tag{29}$$

Next, as $d_t = 0$ and $\lambda = b_0/42$, (26) and (29) give

$$f(t) \le \sum_{i \in \mathcal{J}_t} \frac{p\delta^2\langle a_i^{(2)}, v_t^{(2)}\rangle}{s_i(t)^{p+1}} \cdot \left(-\frac{b_0}{42} + 20(p+1) \cdot \max_{j \in \mathcal{J}_t} s_j(t)^{-1}\right).$$

So if b_0 satisfies equation (14), then $f(t) \le 0$. ◀

Proof of Lemma 12. By the definition of $\mathcal{J}_t$, for any $i \in \mathcal{J}_t$, there are at least $\lfloor n_t/12 \rfloor + 1$ indices j in $\mathcal{I}_t$ such that $s_j(t) \le s_i(t)$. Therefore,

$$\max_{i \in \mathcal{J}_t} \frac{1}{s_i(t)} \le \left(\frac{12\Phi(t)}{n_t}\right)^{\frac{1}{p}} \le \frac{2}{b_0}\left(\frac{48m}{n_t}\right)^{\frac{1}{p}}, \tag{30}$$

where the last inequality follows by the assumption, $\Phi(t) \le 4m(2/b_0)^p$. ◀

Proof of Lemma 13. Using $(a+b)^2 \le 2(a^2+b^2)$, and as $|2\lambda e_{t,i}(j)| = |2\lambda a_i(j)^2 x_t(j)| \le |a_i(j)|$ as $|a_i(j)| \le 1/2\lambda$ for any j and $i \in \mathcal{I}^S$, we have that for any w,

$$\sum_{i \in \mathcal{J}_t} \frac{\langle 2\lambda e_{t,i} - a_i, w\rangle^2}{s_i(t)^{p+1}} \le \sum_{i \in \mathcal{J}_t} \frac{2\langle a_i, w\rangle^2 + 2\langle 2\lambda e_{t,i}, w\rangle^2}{s_i(t)^{p+1}} \le 4 \sum_{i \in \mathcal{J}_t} \frac{\langle a_i, w\rangle^2}{s_i(t)^{p+1}}.$$

Let $W_t = \{w_1, \ldots, w_k\}$ be an orthonormal basis for $\mathcal{Z}_t$ and $k = \dim(\mathcal{Z}_t)$. As $\mathcal{Z}_t \subseteq \mathcal{V}_t$,

$$\sum_{i \in \mathcal{J}_t} \frac{\sum_{j=1}^k \langle a_i, w_j\rangle^2}{s_i(t)^{p+1}} \le \sum_{i \in \mathcal{J}_t} \frac{\sum_{j \in \mathcal{V}_t} a_i(j)^2}{s_i(t)^{p+1}} \le n_t \sum_{i \in \mathcal{J}_t} \frac{h(n_t)}{s_i(t)^{p+1}}.$$

where the second inequality uses that $\sum_{j \in \mathcal{V}_t} a_i(j)^2 \le n_t \cdot h(n_t)$ by the definition of h.

As $k \ge \lceil n_t/2 \rceil$, this gives

$$\frac{1}{k}\sum_{j=1}^{k} \sum_{i \in \mathcal{J}_t} \frac{\langle 2\lambda e_{t,i} - a_i, w_j\rangle^2}{s_i(t)^{p+1}} \le \frac{n_t}{k} \sum_{i \in \mathcal{J}_t} \frac{4h(n_t)}{s_i(t)^{p+1}} \le \sum_{i \in \mathcal{J}_t} \frac{8h(n_t)}{s_i(t)^{p+1}}.$$

The result now follows as v_t in (12) minimizes $\sum_{i \in \mathcal{J}_t} \langle 2\lambda e_{t,i} - a_i, w_j\rangle^2 s_i(t)^{-p-1}$ over all $w_j \in W_t$. ◀

Proof of Lemma 10. To lower bound the dimension of $\mathcal{Z}_t$ we lower bound the dimensions of $\mathcal{U}_t, \mathcal{Y}_t$ and $\mathcal{G}_t$.

