

The Gram–Schmidt Walk: A Cure for the Banaszczyk Blues

Nikhil Bansal* Daniel Dadush† Shashwat Garg‡ Shachar Lovett§

Received July 24, 2018; Revised December 12, 2019; Published December 23, 2019

Abstract: A classic result of Banaszczyk (Random Str. & Algor. 1997) states that given any n vectors in \mathbb{R}^m with ℓ_2 -norm at most 1 and any convex body K in \mathbb{R}^m of Gaussian measure at least half, there exists a ± 1 combination of these vectors that lies in $5K$. Banaszczyk’s proof of this result was non-constructive and it was open how to find such a ± 1 combination in polynomial time. In this paper, we give an efficient randomized algorithm to find a ± 1 combination of the vectors which lies in cK for some fixed constant $c > 0$. This leads to new efficient algorithms for several problems in discrepancy theory.

ACM Classification: F.2.2, G.3

AMS Classification: 68W20

Key words and phrases: discrepancy, random walks

1 Introduction

Let (X, \mathcal{S}) be a finite set system, with $X = \{1, 2, \dots, n\}$ and $\mathcal{S} = \{S_1, S_2, \dots, S_m\}$ a collection of subsets of X . Given a two-coloring $x : X \rightarrow \{-1, 1\}$, the discrepancy of a set S is defined as $|x(S)|$ where

A conference version of this paper appeared in the [Proceedings of the 50th ACM Symposium on Theory of Computing, 2018](#).

*Supported by an ERC consolidator grant 617951 and NWO VICI grant 639.023.812.

†Supported by NWO Veni grant 639.071.510.

‡Supported by NWO grant 022.005.025.

§Supported by an NSF CCF award 1614023.

$x(S) = \sum_{i \in S} x(i)$, and measures the imbalance between the number of elements in S colored -1 and 1 . The discrepancy of the set system (X, \mathcal{S}) is defined as

$$\text{disc}(\mathcal{S}) = \min_{x: X \rightarrow \{-1, 1\}} \max_{S \in \mathcal{S}} |x(S)|$$

and is the minimum imbalance achieved by all the sets in \mathcal{S} over all possible two-colorings.

Discrepancy is defined more generally for any $m \times n$ matrix A as

$$\text{disc}(A) = \min_{x \in \{-1, 1\}^n} \|Ax\|_\infty.$$

That is, minimum achievable ℓ_∞ -norm of the vector Ax over all two-colorings x of the columns of A . This can be seen as a vector balancing problem: given vectors $v_1, \dots, v_n \in \mathbb{R}^m$ (specified by the columns of A), find a two-coloring $x: [n] \rightarrow \{-1, 1\}$ to minimize $\|\sum_{i=1}^n x(i)v_i\|_\infty$. The set system view mentioned earlier corresponds to the special case where A is the incidence matrix of the set system with columns indexed by elements in X and rows by sets in \mathcal{S} .

Discrepancy is a widely studied topic and has applications to many areas in mathematics and computer science. In particular in computer science, it arises naturally in computational geometry, data structure lower bounds, rounding in approximation algorithms, combinatorial optimization, communication complexity, pseudorandomness and differential privacy. For more on these connections we refer the reader to [14, 27, 30].

One of the earliest techniques employed in discrepancy theory was linear algebraic in nature and similar to the well-known iterated-rounding technique [9, 12, 22]. Though this technique gave surprisingly good bounds for some problems in discrepancy theory, there remained a big gap between the bounds obtained and the lower bounds known for these problems.

A huge breakthrough was made in the early 80's with Beck's partial-coloring method [11], that was further refined by Spencer to the entropy method [35]. A similar approach based on ideas from convex geometry was developed independently by Gluskin [19]. Roughly speaking, this method guarantees the existence of a coloring of a constant fraction of the elements where every set in the set system incurs a low discrepancy. This is then repeated $O(\log n)$ times in order to get a coloring of all the n elements. The partial-coloring method led to improved bounds for many problems in discrepancy theory and in particular, it led to the famous "six standard deviations" theorem of Spencer [35]: given a set system with n sets and n points, there exists a coloring of discrepancy at most $6\sqrt{n}$. This bound is tight up to constant factors.

While the original proofs of the partial-coloring method were based on the pigeonhole principle and were non-algorithmic, several new algorithmic versions of the partial-coloring method have been developed over the past few years [3, 25, 34, 20, 16]. In particular, all known applications of partial-coloring [35, 27] can now be made algorithmic. These ideas have also led to new results in approximation algorithms [33, 4, 8, 32].

Despite its huge success, the partial-coloring method gives suboptimal bounds for many problems. The reason is that the partial-coloring step only colors a constant fraction of the points and hence must be repeated $O(\log n)$ times before all the n points are colored. The discrepancies incurred at each of these steps are independent of each other and in fact can add up adversarially. Consider for instance the long-standing Beck–Fiala problem [12] about discrepancy of low-degree set systems, where every point

lies in at most t sets. Here, one round of partial-coloring ensures a discrepancy of $O(\sqrt{t})$ to every set, but over all the $O(\log n)$ rounds the discrepancy ends up being $O(\sqrt{t} \log n)$. On the other hand, the Beck–Fiala conjecture is that the discrepancy of such set systems should be $O(\sqrt{t})$. Such logarithmic-factor gaps also exist in several other problems in discrepancy theory for similar reasons.

Banaszcyk’s method. One of the key results in discrepancy theory is the following result by Banaszcyk [1].

Theorem 1.1 (Banaszcyk). *Given any convex body $K \subseteq \mathbb{R}^m$ of Gaussian measure $\gamma_m(K) \geq 1/2$, and vectors $v_1, \dots, v_n \in \mathbb{R}^m$ with ℓ_2 -norm at most 1, there exists a coloring $x : [n] \rightarrow \{-1, 1\}$ such that $\sum_{i=1}^n x(i)v_i \in cK$, where c is an absolute constant. In particular, $c = 5$ suffices.*

Here the Gaussian measure $\gamma_m(S)$ of any measurable set $S \subseteq \mathbb{R}^m$ is defined as

$$\gamma_m(S) = \Pr[g \in S] = \int_{y \in S} \frac{1}{(2\pi)^{m/2}} e^{-\|y\|^2/2} dy$$

where g is a standard Gaussian random vector in \mathbb{R}^m .

In contrast to the partial-coloring method, [Theorem 1.1](#) gives a full coloring directly, resulting in improved bounds for several problems. For instance, Banaszcyk’s result implies a discrepancy bound of $O(\sqrt{t \log n})$ for the Beck–Fiala problem and an $O(\sqrt{\log m})$ bound for the Komlós problem (defined in [Section 1.2](#)). The latter follows by noting that an $O(\sqrt{\log m})$ scaling of the hypercube $[-1, 1]^m$ has Gaussian volume at least half. This was further improved to $O(\sqrt{\log n})$ in [\[10\]](#) (standard reductions give $n \leq m$ [\[12\]](#)). Matoušek et al. [\[28\]](#) used it to bound the hereditary discrepancy with $O(\sqrt{\log m})$ factor of a certain matrix factorization norm, and Larsen [\[21\]](#) used it to give update-query tradeoffs for dynamic data structures. [Theorem 1.1](#) was also used in a very interesting way in a later work by Banaszcyk [\[2\]](#) to show improved bounds for several variants of the Steinitz problem. Recently, Nikolov [\[31\]](#) used this to obtain improved bounds for Tusnády’s problem.

Banaszcyk’s proof is highly elegant and based on deep ideas from convex geometry. However, his approach is non-algorithmic and finding an algorithmic version of [Theorem 1.1](#) has been a major challenge in discrepancy theory [\[34, 30, 15\]](#). Partial progress was made on this recently [\[5, 7, 23\]](#) and algorithms achieving the same bounds as Banaszcyk for the Komlós problem and for the Steinitz problem in the ℓ_∞ -norm were obtained. Roughly, these results correspond to the case when the convex body K in [Theorem 1.1](#) is a scaling of the hypercube, and the question about general convex bodies remained open.

In a recent paper, Dadush et al. [\[15\]](#) reformulated Banaszcyk’s result in terms of certain sub-Gaussian distributions and reduced the question of finding an algorithmic version of Banaszcyk’s result to constructing such sub-Gaussian distributions. To state this result, we first need some definitions. A random vector Y taking values in \mathbb{R}^m is said to be sub-Gaussian with parameter σ (or σ -sub-Gaussian) if for all $\theta \in \mathbb{R}^m$,

$$\mathbb{E} \left[e^{\langle \theta, Y \rangle} \right] \leq e^{(\sigma^2/2) \|\theta\|_2^2}.$$

Roughly, the entries of Y are approximately independent while having Gaussian-like marginals. Observe that a standard Gaussian random vector is 1-sub-Gaussian.

We can now state the result of [\[15\]](#). Let (P1) denote the following statement:

“Let $v_1, \dots, v_n \in \mathbb{R}^m$ be vectors with ℓ_2 -norm at most 1 and let $x_0 \in [-1, 1]^n$. Then there exists a distribution D on $\{-1, 1\}^n$, such that for x sampled from D , the random variable $\sum_{i=1}^n (x(i) - x_0(i))v_i$ is σ -sub-Gaussian, for some absolute constant $\sigma > 0$. Moreover, for every i for which $x_0(i) \in \{-1, 1\}$ we have $x(i) = x_0(i)$ with probability one.” (P1)

Theorem 1.2 (Dadush et al.). *Theorem 1.1 (up to the exact value of c) is equivalent to the statement (P1).*

More precisely, for $v_1, \dots, v_n \in \mathbb{R}^m$ of ℓ_2 -norm at most 1 and a convex body $K \subseteq \mathbb{R}^m$ of Gaussian measure at least $1/2$, there exists $x_0 \in [-1, 1]^n$ such that for $x(1), x(2), \dots, x(n)$ with properties as promised by (P1).