First, we have $\dim(\mathcal{U}_t) \ge n_t - \dim(\mathcal{C}_t) - 1 \ge \lceil 11n_t/12 \rceil - 1$. Second, at time t, as the sum of ℓ_2-norm square of all columns is at most $2n_t$, we have that $\sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{V}_t} a_i(j)^2 \le 2n_t$. So the number of rows a_i with $\sum_{j \in \mathcal{V}_t} a_i(j)^2 \ge 20$ is at most $\lfloor n_t/10 \rfloor$ and $\dim(\mathcal{Y}_t) \ge n_t - \lfloor n_t/10 \rfloor = \lceil 9n_t/10 \rceil$.

We now bound $\dim(\mathcal{G}_t)$ by applying Lemma 9. Let G denote the matrix with columns j corresponding to variables in $\mathcal{V}_t$ and rows i restricted to $i \in \mathcal{J}_t$ with (i,j) entry $(2\lambda e_{t,i}(j) - a_i(j)) s_i(t)^{-(p+1)/2}$.

Let H be the matrix with entries $a_i(j) \cdot s_i(t)^{-(p+1)/2}$ for $i \in \mathcal{J}_t\}$ and $j \in \mathcal{V}_t$. As $|a_{ij}| \le 1/(2\lambda)$ for $i \in \mathcal{I}_t$, we have

$$|G_{ij}| = |2\lambda a_i(j)^2 x_t(j) - a_j(i)| \le |2\lambda a_i(j)^2 x_t(j)| + |a_j(i)| \le 2|a_j(i)| = 2|H_{ij}|.$$

Let $K = \mathrm{diag}(H^\top H)$. Then, using Lemma 9 with $\alpha = 2$ and $\beta = 1/10$, we get that there is a subspace $\mathcal{G}_t$ with $\dim(\mathcal{G}_t) \ge \lceil 9n_t/10 \rceil$ such that $\mathcal{G}_t = \{w \in \mathcal{W}_t : w^\top G^\top G w \le 40 \cdot w^\top K w\}$, which by the definition of G and H is equivalent to that given by (11).

Putting together the bounds on the dimensions of these subspaces gives,

$$\dim(\mathcal{Z}_t) \ge \dim(\mathcal{U}_t \cap \mathcal{Y}_t \cap \mathcal{G}_t) \ge \lceil 11n_t/12 \rceil - 1 + \lceil 9n_t/10 \rceil + \lceil 9n_t/10 \rceil - 2n_t \ge \lceil 2n_t/3 \rceil. ◀$$

Proof of Theorem 1. Recall that we divide each row a of A as $a = a^S + a^L$. We will bound $\langle a^L, x_T\rangle$ and $\langle a^S, x_T\rangle$ separately.

Let t_1 denote the earliest when the squared norm of a^L (restricted to the alive variables) is at most 20, and let n_1 be number of non-zeros in a^L restricted to the set $\mathcal{V}_{t_1}$. As $|a^L(j)| \geq 1/(2\lambda)$ for each j, the number of non-zero variables n_1 in a^L at time t_1 is at most $80\lambda^2$, as $n_1/(4\lambda^2) \leq \sum_{j\in\mathcal{V}_{t_1}} a^L(j)^2 \leq 20$. Moreover, as a^L incurs zero discrepancy until t_1, the overall discrepancy satisfies

$$|\langle a^L, x_T\rangle| = |\langle a^L, x_{t_1}\rangle| + |\langle a^L, x_T - x_{t_1}\rangle| \leq \sqrt{n_1} \cdot (\sum_{j\in\mathcal{V}_{t_1}} a^L(j)^2)^{1/2} \leq 80\lambda \leq 3b_0. \tag{31}$$

Henceforth, we focus on the rows a^S. We first show that the slacks are always positive. Let $\gamma = b_0/4(4m)^{\frac{1}{p}}$. By Lemma 11, for all $t \in [T]$, $\Phi(t) \leq 4m(2/b_0)^p < \gamma^{-p}$. This implies that $|s_i(t)| \geq \gamma$ for all $i \in \mathcal{I}_t^S$ and $t \in [T]$. In one step of the algorithm,

$$\begin{aligned}|s_i(t) - s_i(t-1)| &\leq \delta^2 d_{t-1} + |\langle a_i, x_t\rangle - \langle a_i, x_{t-1}\rangle| \\ &\leq \delta^2 d_{t-1} + |\delta\langle a_i, v_{t-1}\rangle| \leq 20n\delta \leq 2\gamma.\end{aligned}$$

So, if $s_i(t-1) \geq \gamma$ and $\Phi(t) < \gamma^{-p}$, then $s_i(t) \geq 0$, i.e., the slack $s_i(t)$ cannot go from being greater than γ to less than $-\gamma$ in a single step. So, for every $i \in \mathcal{I}^S$ and $t \in [T]$, $s_i(t) \geq \gamma$ and $\langle a_i, x_T\rangle \leq b_T$. Together with (31) this gives, $|\langle a, x_T\rangle| \leq |\langle a^S, x_T\rangle| + |\langle a^L, x_T\rangle| \leq b_T + 3b_0$.

Let $x \in \{-1,1\}^n$ be obtained from x_T by the rounding $x(j) = \mathrm{sign}(x_T(j))$. As $T = (n-2)/\delta^2$, $\|x_T\|^2 = n-2$ with $|x_\tau(j)| \leq 1$ for all $j \in [n]$. After rounding x_T to x, we have $\|x\|^2 = n$. For any row a of A, the discrepancy is bounded by

$$|\langle a, x\rangle| = |\langle a, x_T\rangle| + |\langle a, x - x_T\rangle| \leq |\langle a, x_T\rangle| + M\sum_{j=1}^n |x(j) - x_T(j)| \leq b_T + 3b_0 + 2M.$$

We now consider the two cases for b_0, d_t, p. If the second case given by (14), then by (30), $b_0 \leq 1680(p+1)\cdot(48m/n_t)^{1/p}/b_0$. As $n_t \geq 1$ for all $t \in [T]$ and $p = \log(2m)$, we have $(48m/n_t)^{1/p} \leq 10e$, and setting $b_0 = 250\sqrt{\log(2m)}$ suffices. Since $d_t = 0$, $b_T = b_0$ and $\|Ax\|_\infty \leq 4b_0 + 2M$.

In the first case given by (13), then by (30), we have $d_t = 8(p+1)(48m)^{\frac{1}{p}} \cdot \frac{h(n_t)}{b_0 n_t^{1/p}}$ for all $t \in [T]$. Summing d_t over t gives

$$b_T - b_0 = \delta^2\sum_{t=0}^{T-1} d_t = 8(p+1)(48m)^{\frac{1}{p}}\delta^2 \cdot \sum_{t=0}^{T-1} h(n_t)/(b_0 n_t^{1/p}).$$

As $n_t > n - \delta^2 t - 1 \geq$ and h is non-increasing, $\delta^2 \cdot \sum_{t=0}^{T-1} h(n_t)n_t^{-1/p} \leq \beta$, so that $b_T \leq b_0 + 8(p+1)(48m)^{1/p}\beta/b_0$. Optimizing $b_0 = (8(p+1)(48m)^{1/p}\beta)^{1/2}$ gives that $b_T = 2b_0$ and thus $\|Ax\|_\infty \leq b_T + 3b_0 + 2M \leq 5b_0 + 2M$, giving the desired result. ◀

Proof of Corollary 8. For a constant h, we have $\beta = \int_0^{n-2}(n-t)^{-1/p}h\,dt \leq n^{1-1/p}h/(1-1/p)$. Choosing $p = \log(2m/n)$ to optimize the first term in (4) gives the result. ◀

A.1 Sub-Gaussian Matrices and Random Matrices

Let X be a random variable with $\mathbb{E}(X) = 0$. X is called Sub-Gaussian with variance σ^2 if its moment generating function satisfies $\mathbb{E}(e^{sX}) \leq e^{\sigma^2 s^2/2}$ for all $s \in \mathbb{R}$. For a Sub-Gaussian random variable, $\mathbb{E}(X^2) \leq 4\sigma^2$.

Proof of Theorem 2. As $a_i(j)$ is a Sub-Gaussian with variance σ^2, $a_i(j)^2 - \mathbb{E}(a_i(j)^2)$ is a mean zero and sub-exponential random variable with parameter $16\sigma^2$ [36].

For any $S \subseteq [n]$ with $|S| = s$, Bernstein's inequality for sub-exponential random variables [36] (Theorem 2.8.1) gives that,

$$\Pr(\sum_{j\in S} a_i(j)^2 - \mathbb{E}(a_i(j)^2) \geq st) \leq \exp(-\min(s^2t^2/16\sigma^4, st/16\sigma^2)). \tag{32}$$

Setting $t = 96\sigma^2 (\log(ne/s) + (\log m)/s)$ and as $\mathbb{E}(a_i(j)^2) \leq 4\sigma^2$, and taking a union bound over all the rows and all possible subsets of s columns, we get that,

$$\sum_{j\in S} a_i^2(j) \leq 100\sigma^2|S| (\log(ne/|S|) + \log m)/|S|)) . \tag{33}$$

for every $S \subseteq [n]$, $i \in [m]$, with probability at least $1 - 1/2m^2$.