(i) $\sum x_0(i)v_i \in K,$

(ii) $\Pr_{x \sim D}[\sum_{i=1}^n x(i)v_i \in cK] \geq 1/2,$

where D is as in (P1) and $c := c(\sigma)$ is an absolute constant.

Furthermore, for a symmetric convex body (i. e., $K = -K$), the choice $x_0 = 0$ suffices, and for a general convex body given by a membership oracle, x_0 can be computed in expected polynomial time.

The above theorem implies that to get a constructive version of [Theorem 1.1](#) for any convex body it suffices to give an algorithm that can efficiently sample from the σ -sub-Gaussian distributions in (P1). Furthermore, restricting the sampler to the choice $x_0 = 0$, gives a *universal algorithm* for finding colorings that land inside any *symmetric convex body* of Gaussian measure at least $1/2$.

The idea behind [Theorem 1.2](#) is the following. Suppose first that K is symmetric about the origin. Then as $\gamma_m(K) \geq 1/2$, a random Gaussian vector G satisfies $\Pr[G \in K] \geq 1/2$, or equivalently, $\Pr[\|G\|_K \leq 1] \geq 1/2$ where $\|\cdot\|_K$ is the norm having K as its unit ball. By standard tail bounds, this gives that $\mathbb{E}[\|G\|_K] = O(1)$. [Theorem 1.2](#) now follows with $x_0 = 0$ by the following theorem of Talagrand, together with Markov’s inequality.

Theorem 1.3 (Talagrand [38]). *Let $K \subset \mathbb{R}^m$ be a symmetric convex body and $Y \in \mathbb{R}^m$ be a σ -sub-Gaussian random vector. Then for the standard Gaussian in $G \in \mathbb{R}^m$, we have that*

$$\mathbb{E}[\|Y\|_K] \leq \sigma \cdot c \cdot \mathbb{E}[\|G\|_K]$$

where c is an absolute constant.

Next, suppose that K is non-symmetric but satisfies the slightly stronger condition¹ $\gamma_m(K) \geq 3/4$. Then the body $K \cap (-K)$ is convex and symmetric and has Gaussian measure at least $1/2$ (as $\gamma_m(K^c) = \gamma((-K)^c) \leq 1/4$), and we can apply the argument above to $K \cap (-K)$. For non-symmetric bodies K with Gaussian measure $1/2$, Dadush et al. [15] show that one can compute a well chosen center x_0 in K and reduce to the symmetric case by symmetrizing K around x_0 .

¹As far as we know, for all known applications of Banaszczyk’s result, it makes no qualitative difference whether $\gamma_m(K) \geq 3/4$ or $\gamma_m(K) \geq 1/2$.

1.1 Main result

In this paper we give the first efficient algorithm for obtaining Banaszczyk’s result ([Theorem 1.1](#)), which also yields a new constructive proof, by providing a new random walk procedure to sample from the requisite sub-Gaussian coloring distributions. We dub this procedure the Gram–Schmidt walk and state its guarantees below:

Theorem 1.4 (Gram–Schmidt Walk). *There is a polynomial-time randomized algorithm that, given as input vectors $v_1, \dots, v_n \in \mathbb{R}^m$ of ℓ_2 -norm at most 1 and $x_0 \in [-1, 1]^n$, outputs a coloring $x \in \{-1, 1\}^n$ such that the random variable $\sum_{i=1}^n (x(i) - x_0(i))v_i$ is sub-Gaussian with parameter $\sigma = \sqrt{40} \approx 6.32$. Moreover, if $|x_0(i)| = 1$ then $x(i) = x_0(i)$ with probability 1.*

Our algorithm in fact runs in time of $O(n \cdot (n + m)^\omega)$, where ω is the exponent of matrix multiplication. In particular, it runs in n iterations (one variable gets colored at each iteration), and in each iteration the most expensive step is solving a linear system in n variables and m equations, which can be done in time $O((n + m)^\omega)$.

Previous works on the problem either achieved $O(1)$ -sub-Gaussianity only for coordinate directions [[5](#), [7](#)], or achieved $O(\sqrt{\log n})$ -sub-Gaussianity in all directions [[15](#)]. In contrast, [Theorem 1.4](#) gives $O(1)$ -sub-Gaussianity guarantee in all directions. In contrast to the algorithms in [[5](#), [7](#)], where the steps for the walk were generated by solving a semidefinite program at each step, the walk steps for our sampler require only basic linear algebra, namely, the Gram–Schmidt orthogonalization. The idea for the walk was inspired by the constructive proof of Dadush et al. [[15](#)] for the existence of solutions to the Komlós vector coloring program of Nikolov [[29](#)], where the Gram–Schmidt orthogonalization plays a crucial role in the analysis.

1.2 Other results

[Theorem 1.4](#) directly gives an algorithm for previous applications of Banaszczyk’s result.

Komlós problem. This is a generalization of the Beck–Fiala problem and is defined as follows: given an $m \times n$ matrix A with columns of ℓ_2 -norm at most one, find a coloring $x \in \{-1, 1\}^n$ to minimize $\text{disc}(A) = \|Ax\|_\infty$. The Komlós conjecture [[36](#)] states that $\text{disc}(A) = O(1)$ and is a generalisation of the Beck–Fiala conjecture. [Theorem 1.1](#) directly gives an $O(\sqrt{\log n})$ bound for the Komlós problem [[1](#)]. While algorithms to find such a coloring were recently given in [[5](#), [7](#), [23](#)], [Theorem 1.4](#) gives another, more direct, algorithm to find such a coloring.

ℓ_p discrepancy. For $p \in [1, \infty)$, the ℓ_p discrepancy of an $m \times n$ matrix A under a coloring x is defined as

$$\left(\frac{1}{m} \|Ax\|_p^p \right)^{1/p}.$$

Matoušek [[26](#)] showed an ℓ_p discrepancy bound of $O(pt^{1/2})$ for the Beck–Fiala problem, using partial-coloring methods. An improved bound of $O(p^{1/2})$ for the more general Komlós setting follows directly from Banaszczyk’s result (and a standard estimate of the Gaussian measure of the ℓ_p -ball). Our result gives an algorithmic version of this bound.

Corollary 1.5. *There is an efficient randomized algorithm which given a matrix A with n columns of ℓ_2 -norm at most 1, and $p \in [1, \infty)$, finds a $\{-1, 1\}^n$ coloring with expected ℓ_p discrepancy $O(\sqrt{p})$.*

Remark 1.6. The single algorithm given by [Theorem 1.4](#) produces a (random) coloring that simultaneously achieves this bound for every $p \in [1, \infty)$. In contrast, the coloring in Banaszczyk’s approach depends on the body K which is different for different values of p .

Discrepancy relative to γ_2 -norm. Given an $m \times n$ matrix A and a set $J \subseteq [n]$, let $A_{|J}$ denote the $m \times |J|$ matrix restricted to columns of A in J . The *hereditary discrepancy* of A is defined as

$$\text{herdisc}(A) = \max_{J \subseteq [n]} \text{disc}(A_{|J}).$$

Hereditary discrepancy is often a better measure of the complexity of a set system than the discrepancy. It is also rather well behaved; while no polynomial time algorithm can distinguish between set systems with zero discrepancy and set systems with $O(\sqrt{n})$ discrepancy (assuming $P \neq NP$) [[13](#)], hereditary discrepancy can be efficiently approximated. It was shown in [[28](#)] that for any matrix A ,

$$\Omega(\gamma_2(A)/\log m) \leq \text{herdisc}(A) \leq O(\gamma_2(A)\sqrt{\log m}), \tag{1.1}$$

where $\gamma_2(A)$ is a matrix factorization norm that can be efficiently computed by a semidefinite program.

The proof for the upper bound above was non-constructive in the sense that, given any $J \subseteq [n]$, it was not known how to efficiently find a coloring $x : J \rightarrow \{-1, 1\}$ with discrepancy of $A_{|J}$ bounded by the right hand side of [\(1.1\)](#). Using [Theorem 1.4](#), we get an algorithm to find such a coloring.

Corollary 1.7. *There exists an efficient randomized algorithm that, given any $m \times n$ matrix A and $J \subseteq [n]$, returns a coloring $x : J \rightarrow \{-1, 1\}$ such that with constant probability,*

$$\|A_{|J}x\|_\infty = O(\gamma_2(A)\sqrt{\log m}).$$

A generalization of Banaszczyk’s result. We also give a generalization of [Theorem 1.1](#). Let \mathbb{B}_2^m denote the Euclidean ball in \mathbb{R}^m of radius 1 and centered at the origin.

Theorem 1.8. *Let S_1, S_2, \dots, S_n be sets such that for each $i \in [n]$, $S_i \subseteq \mathbb{B}_2^m$ and 0 lies in the convex hull of S_i . Then for any convex body K with $\gamma_m(K) \geq 1/2$, there exist vectors $v_i \in S_i$ such that $\sum_{i=1}^n v_i \in cK$, where $c > 0$ is an absolute constant. Moreover, there is an efficient algorithm to find these vectors.*

[Theorem 1.1](#) is the special case of the above theorem, obtained by taking the sets $S_i = \{-v_i, v_i\}$ for each $i \in [n]$. We will constructively reduce [Theorem 1.8](#) to [Theorem 1.1](#), which implies that an algorithm for the latter also gives an algorithm for the former. Similar generalizations for several other problems in discrepancy theory are mentioned in [[9](#)].

1.3 Organization of the paper

We state the algorithm for [Theorem 1.4](#) in [Section 2](#). The analysis is in [Section 3](#), with a sketch of the main ideas in [Section 3.3](#). The applications are discussed in [Section 4](#).

2 Algorithm description: The Gram–Schmidt walk

The algorithm will proceed in time steps $t = 1, 2, \dots, n$, and will maintain a fractional coloring $x_t \in [-1, 1]^n$ at all time steps. Recall that x_0 is an arbitrary initial fractional coloring. Let x_t denote the coloring at the end of time step t . An element i is called *alive* at time t if $|x_{t-1}(i)| < 1$ and *frozen* (or fixed) otherwise. Let $A_t \subseteq [n]$ denote the set of alive elements at time t .