Similarly, as $a_i(j)$ is sub-Gaussian with mean 0 and variance σ^2, with probability at least $1 - 1/2m^2$, we have $|a_i(j)| \leq 3\sigma\sqrt{\log(mn)}$ for all $i \in [m], j \in [n]$, and thus the ℓ_2-norm of a column is at most $L = 3\sqrt{m}\sigma\sqrt{\log(mn)}$ and $M = 3\sigma\sqrt{\log mn}$. By (33), we can set

$$h(t) = 100\sigma^2 \left(\log\left(\frac{ne}{t}\right) + \frac{\log m}{t}\right).$$

A direct computation gives $\beta = \int_0^{n-2} h(n-t)(n-t)^{-1/p}dt = O(\sigma^2(n^{1-1/p} + p\log m))$. Using Theorem 1 with $p = 2\lceil\log(2m/n)\rceil$, gives $b_0 = O(\sigma(p(m/n)^{1/p}(n + n^{1/p}p\log m))^{1/2}) = O(\sigma n^{1/2}\log(2m/n))$.

Thus, with high probability $\|Ax\|_\infty \leq (5b_0 + 2M) = O(\sigma\sqrt{n\log(2m/n)})$. ◀

Proof of Theorem 3. Consider a random vector X chosen uniformly at random from the unit ball, $\{x \in \mathbb{R}^m : \|x\|_2 \leq 1\}$. Then every coordinate of X is sub-Gaussian with variance $\sigma^2 = C/\sqrt{m}$, where C is a constant [36] (Theorem 3.4.6, Ex 3.4.7). The result now follows from Theorem 5. ◀

A.2 Subadditive Stochastic Discrepancy

Proof of Theorem 4. Let $\Phi_1(t)$, $\Phi_2(t)$ be the potential functions corresponding to A and B, respectively. Let the parameters for Algorithm 2 on A be $b_0^1, p_1, d_t^1, h_1(\cdot)$ and for B be $b_0^2, p_2, d_t^2, h_2(\cdot)$.

Note that it might not be possible to select an update v_t at time t, that ensures that both $\Phi_1(t+1) \leq \Phi_1(t)$ and $\Phi_2(t+1) \leq \Phi_2(t)$ hold, but we can find a v_t for which a weighted sum of $\Phi_1(t)$ and $\Phi_2(t)$ decreases at every step.

Consider the potential function $\Phi(t) = \left(b_0^1/2\right)^{p_1}\Phi_1(t) + (b_0^2/2)^{p_2}\Phi_2(t)$. We apply the same algorithmic framework. For $t = 1, \ldots, T$, select v_t such that $\mathbb{E}(\Phi(t+1)) \leq \Phi(t)$, and select the sign of ε for which $\Phi(t+1) \leq \Phi(t)$, and set $x_{t+1} = x_t + \epsilon\delta v_t$. To this end, it suffices to find a v_t such that $\mathbb{E}(\Phi_1(t+1)) \leq \Phi_1(t)$ and $\mathbb{E}(\Phi_2(t+1)) \leq \Phi_2(t)$.

Let $\mathcal{Z}_t^1$ and $\mathcal{Z}_t^2$ be the feasible subspaces at step t for A and B respectively from Algorithm 2. We will search for v_t in $\mathcal{Z}_t = \mathcal{Z}_t^1 \cap \mathcal{Z}_t^2$. By Lemma 10, $\dim(\mathcal{Z}_t^1), \dim(\mathcal{Z}_t^2) \geq \lceil 2n_t/3\rceil$. Therefore, $\dim(\mathcal{Z}_t) = \dim(\mathcal{Z}_t^1 \cap \mathcal{Z}_t^2) \geq \lceil 2n_t/3\rceil + \lceil 2n_t/3\rceil - n_t \geq n_t/3$.