We now give some notation to describe the update step at each time t . Let $n(t) \in A_t$ denote the largest indexed element that is alive at time t . We call $n(t)$ the *pivot* element at time t . Let V_t denote the subspace spanned by the vectors $\{v_i : i \in A_t, i \neq n(t)\}$ and let $v^\perp(t) := \Pi_{V_t^\perp} v_{n(t)}$ denote the projection of $v_{n(t)}$ orthogonal to V_t .

We now describe the algorithm formally.

Algorithm description:

Input: vectors $v_1, \dots, v_n \in \mathbb{R}^m$ of ℓ_2 -norm at most one; an initial coloring $x_0 \in [-1, 1]^n$.

Output: coloring $x \in \{-1, 1\}^n$.

1. Given the initial coloring x_0 , initialize $A_1 = \{i \in [n] : |x_0(i)| < 1\}$, $n(1) = \max\{i \in A_1\}$ and $t = 1$.

2. While $A_t \neq \emptyset$, do the following:

(a) Compute an update direction $u_t = (u_t(1), \dots, u_t(n)) \in \mathbb{R}^n$ as follows:

- if $i \notin A_t$ set $u_t(i) = 0$.
- if $i = n(t)$ set $u_t(i) = 1$.
- for $i \in A_t \setminus \{n(t)\}$ set $u_t(i)$ to satisfy $v^\perp(t) = v_{n(t)} + \sum_{i \in A_t \setminus \{n(t)\}} u_t(i) v_i$, where $v^\perp(t)$ is the projection of $v_{n(t)}$ on the subspace orthogonal to the span of the vectors $\{v_i : i \in A_t, i \neq n(t)\}$.

(b) Let $\delta_t^- < 0 < \delta_t^+$ be the unique negative and positive solutions for δ , respectively, for the equation $\max_{i \in A_t} |x_{t-1}(i) + \delta u_t(i)| = 1$. Update the coloring x_{t-1} randomly as

$$x_t = x_{t-1} + \delta_t u_t$$

where $\delta_t \in \{\delta_t^-, \delta_t^+\}$ is chosen randomly as

$$\delta_t = \begin{cases} \delta_t^- & \text{with probability } \frac{\delta_t^+}{\delta_t^+ - \delta_t^-}, \\ \delta_t^+ & \text{with probability } \frac{-\delta_t^-}{\delta_t^+ - \delta_t^-}. \end{cases}$$

(c) Update $A_{t+1} = \{i \in [n] : |x_t(i)| < 1\}$, $n(t+1) = \max\{i \in A_{t+1}\}$ and $t \leftarrow t+1$.

3. Output x_t .

Note that $v^\perp(t)$ depends on both A_t and $v_{n(t)}$, and may change between time steps (if A_t changes) even if $n(t)$ remains the same. From the perspective of the walk, $v^\perp(t)$ will correspond to the direction

that the output of the random walk will move in during timestep t . In other words, the update direction u_t , in the step 2(a), ensures that $Au_t = v^\perp(i)$. The step 2(b) is the standard randomized pipage-rounding step [37, 18], where the algorithm updates the coloring randomly along the direction u_t , while ensuring that the update has mean-zero, and x_t stays in $[-1, 1]^n$ and at least one additional coordinate of x_t reaches -1 or 1 . To justify the name of the walk, note that $v^\perp(t)$ is simply the last vector of the Gram–Schmidt orthogonalization of the ordered sequence of vectors $(v_i)_{i \in A_t}$.

3 Algorithm analysis

We proceed now to the analysis of our algorithm. In the first subsection we develop some preliminaries which will be helpful later. The main ideas and the roadmap of the analysis are in [Section 3.3](#).

3.1 Preliminaries

To bound the discrepancy, we will use a concentration inequality which is a variant of Freedman’s inequality for martingales [17]. The following lemma will be useful.

Lemma 3.1. *Let X be a random variable such that $X \leq 1$. Then for any $\lambda > 0$,*

$$\mathbb{E}[e^{\lambda X}] \leq \exp\left(\lambda \mathbb{E}[X] + (e^\lambda - \lambda - 1)\mathbb{E}[X^2]\right).$$

Proof. Let

$$f(x) = \frac{e^{\lambda x} - \lambda x - 1}{x^2}$$

where we set $f(0) = \lambda^2/2$. It can be verified that $f(x)$ is increasing for all x . This implies $e^{\lambda x} \leq f(1)x^2 + 1 + \lambda x$ for any $x \leq 1$. Taking expectations, this becomes

$$\mathbb{E}[e^{\lambda X}] \leq 1 + \mathbb{E}[\lambda X] + f(1)\mathbb{E}[X^2] = 1 + \lambda \mathbb{E}[X] + (e^\lambda - \lambda - 1)\mathbb{E}[X^2] \leq e^{\lambda \mathbb{E}[X] + (e^\lambda - \lambda - 1)\mathbb{E}[X^2]}$$

where the last inequality uses the fact that $1 + x \leq e^x$. □

We will use the following concentration inequality to bound the discrepancy, which is a slight variant of Freedman’s inequality for martingales [17].

Lemma 3.2. *Let X_1, \dots, X_n and Z_0, Z_1, \dots, Z_n be random variables that satisfy*

1. Z_0 is deterministic,
2. $Z_t - Z_{t-1} \leq X_t$ for all $t = 1, \dots, n$ with probability one,
3. $X_t \leq 1$ for all $t = 1, \dots, n$ with probability one, and
4. For all $t = 1, \dots, n$, and for all admissible values z_1, \dots, z_{t-1} for Z_1, \dots, Z_{t-1} , with probability one it holds that $\mathbb{E}[X_t + X_t^2 \mid Z_1 = z_1, \dots, Z_{t-1} = z_{t-1}] \leq 0$.

Then

$$\mathbb{E}[e^{Z_n}] \leq e^{Z_0}.$$

Proof. Let $\lambda > 0$ be a real number to be determined later. We shorthand as $\mathbb{E}_{t-1}[\cdot]$ the conditional expectation $\mathbb{E}[\cdot | Z_1, \dots, Z_{t-1}]$. We first bound $\mathbb{E}_{t-1}[e^{\lambda Z_t}]$ for which we observe the following:

$$\begin{aligned} \mathbb{E}_{t-1}[e^{\lambda Z_t}] &= e^{\lambda Z_{t-1}} \mathbb{E}_{t-1}[e^{\lambda(Z_t - Z_{t-1})}] \\ &\leq e^{\lambda Z_{t-1}} \mathbb{E}_{t-1}[e^{\lambda X_t}] \\ &\leq e^{\lambda Z_{t-1}} \exp\left(\lambda \mathbb{E}_{t-1}[X_t] + (e^\lambda - \lambda - 1) \mathbb{E}_{t-1}[X_t^2]\right) \quad (\text{using Lemma 3.1}) \\ &\leq e^{\lambda Z_{t-1}} \exp\left((e^\lambda - 2\lambda - 1) \mathbb{E}_{t-1}[X_t^2]\right) \quad (\text{using } \mathbb{E}_{t-1}[X_t] \leq -\mathbb{E}_{t-1}[X_t^2]). \end{aligned} \quad (3.1)$$

Set $\lambda > 0$ to be the (unique) solution of $e^\lambda - 2\lambda - 1 = 0$. We first claim that $\lambda \geq 1$. This follows as $e^x \leq 1 + x + x^2$ for $x \leq 1$, and hence for $0 < \lambda < 1$, $e^\lambda - 2\lambda - 1 \leq \lambda^2 - \lambda < 0$.

Using this λ in (3.1) gives $\mathbb{E}_{t-1}[e^{\lambda Z_t}] \leq e^{\lambda Z_{t-1}}$. As

$$\mathbb{E}[e^{\lambda Z_n}] = \mathbb{E}_0 \mathbb{E}_1 \dots \mathbb{E}_{n-1}[e^{\lambda Z_n}],$$

by induction, we have that $\mathbb{E}[e^{\lambda Z_n}] \leq \mathbb{E}[e^{\lambda Z_0}] = e^{\lambda Z_0}$, as Z_0 is deterministic.

As $\lambda \geq 1$, by Jensen's inequality $\mathbb{E}[e^{Z_n}] \leq \mathbb{E}[e^{\lambda Z_n}]^{1/\lambda} \leq e^{Z_0}$, as needed. \square

Remark 3.3. Condition (4) in Lemma 3.2 can be improved to $\mathbb{E}_{t-1}[X_t + cX_t^2] \leq 0$ for $c \geq e - 2$, as the analysis only needs that the non-negative solution λ of $e^\lambda - (c+1)\lambda - 1 = 0$ should satisfy $\lambda \geq 1$.

3.2 Notation and preliminary observations

We can now start with the analysis of the algorithm. Since at the end of each time step t , at least one more element gets colored -1 or 1 , the algorithm clearly terminates in at most n steps. To simplify notation, if the algorithm terminates after $t < n$ steps and outputs x_t , we set $x_{t+1} = \dots = x_n := x_t$. As we maintain $\|x_t\|_\infty \leq 1$ for all time steps t , we see that $x_n \in \{-1, 1\}^n$.