Using Lemma 13 on A and B, along with Markov's inequality implies that there exists a vector $w \in \mathcal{Z}_t$ such that

$$\sum_{i\in\mathcal{I}_t^1} \frac{\langle 2cb_0^1 e_{t,i} - a_i, w\rangle^2}{s_i(t)^{p_1+1}} \leq \sum_{i\in\mathcal{I}_t^1} \frac{25h_1(n_t)}{s_i(t)^{p_1+1}} \quad \text{and} \quad \sum_{i\in\mathcal{I}_t^2} \frac{\langle 2cb_0^2 e_{t,i} - a_i, w\rangle^2}{s_i(t)^{p_2+1}} \leq \sum_{i\in\mathcal{I}_t^2} \frac{25h_2(n_t)}{s_i(t)^{p_2+1}}.$$

$$(34)$$

Comparing (34) with (15), the functions $h_1(\cdot)$ and $h_2(\cdot)$ only increase by a constant factor when compared to running Algorithm 2 on A and B independently. So it suffices to multiply d_t^1 and d_t^2 by 4 to ensure that by Lemma 11,

$$\mathbb{E}[\Phi_1(t)] - \Phi_1(t-1) \leq \frac{1}{Tn(b_0^1)^{p_1}} \quad \text{and} \quad \mathbb{E}[\Phi_2(t)] - \Phi_2(t-1) \leq \frac{1}{Tn(b_0^2)^{p_2}}. \tag{35}$$

Plugging (35) in the definition of $\Phi(t)$, we get $\mathbb{E}[\Phi(t)] - \Phi(t-1) \leq 2/(Tn)$. So one of the two choices of x_t gives $\Phi(t) - \Phi(t-1) \leq 2/(Tn)$. Summing over t,

$$\Phi(t) \leq \Phi(0) + \frac{2}{n} \leq \left(\frac{b_0^1}{2}\right)^{p_1} \Phi_1(0) + \left(\frac{b_0^2}{2}\right)^{p_2} \Phi_2(0) + \frac{2}{n}.$$

By Lemma 19, $\Phi_1(0) \leq 2m\cdot(2/b_0^1)^{p_1}$ and $\Phi_2(0) \leq 2m\cdot(2/b_0^2)^{p_2}$, thus $\Phi(t) \leq \Phi(0)+2/n \leq 5m$. For a row $i \in \mathcal{J}_t^\ell$ for $\ell \in \{1,2\}$, we have $(\lfloor n_t/12\rfloor + 1)\cdot(b_0^\ell/2)^{p_\ell}\cdot s_i(t)^{-p_\ell} \leq \Phi(t) \leq 5m$, which implies that for any t, and $\ell \in \{1,2\}$,

$$\max_{i\in\mathcal{J}_t^\ell} s_i(t)^{-1} \leq \frac{2}{b_0^\ell}\left(\frac{60m}{n_t}\right)^{\frac{1}{p_\ell}}. \tag{36}$$

Upon comparing (36) with (30), notice that $\max_{k\in\mathcal{J}_t^1} s_k(t)^{-1}$ and $\max_{k\in\mathcal{J}_t^2} s_k(t)^{-1}$ are only a constant factor larger when compared to running Algorithm 2 on A and B separately, and hence the discrepancies for both A and B are only a constant factor larger. ◀

B Appendix: Bounding the step size

▶ **Lemma 19.** *For $A \in \mathbb{R}^{m\times n}$,*

- $\Phi(0) + \sum_t |\mathcal{I}_{t+1}\backslash\mathcal{I}_t| \cdot \left(\frac{2}{b_0}\right)^p \leq 2m \cdot \left(\frac{2}{b_0}\right)^p$.
- *For all $t \in \{0,1,\ldots,T-1\}$, if $\Phi(t) \leq 2^7m^2\left(\frac{2}{b_0}\right)^p$ and $d_t = O(pn\cdot\max_{i\in\mathcal{J}_t} s_i(t)^{-1})$, then for step size $\delta^2 \leq (Cn^2m^6p^4)^{-1}$,*

$$\mathbb{E}(\Phi(t+1)) - \Phi(t) \leq f(t) + \frac{1}{Tnb_0^p} + |\mathcal{I}_{t+1}\backslash\mathcal{I}_t|\cdot\left(\frac{2}{b_0}\right)^p, \text{ where}$$

$$f(t) = -p\delta^2\sum_{i\in\mathcal{I}_t}\frac{d_t + cb_0\langle a_i^{(2)}, v_t^{(2)}\rangle}{s_i(t)^{p+1}} + \frac{p(p+1)\delta^2}{2}\sum_{i\in\mathcal{I}_t}\frac{\langle 2cb_0e_{t,i} - a_i, v_t\rangle^2}{s_i(t)^{p+2}}.$$