It should also be noted in step (2.a) of the algorithm that such a direction u_t always exists. This is because $v_{n(t)} - v^\perp(t)$ lies in the subspace V_t and hence there always exist $u_t(i)$'s such that

$$v_{n(t)} - v^\perp(t) = - \sum_{i \in A_t \setminus \{n(t)\}} u_t(i) v_i.$$

This then gives $v^\perp(t) = v_{n(t)} + \sum_{i \in A_t \setminus \{n(t)\}} u_t(i) v_i = \sum_{i \in A_t} u_t(i) v_i$. As $u_t(i) = 0$ if $i \notin A_t$, we obtain that

$$v^\perp(t) = \sum_{i=1}^n u_t(i) v_i. \quad (3.2)$$

We now focus on showing the sub-Gaussian bound in Theorem 1.4. Henceforth, we fix a vector $\theta \in \mathbb{R}^m$ with respect to which we want to show sub-Gaussianity. Let A be the $m \times n$ matrix whose columns are given by v_1, \dots, v_n . Define $Y := \sum_{i=1}^n (x(i) - x_0(i)) v_i$ and let

$$\text{disc}(\theta) = \langle \theta, Y \rangle = \langle \theta, \sum_{i=1}^n (x(i) - x_0(i)) v_i \rangle = \sum_{i=1}^n (x(i) - x_0(i)) \langle \theta, v_i \rangle$$

and hence $\langle \theta, Y \rangle$ can be seen as the discrepancy of the “row” $\theta^T A$, which has $\langle \theta, v_i \rangle$ as its i -th entry.

Let us denote the respective signed discrepancy at the end of time step t by

$$\text{disc}_t := \langle \theta, \sum_{i=1}^n (x_t(i) - x_0(i)) v_i \rangle = \sum_{i=1}^n (x_t(i) - x_0(i)) \langle \theta, v_i \rangle,$$

let

$$\Delta_t x := x_t - x_{t-1} = \delta_t u_t$$

denote the coloring update at time t , and let

$$\Delta_t \text{disc} := \text{disc}_t - \text{disc}_{t-1} = \sum_{i=1}^n \langle \theta, v_i \rangle \Delta_t x(i) = \delta_t \sum_{i=1}^n \langle \theta, v_i \rangle u_t(i)$$

be the change in discrepancy at time t . Our end goal is to show that $\mathbb{E}[e^{\text{disc}_n}] \leq e^{(\sigma^2/2)\|\theta\|_2^2}$, for $\sigma = O(1)$ to be computed later.

A key observation is that the change in discrepancy at time t depends only on the vector $v^\perp(t)$.

Lemma 3.4. *At each time step t , $\Delta_t \text{disc} = \delta_t \langle \theta, v^\perp(t) \rangle$.*

Proof. The change in discrepancy at time t is

$$\Delta_t \text{disc} = \delta_t \sum_{i=1}^n \langle \theta, v_i \rangle u_t(i) = \delta_t \langle \theta, \sum_{i=1}^n u_t(i) v_i \rangle = \delta_t \langle \theta, v^\perp(t) \rangle,$$

where the last equality follows by (3.2). □

3.3 Main ideas

Before we give the technical details, we describe the main ideas of the analysis, which are simple, and then give a roadmap of the analysis.

First, as our algorithm can start with any initial coloring x_0 , we may assume without loss of generality that the vectors v_1, \dots, v_n are linearly independent. Otherwise, we can apply a standard preprocessing step that removes linear dependencies, by finding some linear dependency among the vectors and making an update step which incurs zero discrepancy.

Our algorithm can be viewed as a randomized extension of the above dependent rounding approach. Suppose for a moment that the element that got colored at each time step was the pivot. That is, the elements got colored in the order $n, n-1, \dots, 1$. Then, at time t , the pivot is $n(t) = n - t + 1$ and the vectors $v^\perp(t)$ can be described as follows. Let w_1, \dots, w_n be the orthonormal vectors obtained by applying the Gram–Schmidt orthonormalization procedure (GS) on the vectors v_1, \dots, v_n in that order. That is, $w_1 = v_1 / \|v_1\|$ and for $i > 1$, w_i is the projection of v_i orthogonal to v_1, \dots, v_{i-1} , normalized to have unit norm. By our assumption that the vectors v_i are linearly independent, each w_i is non-zero. It is easily checked that $v^\perp(t) = \langle v_{n(t)}, w_{n(t)} \rangle w_{n(t)}$. Another observation that will be useful later is that $w_{n(t)}$ (and hence $v^\perp(t)$) depends only on the set $\{v_1, \dots, v_{n(t)-1}\}$ and not the particular order in which GS is applied to this set.

Now δ_t is a mean-zero random variable which is independently chosen at each time t . Also $|\delta_t| \leq 2$ for all t (see Lemma 3.7). This suggests that the moment generating function of the discrepancy is

$$\mathbb{E} \left[e^{\text{disc}(\theta)} \right] = \mathbb{E} \left[e^{\sum_{t=1}^n \delta_t \langle \theta, v^\perp(t) \rangle} \right] \leq e^{O(1) \cdot \sum_{t=1}^n \langle \theta, v^\perp(t) \rangle^2}.$$

But this is at most $e^{O(1) \cdot \|\theta\|_2^2}$ as

$$\sum_t \langle \theta, v^\perp(t) \rangle^2 = \sum_t \langle \theta, \langle v_{n(t)}, w_{n(t)} \rangle w_{n(t)} \rangle^2 \leq \sum_t \langle \theta, w_{n(t)} \rangle^2 \leq \|\theta\|_2^2.$$

Here the first inequality follows as $|\langle v_{n(t)}, w_{n(t)} \rangle| \leq 1$ as the w_i 's are orthonormal unit vectors and $\|v_i\|_2 \leq 1$ for each i . The second inequality follows as $\sum_i \langle \theta, w_i \rangle^2 \leq \|\theta\|_2^2$ for any orthonormal collection of vectors w_i .

There are two issues that need to be addressed in the simplified description above. First, non-pivot elements may get colored sometimes. That is, the variables may not get colored in the order $n, n-1, \dots, 1$. The key point of the analysis is to show that this only makes the problem easier. To make this a bit more precise, let us view $\sum_i \langle \theta, w_i \rangle^2 \leq \|\theta\|_2^2$ as the energy budget initially available to us. If some non-pivot element x_k is colored at some time t , then the GS procedure (without v_k) will produce a different set of orthonormal vectors $\{w'_i\}$. However we can bound the increase

$$\langle \theta, w'_{n(t)} \rangle^2 - \langle \theta, w_{n(t)} \rangle^2$$

in the pivot's energy by the amount $\langle \theta, w_k \rangle^2$ which was available to us, but we will never use it anymore as k will never be a pivot once it is colored. The formal analysis later on will divide the time steps into *phases* where the pivot element remains the same during a phase and consider the evolution of $v^\perp(t)$ in each phase.

A second technical issue is that to obtain a sub-Gaussian distribution for all $\theta \in \mathbb{R}^m$, we need to control the variance of the energy which requires that $|\langle \theta, v^\perp(t) \rangle| = O(1)$. Let us refer to the times when this does not happen as bad times. To get around this issue we break the discrepancy contribution into two parts: a deterministic part due to the bad times, and a random contribution due to the other time steps. Again by considering the dynamics of how $v^\perp(t)$ evolves during a phase, one can show that the cumulative discrepancy due to all the bad times can be at most $O(\|\theta\|_2^2)$.

3.4 Phases and dynamics of pivots

To analyze the process, we need to set up some more notation. Let T be the first time at the end of which all the elements are colored. Recall that at the beginning of each time step $t \in [T]$, there is some pivot element $n(t)$ and some set of alive elements A_t . Recall that $V_t = \text{span}\{v_i : i \in A_t, i \neq n(t)\}$. After the update at time t , some element is frozen. By convention, we define $V_0 := \text{span}\{v_1, \dots, v_n\}$. Note that for any time $t \in [T]$, we have $V_t \subseteq V_{t-1}$.

We divide the time steps $t \in [T]$ into $\kappa \in [T]$ disjoint phases, where a phase is a maximal sequence of time steps with the same pivot element (note that κ is a random variable). Let t_k^b be the time when phase $k \in [\kappa]$ begins and t_k^e be the time when this phase ends. Note that $n(t) = n(t_k^b)$ for all $t \in [t_k^b, t_k^e]$ and the pivot element $n(t_k^b)$ gets frozen after the update at time step t_k^e . Thus, the first phase, with $n(1)$ as pivot,

begins at the beginning of time step $t_1^b = 1$ and ends at the end of time step $t_1^e = t_2^b - 1$. Given a time $t \in [T]$, we define $f_t = \min \{t' \in [t] : n(t') = n(t)\}$ to be the beginning time of the phase that t belongs to. Note that if $t \in [t_k^b, t_k^e]$ then $f_t = t_k^b$.

We will now give a useful characterization of the discrepancy vector $v^\perp(t)$ at any time step $t \in [T]$. For any subspace $W \subseteq \mathbb{R}^m$, let W^\perp denote the subspace orthogonal to W and let $\Pi_W(\cdot)$ denote the projection operator on W . For two linear subspaces W_1, W_2 satisfying $W_1 \supseteq W_2$, we use the notation $W_1/W_2 := W_1 \cap W_2^\perp$.

Lemma 3.5. *At each time step $t \in [T]$ of the algorithm the following holds:*

$$v^\perp(t) = \sum_{i=f_t}^t \Pi_{V_{i-1}/V_i}(v_{n(f_i)}), \quad (3.3)$$

where the subspaces V_{i-1}/V_i are mutually orthogonal.

Lemma 3.5 follows directly from the following useful fact.

Lemma 3.6. *Let $\mathbb{R}^m \supseteq W_0 \supseteq W_1 \supseteq \dots \supseteq W_t$, $t \geq 1$, denote a non-increasing sequence of linear subspaces and let $y \in W_0$. Then the subspaces W_{i-1}/W_i for $i \in [t]$ are mutually orthogonal and are also orthogonal to W_t , and*

$$\Pi_{W_t^\perp}(y) = \sum_{i=1}^t \Pi_{W_{i-1}/W_i}(y). \quad (3.4)$$

Proof. We first prove orthogonality. Let $W_{t+1} = \emptyset$ denote the empty subspace. Then $W_t = W_t/W_{t+1}$. Now for $1 \leq i < j \leq t+1$, we need to show that W_{i-1}/W_i and W_{j-1}/W_j are orthogonal. This follows directly since W_{i-1}/W_i is orthogonal to the subspace W_i and $W_{j-1}/W_j \subseteq W_{j-1} \subseteq W_i$ since $j-1 \geq i$.

To prove (3.4), notice that $\text{span}\{W_i \cup (W_{i-1}/W_i)\} = W_{i-1}$ and thus by induction,

$$W_0 = \text{span} \left\{ \bigcup_{i=1}^{t+1} (W_{i-1}/W_i) \right\}$$

where each of the subspaces W_{i-1}/W_i are mutually orthogonal. This implies that for $y \in W_0$,

$$y = \Pi_{W_0}(y) = \sum_{i=1}^{t+1} \Pi_{W_{i-1}/W_i}(y) = \sum_{i=1}^t \Pi_{W_{i-1}/W_i}(y) + \Pi_{W_t}(y).$$

The lemma follows now by observing that $\Pi_{W_t^\perp}(\Pi_{W_t}) = 0$ and $\Pi_{W_t^\perp}(\Pi_{W_{i-1}/W_i}) = \Pi_{W_{i-1}/W_i}$ as $W_{i-1}/W_i \subseteq W_t^\perp$. \square

Proof of Lemma 3.5. First, by definition of f_t , note that $n(t) = n(f_t)$ and hence

$$v^\perp(t) = \Pi_{V_t^\perp}(v_{n(t)}) = \Pi_{V_t^\perp}(v_{n(f_t)}).$$

We now check that $v_{n(f_t)} \in V_{f_t-1}$. If $f_t = 1$, this is immediate since $V_0 = \text{span}\{v_1, \dots, v_n\}$ by convention. Otherwise, since $n(f_t)$ is both alive and not the pivot at time $f_t - 1 \geq 1$, we have that $n(f_t) \in A_{f_t-1} \Rightarrow v_{n(f_t)} \in V_{f_t-1}$. The main statement now follows directly by applying **Lemma 3.6** on the subspaces $V_{f_t-1} \supseteq V_{f_t} \supseteq \dots \supseteq V_t$ and the vector $v_{n(f_t)}$. \square

3.5 Discrepancy in a phase

We now bound the discrepancy incurred during any subinterval of a phase. We first need the following simple but very useful fact.

Lemma 3.7. *Let $p \leq q$ be any two time steps in $[t_k^b, t_k^e]$ during a phase $k \in [\kappa]$. Then $|\sum_{t=p}^q \delta_t| \leq 2$.*

Proof. At any time $t \in [p, q]$, note that $n(t) = n(p)$, i. e., the pivot remains unchanged. The color of the pivot element $n(p)$ at time $t \in [p, q]$ is updated by $\delta_t u_t(n(p)) = \delta_t$ and hence

$$x_q(n(p)) - x_{p-1}(n(p)) = \sum_{t=p}^q \delta_t.$$

As $|x_t(n(p))| \leq 1$ for all $t \in [T]$, we have that $|x_q(n(p)) - x_{p-1}(n(p))| \leq 2$ as needed. \square

Lemma 3.8. *Let $p \leq q$ be any two time steps in $[t_k^b, t_k^e]$ during a phase $k \in [\kappa]$. The discrepancy $|\text{disc}_q - \text{disc}_{p-1}|$ incurred during the time interval $[p, q]$ is at most $2\|\theta^{(k)}\|_2$, where*

$$\theta^{(k)} := \Pi_{V_{t_k^b-1}/V_{t_k^e}}(\theta).$$

Furthermore, the subspaces $V_{t_k^b-1}/V_{t_k^e}$ for $k \in [\kappa]$ are mutually orthogonal.

Proof. We first prove orthogonality. Take $k_1, k_2 \in [\kappa]$, where $k_1 < k_2$. We must show that the subspaces

$$V_{t_{k_1}^b-1}/V_{t_{k_1}^e} \quad \text{and} \quad V_{t_{k_2}^b-1}/V_{t_{k_2}^e}$$

are orthogonal. This follows since the first subspace is orthogonal to $V_{t_{k_1}^e}$ and the second subspace is contained in $V_{t_{k_2}^b-1} \subseteq V_{t_{k_1}^e}$ since $t_{k_2}^b - 1 \geq t_{k_1}^e$. We now prove the main statement. First, note that

$$\text{disc}_q - \text{disc}_{p-1} = \sum_{t=p}^q \delta_t \langle \theta, v^\perp(t) \rangle. \quad (3.5)$$

By Lemma 3.5, for $t \in [p, q] \subseteq [t_k^b, t_k^e]$, letting $h := n(t_k^b) = n(t)$ denote the pivot index, we have that

$$v^\perp(t) := \Pi_{V_t^\perp}(v_h) = \sum_{i=t_k^b}^t \Pi_{V_{i-1}/V_i}(v_h). \quad (3.6)$$

Combining (3.5), (3.6) we get that

$$\begin{aligned} |\text{disc}_q - \text{disc}_{p-1}| &= \left| \sum_{t=p}^q \delta_t \cdot \sum_{i=t_k^b}^t \langle \Pi_{V_{i-1}/V_i}(v_h), \theta \rangle \right| \\ &= \left| \sum_{i=t_k^b}^q \langle \Pi_{V_{i-1}/V_i}(v_h), \theta \rangle \cdot \sum_{t=\max\{i,p\}}^q \delta_t \right| \end{aligned}$$

$$\begin{aligned}
 &\leq \sum_{i=t_k^b}^q \left| \langle \Pi_{V_{i-1}/V_i}(v_h), \theta \rangle \right| \cdot \left| \sum_{t=\max\{i,p\}}^q \delta_t \right| \\
 &\leq 2 \sum_{i=t_k^b}^q \left| \langle \Pi_{V_{i-1}/V_i}(v_h), \theta \rangle \right| && \text{(by Lemma 3.7)} \\
 &= 2 \sum_{i=t_k^b}^q \left| \langle \Pi_{V_{i-1}/V_i}(v_h), \Pi_{V_{i-1}/V_i}(\theta) \rangle \right|. \tag{3.7}
 \end{aligned}$$

Now applying the Cauchy-Schwarz inequality twice, we get

$$\begin{aligned}
 &\sum_{i=t_k^b}^q \left| \langle \Pi_{V_{i-1}/V_i}(v_h), \Pi_{V_{i-1}/V_i}(\theta) \rangle \right| && \tag{3.8} \\
 &\leq \sum_{i=t_k^b}^q \|\Pi_{V_{i-1}/V_i}(v_h)\|_2 \|\Pi_{V_{i-1}/V_i}(\theta)\|_2 \\
 &\leq \left(\sum_{i=t_k^b}^q \|\Pi_{V_{i-1}/V_i}(v_h)\|_2^2 \right)^{1/2} \cdot \left(\sum_{i=t_k^b}^q \|\Pi_{V_{i-1}/V_i}(\theta)\|_2^2 \right)^{1/2} \\
 &\leq \|\Pi_{V_{t_k^b-1}/V_q}(v_h)\|_2 \cdot \|\Pi_{V_{t_k^b-1}/V_q}(\theta)\|_2 \\
 &\leq \|\Pi_{V_{t_k^b-1}/V_q}(\theta)\|_2 && \text{(since } \|v_h\|_2 \leq 1) \\
 &\leq \|\Pi_{V_{t_k^b-1}/V_{t_k^e}}(\theta)\|_2 && \text{(since } V_q \supseteq V_{t_k^e}) \tag{3.9}
 \end{aligned}$$

where the third inequality follows by orthogonality of the subspaces V_{i-1}/V_i for all $i \in [t_k^b, t_k^e]$, and their containment inside $V_{t_k^b-1}/V_q$. The lemma now follows by combining (3.7), (3.9). \square

Notice that discrepancy $\{\text{disc}_t : t \geq 0\}$ is a martingale. For the rest of the analysis, we will define another closely related martingale $\{Y_t\}$ and show that the sub-Gaussianity of $\{\text{disc}_t\}$ follows from the sub-Gaussianity of $\{Y_t\}$. Then, we will show the sub-Gaussianity of $\{Y_t\}$.

Define the random process $\{Y_t : t \geq 0\}$ as

$$Y_t := \begin{cases} 0 & \text{if } t = 0, \\ Y_{t-1} & \text{if } \|\Pi_{V_{t-1}/V_t}(\theta)\|_2 > 1/8, \\ Y_{t-1} + \Delta_t \text{ disc} & \text{otherwise.} \end{cases}$$

Note that $\mathbb{E}[Y_t | Y_1, \dots, Y_{t-1}] = Y_{t-1}$ and thus $\{Y_t\}$ is a martingale. To prove [Theorem 1.4](#), we will bound the exponential moment of disc_n .

Theorem 3.9. $\mathbb{E}[e^{\text{disc}_n}] \leq e^{20\|\theta\|_2^2}$.

Note that this directly gives that the walk is $\sqrt{40} \approx 6.32$ -sub-Gaussian.

To control this exponential moment, we will split

$$\text{disc}_n = \text{disc}_T = (\text{disc}_T - Y_T) + Y_T,$$

where $(\text{disc}_T - Y_T)$ will correspond to the discrepancy incurred during “bad” times and Y_T corresponds to the “good” times. More precisely, call a time step t *good* if $\|\Pi_{V_{f_{t-1}}/V_t}(\theta)\|_2 \leq 1/8$ and *bad* otherwise. Let

$$\mathcal{B} := \left\{ t \in [T] : \|\Pi_{V_{f_{t-1}}/V_t}(\theta)\|_2 > 1/8 \right\}$$

denote the set of bad times. Note that

$$\text{disc}_T = Y_T + \sum_{t \in \mathcal{B}} \Delta_t \text{disc}.$$

The following lemma now shows that the discrepancy incurred during the bad times can be upper bounded deterministically.

Lemma 3.10. $|\sum_{t \in \mathcal{B}} \Delta_t \text{disc}| < 16\|\theta\|_2^2$.

Proof. For each $k \in [\kappa]$, let \mathcal{B}_k denote the set of bad time steps in phase k . Notice that in a phase, once a time step becomes bad, all subsequent time steps in that phase are bad as the length of $\Pi_{V_{f_{t-1}}/V_t}(\theta)$ is a non-decreasing function of t during a phase (as $V_{t+1} \subseteq V_t$ for every t). Thus, the set \mathcal{B}_k , if non-empty, forms a consecutive interval (ending at the end of phase k), and hence [Lemma 3.8](#) can be applied to it.

Recall that we denoted

$$\theta^{(k)} := \Pi_{V_{t_k^b-1}/V_{t_k^e}}(\theta) \quad \text{for } k \in [\kappa].$$

Note that the vectors $\theta^{(1)}, \dots, \theta^{(\kappa)}$ are mutually orthogonal. Furthermore, by the argument above, for any phase $k \in [\kappa]$ with $\mathcal{B}_k \neq \emptyset$ we have that $\|\theta^{(k)}\| \geq 1/8$.

The discrepancy incurred during bad times can now be bounded as follows

$$\begin{aligned} \left| \sum_{t \in \mathcal{B}} \Delta_t \text{disc} \right| &\leq \sum_{k \in [\kappa]} \left| \sum_{t \in \mathcal{B}_k} \Delta_t \text{disc} \right| \\ &\leq \sum_{k \in [\kappa], \mathcal{B}_k \neq \emptyset} 2\|\theta^{(k)}\|_2 && \text{(by Lemma 3.8)} \\ &\leq \sum_{k \in [\kappa], \mathcal{B}_k \neq \emptyset} 16\|\theta^{(k)}\|_2^2 && \left(\|\theta^{(k)}\| \geq 1/8 \text{ if } \mathcal{B}_k \neq \emptyset \right) \\ &\leq 16\|\theta\|_2^2 && \left(\text{by orthogonality of } \theta^{(1)}, \dots, \theta^{(\kappa)} \right). \quad \square \end{aligned}$$

3.6 Bounding the discrepancy

We now prove [Theorem 3.9](#). To this end, we will define potential functions that capture how the variance of Y_t increases over time. For $t \in [T]$, define the subspace $R_t := \text{span}\{v_i : i \in A_t\}$, i. e., the span of the active vectors. Note that at any timestep $t \in [T]$, the discrepancy direction $v^\perp(t) \in R_t/V_t$. Consider the potentials

$$\Phi_t^b := \begin{cases} \|\Pi_{V_t}(\theta)\|^2 + (1 - x_{t-1}(n(t))^2) \|\Pi_{R_t/V_t}(\theta)\|_2^2 & \text{if } t \text{ is good,} \\ \|\Pi_{V_t}(\theta)\|^2 & \text{if } t \text{ is bad,} \end{cases}$$

and

$$\Phi_t^e := \begin{cases} \|\Pi_{V_t}(\theta)\|^2 + (1 - x_t(n(t))^2) \|\Pi_{R_t/V_t}(\theta)\|_2^2 & \text{if } t \text{ is good,} \\ \|\Pi_{V_t}(\theta)\|^2 & \text{if } t \text{ is bad.} \end{cases}$$

Here we think of Φ_t^b as the potential at the beginning of time step t and Φ_t^e as the potential at the end of time step t . By convention, we define $\Phi_0^e = \Phi_0^b = \|\theta\|_2^2$.

As by [Lemma 3.10](#) we have deterministic control over the discrepancy incurred during the bad times, our goal is now to control the discrepancy during the good times. Namely, we need to bound the process Y_0, Y_1, \dots, Y_T . Here, the main idea is to charge the drop in potential to the increase in discrepancy. To this end, we define the potential weighted discrepancy process $\{Z_t\}$ for $t \geq 0$ as

$$Z_t := Y_t + 4\Phi_t^e.$$

Note that $Y_0 = 0, Z_0 = 4\Phi_0^e = 4\|\theta\|_2^2$ and $Z_t \geq Y_t$ for all $t \geq 0$. For $t \in [T+1, n]$ we define $Z_t := Z_T$. Our goal is to now show that the process Z_0, \dots, Z_T has a strong “negative” drift.

The increments of $Z_t - Z_{t-1}$ for $t \in [T]$ can be expressed as follows:

$$Z_t - Z_{t-1} = Y_t - Y_{t-1} + 4(\Phi_t^e - \Phi_{t-1}^e) = (Y_t - Y_{t-1} + 4(\Phi_t^e - \Phi_t^b)) + 4(\Phi_t^b - \Phi_{t-1}^e). \quad (3.10)$$

We decompose this increment into a “predictable” part

$$\begin{aligned} X_t := Y_t - Y_{t-1} + 4(\Phi_t^e - \Phi_t^b) &= \begin{cases} \delta_t \langle v^\perp(t), \theta \rangle - 4(x_t(n(t))^2 - x_{t-1}(n(t))^2) \|\Pi_{R_t/V_t}(\theta)\|_2^2 & \text{if } t \text{ is good,} \\ 0 & \text{if } t \text{ is bad,} \end{cases} \\ &= \begin{cases} \delta_t \langle v^\perp(t), \theta \rangle - 4\delta_t(\delta_t + 2x_{t-1}(n(t))) \|\Pi_{R_t/V_t}(\theta)\|_2^2 & \text{if } t \text{ is good,} \\ 0 & \text{if } t \text{ is bad,} \end{cases} \end{aligned} \quad (3.11)$$

over which we will be able to get stochastic control, and a “free” part

$$4(\Phi_t^b - \Phi_{t-1}^e)$$

which we show is always non-positive. The following crucial lemma shows that $4(\Phi_t^b - \Phi_{t-1}^e)$ indeed gives us a “free drop” in potential.

Lemma 3.11. *For all $t \in [T]$, $\Phi_t^b \leq \Phi_{t-1}^e$. Hence, the increments satisfy $Z_t - Z_{t-1} \leq X_t$ for all $t \in [T]$.*

Proof. For $t = 1$, the statement is trivial since $\Phi_0^e = \|\theta\|_2^2$ and clearly $\Phi_1^b \leq \|\theta\|_2^2$ for all t . Thus, we may assume $t \geq 2$. If t is bad, then using $V_t \subseteq V_{t-1}$ we get $\Phi_t^b = \|\Pi_{V_t}(\theta)\|_2^2 \leq \|\Pi_{V_{t-1}}(\theta)\|_2^2 \leq \Phi_{t-1}^e$. So, we may assume from now on that t is good.

If the time step t is the beginning of a new phase, then

$$\begin{aligned}\Phi_t^b &= \|\Pi_{V_t}(\theta)\|_2^2 + (1 - x_{t-1}(n(t))^2)\|\Pi_{R_t/V_t}(\theta)\|_2^2 \\ &\leq \|\Pi_{V_t}(\theta)\|_2^2 + \|\Pi_{R_t/V_t}(\theta)\|_2^2 \\ &= \|\Pi_{R_t}(\theta)\|_2^2 \leq \|\Pi_{V_{t-1}}(\theta)\|_2^2 \leq \Phi_{t-1}^e.\end{aligned}$$

Here the inequality $\|\Pi_{R_t}(\theta)\|_2 \leq \|\Pi_{V_{t-1}}(\theta)\|_2$ follows by our assumption that t is the beginning of a new phase and hence $A_t \subseteq A_{t-1} \setminus \{n(t-1)\}$ i. e., $R_t \subseteq V_{t-1}$.

Lastly, if t is good and not the beginning of a new phase, then note that $t-1$ is also good and that $n(t) = n(t-1)$. Thus

$$\begin{aligned}\Phi_t^b &= \|\Pi_{V_t}(\theta)\|_2^2 + (1 - x_{t-1}(n(t))^2)\|\Pi_{R_t/V_t}(\theta)\|_2^2 \\ &\leq \|\Pi_{V_t}(\theta)\|_2^2 + (1 - x_{t-1}(n(t))^2)\|\Pi_{R_{t-1}/V_t}(\theta)\|_2^2 && \text{(using } R_t \subseteq R_{t-1}\text{)} \\ &= \|\Pi_{V_t}(\theta)\|_2^2 + (1 - x_{t-1}(n(t))^2) (\|\Pi_{V_{t-1}/V_t}(\theta)\|_2^2 + \|\Pi_{R_{t-1}/V_{t-1}}(\theta)\|_2^2) \\ &\leq (\|\Pi_{V_t}(\theta)\|_2^2 + \|\Pi_{V_{t-1}/V_t}(\theta)\|_2^2) + (1 - x_{t-1}(n(t))^2)\|\Pi_{R_{t-1}/V_{t-1}}(\theta)\|_2^2 \\ &= \|\Pi_{V_{t-1}}(\theta)\|_2^2 + (1 - x_{t-1}(n(t))^2)\|\Pi_{R_{t-1}/V_{t-1}}(\theta)\|_2^2 = \Phi_{t-1}^e\end{aligned}$$

as needed. \square

We now show the increment bounds X_t satisfy the negative drift requirements of [Lemma 3.2](#).

Lemma 3.12. *For $t \in [T]$, $|X_t| \leq 1$ and $\mathbb{E}[X_t + X_t^2 \mid Z_1, \dots, Z_{t-1}] \leq 0$.*

Proof. Clearly, we may assume that t is good, since otherwise $X_t = 0$ and the statement is trivial. For simplicity of notation, let Ω_{t-1} denote all the random choices made by the algorithm in the first $t-1$ time steps. Note that Ω_{t-1} determines in particular Z_1, \dots, Z_{t-1} . We adopt the shorthand $\mathbb{E}_{t-1}[\cdot] := \mathbb{E}[\cdot \mid \Omega_{t-1}]$. Let us further denote $\theta_t := \langle v^\perp(t), \theta \rangle$, $\bar{\theta}_t = \|\Pi_{R_t/V_t}(\theta)\|_2$, and $x := x_{t-1}(n(t))$. With this notation, we have that

$$X_t = \delta_t \theta_t - 4\delta_t(\delta_t + 2x)\bar{\theta}_t^2.$$

Since $v^\perp(t) \in R_t/V_t$, note that $|\theta_t| \leq \|v^\perp(t)\|_2 \bar{\theta}_t \leq \bar{\theta}_t$. By definition of t being good, we have that $\bar{\theta}_t \leq 1/8$. Since $x \in [-1, 1]$ and $x + \delta_t \in [-1, 1]$, we deduce the following simple bounds

$$|\delta_t| \leq 2, \quad |\delta_t(\delta_t + 2x)| = |(\delta_t + x)^2 - x^2| \leq 1, \quad |\delta_t + 2x| = |(\delta_t + x) + x| \leq 2. \quad (3.12)$$

Using (3.12), we see that

$$|X_t| \leq |\delta_t||\theta_t| + 4|\delta_t(\delta_t + 2x)|\bar{\theta}_t^2 \leq 2(1/8) + 4(1)(1/8^2) \leq 1.$$

Next, we have that $\mathbb{E}_{t-1}[X_t] = -4\mathbb{E}_{t-1}[\delta_t^2]\bar{\theta}_t^2$ since $\mathbb{E}_{t-1}[\delta_t] = 0$. Lastly, using the inequality $(a+b)^2 \leq 2a^2 + 2b^2$, we get that

$$\begin{aligned}X_t^2 &\leq 2\delta_t^2\theta_t^2 + 2(16)\delta_t^2(\delta_t + 2x)^2\bar{\theta}_t^4 \\ &\leq 2\delta_t^2\bar{\theta}_t^2 + 2(16)\delta_t^2(4)(1/8)^2\bar{\theta}_t^2 = 4\delta_t^2\bar{\theta}_t^2.\end{aligned}$$

Thus $\mathbb{E}_{t-1}[X_t^2] \leq 4\mathbb{E}_{t-1}[\delta_t^2]\bar{\theta}_t^2$ and hence $\mathbb{E}_{t-1}[X_t + X_t^2] \leq 0$. As Z_1, \dots, Z_{t-1} are determined by Ω_{t-1} , this implies that $\mathbb{E}[X_t + X_t^2 \mid Z_1, \dots, Z_{t-1}] \leq 0$, as needed. \square

We now prove the main moment bound.

Proof of Theorem 3.9. To begin we see that

$$\begin{aligned} \mathbb{E} [e^{\text{disc}_n}] &\leq \mathbb{E} [e^{Y_T + |\sum_{t \in \mathcal{B}} \Delta_t \text{disc}|}] \\ &\leq \mathbb{E} [e^{Y_T + 16\|\theta\|_2^2}] && \text{(by Lemma 3.10)} \\ &\leq \mathbb{E} [e^{Z_n}] e^{16\|\theta\|_2^2} && \text{(as } Y_T \leq Z_T = Z_n \text{).} \end{aligned}$$

Recall that $Z_0 = 4\|\theta\|_2^2$. Furthermore, by Lemma 3.11 we have $Z_t - Z_{t-1} \leq X_t$, by Lemma 3.12 we have $|X_t| \leq 1$ and $\mathbb{E} [X_t + X_t^2 \mid Z_1, \dots, Z_{t-1}] \leq 0$ for all $t \in [T]$, and for $t \in [T + 1, n]$ we have $Z_t - Z_{t-1} = 0$ by definition. Therefore, applying Lemma 3.2 gives that

$$\mathbb{E} [e^{Z_n}] \leq e^{Z_0} = e^{4\|\theta\|_2^2}.$$

Thus, combining everything together, we get $\mathbb{E} [e^{\text{disc}_n}] \leq \mathbb{E} [e^{Z_n}] e^{16\|\theta\|_2^2} \leq e^{20\|\theta\|_2^2}$ as needed. \square

4 Applications

In this section, we list some applications of our main result.

4.1 The ℓ_p variant of the Komlós problem

Corollary 1.5 (restated). *There is an efficient randomized algorithm which given an $m \times n$ matrix A having all columns of ℓ_2 -norm at most one and $p \in [1, \infty)$, finds a coloring $x \in \{-1, 1\}^n$ with expected ℓ_p discrepancy $O(\sqrt{p})$.*

Proof. Let $Y = Ax$. By Theorem 1.4, we get that Y is a σ -sub-Gaussian random vector and hence each component of Y is a σ -sub-Gaussian random variable. Letting Y_i denote the i -th component of Y , we get

$$\mathbb{E} [\|Y\|_p^p] = \sum_{i=1}^m \mathbb{E} [|Y_i|^p] \leq mC^p p^{p/2}$$

for a constant $C = C(\sigma)$ depending only on σ . The inequality above follows from the standard fact that the p -th moment of an $O(1)$ -sub-Gaussian random variable is at most $C^p p^{p/2}$. \square

4.2 Discrepancy relative to the γ_2 -norm

Corollary 1.7 (restated). *There exists an efficient randomized algorithm that, given any $m \times n$ matrix A and $J \subseteq [n]$, returns a coloring $x : J \rightarrow \{-1, 1\}$ such that with constant probability,*

$$\|A_{Jx}\|_\infty = O(\gamma_2(A) \sqrt{\log m}).$$

Proof. It was shown in [28] that for any matrix A , there exists an efficiently computable quantity $\gamma_2(A)$ such that

$$\Omega(\gamma_2(A)/\log m) \leq \text{herdisc}(A) \leq O(\gamma_2(A)\sqrt{\log m}).$$

The upper bound above was proved using [Theorem 1.1](#) (roughly) as follows: given any set $J \subseteq [n]$, factorize the matrix $A|_J$ as $A|_J = BC$ where every row of the matrix B has ℓ_2 -norm at most $\gamma_2(A)$ and every column of the matrix C has ℓ_2 -norm at most one. Such a factorization exists and can be computed by solving an appropriate semidefinite program. We refer the reader to [28] for more details.

Then using the matrix B , a convex body K is defined as follows:

$$K = \{y \in \mathbb{R}^m : \|By\|_\infty \leq c\gamma_2(A)\sqrt{\log m}\}.$$

For c a large enough constant, $\gamma_m(K) \geq 1/2$. This follows by standard Gaussian tail bounds and union bound. Now finding a coloring $x \in \{-1, 1\}^n$ of the columns of $A|_J$ such that the ℓ_∞ -norm of $A|_J x = BCx$ is $O(\gamma_2(A)\sqrt{\log m})$ is equivalent to finding a coloring x of the column vectors of C such that Cx lies in K . As C has all columns of length at most one, by [Theorem 1.1](#) there exists a coloring x such that the discrepancy of $A|_J$ is $O(\gamma_2(A)\sqrt{\log m})$.

Now [Theorem 1.4](#) (with [Theorem 1.2](#)) directly gives an efficient algorithm to find such a coloring. \square

4.3 A generalization of Banaszczyk’s result

In this subsection we prove the following generalization of Banaszczyk’s result. The proof follows along similar lines as the proof of [24] showing that linear discrepancy is at most twice the hereditary discrepancy.

Theorem 1.8 (restated). *Let S_1, S_2, \dots, S_n be sets such that for each $i \in [n]$, $S_i \subseteq \mathbb{B}_2^m$ and 0 lies in the convex hull of S_i . Then for any convex body K with $\gamma_m(K) \geq 1/2$, there exist vectors $v_i \in S_i$ such that $\sum_{i=1}^n v_i \in cK$, where $c > 0$ is an absolute constant. Moreover, there is an efficient algorithm to find these vectors.*

Proof. For technical reasons, we give the proof for the case when $\gamma_m(K) \geq 1/2 + \varepsilon$ for some $\varepsilon > 0$. The running time of the algorithm will be proportional to $\log(1/\varepsilon)$. The general case, for instance when $\gamma_m(K) = 1/2$, is slightly more complicated but can be handled by combining our proof with the techniques in [15]. We provide a sketch of how to do this towards the end of the proof.

For each $i \in [n]$, as 0 lies in the convex hull of S_i , there exist at most $m + 1$ vectors in S_i and a convex combination of them that equals 0 . That is, there exist $m + 1$ vectors $v_{i,j} \in S_i$ and real numbers $x_{i,j} \geq 0$ such that

$$\forall i \in [n] : \sum_{j=1}^{m+1} x_{i,j} = 1 \quad \text{and} \quad \sum_{j=1}^{m+1} x_{i,j} v_{i,j} = 0. \tag{4.1}$$

Our goal will be to round each collection $\{x_{i,j} : j \in [m + 1]\}$ such that exactly one of them is 1 and the rest are 0, without incurring too much discrepancy.

Assume for now that each of $x_{i,j}$ can be expressed using at most k digits in binary, for some finite k ; that is, $x_{i,j} \in 2^{-k}\mathbb{Z}$ for all i, j . We will see later how to get rid of this assumption. The main step will

to be to reduce k to $k - 1$ by rounding the k -th bits in each $x_{i,j}$ either up or down, in such a way that discrepancy stays bounded. Then, we will repeat this operation until we obtain $k = 0$, which means that $x_{i,j} \in \{0, 1\}$. We will maintain the property that $x_{i,j} \geq 0$ and that $\sum_j x_{i,j} = 1$ for all i , which then implies that there exists $j_i \in [m + 1]$ for each i such that $x_{i,j_i} = 1$ and $x_{i,j} = 0$ if $j \neq j_i$. The claim will follow by choosing $v_{i,j_i} \in S_i$.

Setting notation, let $x_{i,j}^{(k)} = x_{i,j}$. Assume that we already computed for $\ell \in [k]$ a choice of

$$x_{i,j}^{(\ell)} \in 2^{-\ell} \mathbb{Z}$$

that satisfies $\sum_j x_{i,j}^{(\ell)} = 1$ for all $i \in [n]$. If $x_{i,j}^{(\ell)} \in \{0, 1\}$ for all j then there is nothing more to do. Otherwise, let

$$x_{i,j}^{(\ell)} = 0.b_{i,j,1}^{(\ell)} b_{i,j,2}^{(\ell)} \dots b_{i,j,\ell}^{(\ell)}$$

be the binary expansion of $x_{i,j}^{(\ell)} < 1$. Let

$$A^{(\ell)} = \{(i, j) \in [n] \times [m + 1] : b_{i,j,\ell}^{(\ell)} = 1\}$$

be the indices where the ℓ -th bit of $x_{i,j}$ is 1. We will construct a coloring $\chi^{(\ell)} : A^{(\ell)} \rightarrow \{-1, 1\}$ which satisfies

$$\sum_{(i,j) \in A^{(\ell)}} \chi^{(\ell)}(i, j) v_{i,j} \in cK \quad \text{and} \quad \forall i \in [n], \quad \sum_{j: (i,j) \in A^{(\ell)}} \chi^{(\ell)}(i, j) = 0 \quad (4.2)$$

for some absolute constant $c > 0$. Given such a coloring $\chi^{(\ell)}$, we compute the value of $x_{i,j}^{(\ell-1)}$ as follows:

$$x_{i,j}^{(\ell-1)} = \begin{cases} x_{i,j}^{(\ell)} & \text{if } b_{i,j,\ell}^{(\ell)} = 0, \\ x_{i,j}^{(\ell)} + 2^{-\ell} & \text{if } b_{i,j,\ell}^{(\ell)} = 1 \text{ and } \chi^{(\ell)}(i, j) = 1, \\ x_{i,j}^{(\ell)} - 2^{-\ell} & \text{if } b_{i,j,\ell}^{(\ell)} = 1 \text{ and } \chi^{(\ell)}(i, j) = -1. \end{cases}$$

Observe that this operation zeros the ℓ -th bit of all $x_{i,j}^{(\ell)}$, namely $x_{i,j}^{(\ell-1)} \in 2^{-(\ell-1)} \mathbb{Z}$; it maintains the property that for all $i \in [n]$, $\sum_j x_{i,j}^{(\ell-1)} = 1$ and $x_{i,j}^{(\ell-1)} \geq 0$; and it satisfies that

$$\sum_{i,j} \left(x_{i,j}^{(\ell-1)} - x_{i,j}^{(\ell)} \right) v_{i,j} \in 2^{-\ell} cK.$$

Thus, repeating this rounding operation for $\ell = k, k - 1, \dots, 1$ will result in a choice of $x_{i,j} \in \{0, 1\}$ that satisfy $\sum_j x_{i,j} = 1$ and

$$\sum_{i,j} x_{i,j} v_{i,j} \in cK.$$

That is, there is a choice of $j_i \in [m + 1]$ for all $i \in [n]$ such that $\sum_{i,j} v_{i,j_i} \in cK$, as claimed.

It remains to show how to find a coloring satisfying (4.2). Let $A_i^{(\ell)} = \{j : (i, j) \in A^{(\ell)}\}$ be the elements being colored in the set S_i when we round the ℓ -th bits. We claim that $|A_i^{(\ell)}|$ must be even. This is since for every $i \in [n]$, the number of elements $x_{i,j}^{(\ell)}$ for which $b_{i,j,\ell}^{(\ell)} = 1$ must be even, as $\sum x_{i,j}^{(\ell)} = 1$. We will pair

up the elements of $A_i^{(\ell)}$ in an arbitrary way, and only consider colorings $\chi^{(\ell)}$ which give opposite colors to elements in each pair. In such a way, such a coloring automatically satisfies that $\sum_j \chi^{(\ell)}(i, j) = 0$ for all $i \in [n]$.

For simplicity of notation, denote the vectors $\{v_{i,j} : j \in A_i^{(\ell)}\}$ by $\{u_{i,1}, \dots, u_{i,2q_i}\}$ for some integer $q_i = |A_i^{(\ell)}|/2$. Define new vectors $w_{i,j} = (u_{i,2j-1} - u_{i,2j})/2$ for $j \in [q_i]$. Observe that $\|w_{i,j}\|_2 \leq 1$. Apply [Theorem 1.1](#) to the vectors $w_{i,j}$ and the convex body K . This gives a coloring $\chi'(i, j)$ for each vector $w_{i,j}$ such that

$$\sum_{i,j} \chi'(i, j) w_{i,j} \in cK.$$

We now define the coloring $\chi^{(\ell)}$ to give the color $\chi'(i, j)$ to $u_{i,2j-1}$ and the color $-\chi'(i, j)$ to $u_{i,2j}$. Clearly this satisfies both the conditions in [\(4.2\)](#) with constant $2c$.

We will now show that we can assume the binary expansion to be finite. Concretely, we will show that a preliminary rounding procedure can allow us to assume that $k \leq \log(2mn/\varepsilon)$. This is since by truncating each $x_{i,j}$ after $\log(2mn/\varepsilon)$ bits, the sum $\sum_{i,j} x_{i,j} v_{i,j}$ changes by at most

$$\sum_{i,j} \frac{1}{2^{\log(2mn/\varepsilon)}} v_{i,j} \in \varepsilon \mathbb{B}_2^m \subseteq K$$

and thus increases the final value of c by at most 1. The last containment follows by our assumption that $\gamma_m(K) \geq 1/2 + \varepsilon$, and thus K must contain a Euclidean ball of radius r which satisfies $\gamma_1([0, r]) \leq \varepsilon$. Clearly this is true for $r = \varepsilon$.

We mention briefly now on how to tackle the case when $\gamma_m(K) < 1/2 + \varepsilon$. This proceeds along similar lines as [Theorem 40](#) from [\[15\]](#). The main idea is that we can find a point p such that

$$p \in K \cap \left(\sum_i \text{conv}(S_i) \right)$$

where S_i denotes the convex hull of S_i and the summation operator used is Minkowski addition. We then instead solve a new problem on the instance given by the convex body $K' := \alpha(K - p)$ and sets S'_i such that $\sum_i \text{conv}(S'_i) = \alpha(\sum_i \text{conv}(S_i) - p)$ for some constant scaling factor $\alpha > 0$. A solution of our original problem is recoverable from a solution of this. p and α moreover satisfy the property that $\gamma_m(K') \geq 1/2 + \varepsilon$, and we already know how to solve this case. \square

5 Conclusion and open questions

We gave efficient algorithms for several problems to find colorings with discrepancy bounds similar to those achievable using Banaszczyk’s result, [Theorem 1.1](#). However there are still some problems that use Banaszczyk’s techniques in a non-trivial iterative way, for which we are unable to obtain an efficient algorithm.

One such problem is Tusnády’s problem concerning the discrepancy of axis-parallel boxes in \mathbb{R}^d . Nikolov [\[31\]](#) used Banaszczyk’s technique to prove that the discrepancy is $O_d(\log^{d-1/2} n)$, where $O_d(\cdot)$

hides factors depending only on d . Our techniques do not seem to apply here and the best known algorithmic bound is $O_d(\log^d n)$ [7].

Another such problem is the Steinitz problem in the ℓ_2 -norm. Here we are given n vectors $v_1, \dots, v_n \in \mathbb{B}_2^m$ such that $\sum_i v_i = 0$, and the goal is to find a rearrangement of these vectors such that the ℓ_2 -norm of the sum of vectors in any prefix along the rearrangement is small. That is, we want to find a permutation $\pi : [n] \rightarrow [n]$ to minimize

$$\max_{k=1, \dots, n} \left\| \sum_{i=1}^k v_{\pi(i)} \right\|_2.$$

Banaszczyk showed in [2], using techniques from [1], that there exists a permutation for which the above expression is at most $O(\sqrt{m} + \sqrt{\log n})$, whereas the best known algorithmic bound is $O(\sqrt{m \log n})$ [7].

Acknowledgments

The authors would like to thank Aleksandar Nikolov for several discussions related to the topic of this paper. N. B. and S. L. would like to thank the Dagstuhl seminar “Computational Complexity of Discrete Problems” where this work was initiated. We would also like to thank the referees for various comments that improved the presentation of the paper.

References

- [1] WOJCIECH BANASZCZYK: Balancing vectors and Gaussian measures of n -dimensional convex bodies. *Random Structures Algorithms*, 12(4):351–360, 1998. [[doi:10.1002/\(SICI\)1098-2418\(199807\)12:4<351::AID-RSA3>3.0.CO;2-S](https://doi.org/10.1002/(SICI)1098-2418(199807)12:4<351::AID-RSA3>3.0.CO;2-S)] 3, 5, 22
- [2] WOJCIECH BANASZCZYK: On series of signed vectors and their rearrangements. *Random Structures Algorithms*, 40(3):301–316, 2012. [[doi:10.1002/rsa.20373](https://doi.org/10.1002/rsa.20373)] 3, 22
- [3] NIKHIL BANSAL: Constructive algorithms for discrepancy minimization. In *Proc. 51st FOCS*, pp. 3–10. IEEE Comp. Soc. Press, 2010. [[doi:10.1109/FOCS.2010.7](https://doi.org/10.1109/FOCS.2010.7), [arXiv:1002.2259](https://arxiv.org/abs/1002.2259)] 2
- [4] NIKHIL BANSAL, MOSES CHARIKAR, RAVISHANKAR KRISHNASWAMY, AND SHI LI: Better algorithms and hardness for broadcast scheduling via a discrepancy approach. In *Proc. 25th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA’14)*, pp. 55–71. SIAM, 2014. [[doi:10.1137/1.9781611973402.5](https://doi.org/10.1137/1.9781611973402.5)] 2
- [5] NIKHIL BANSAL, DANIEL DADUSH, AND SHASHWAT GARG: An algorithm for Komlós conjecture matching Banaszczyk’s bound. *SIAM J. Comput.*, 48(2):534–553, 2019. Preliminary version in *FOCS’16*. [[doi:https://doi.org/10.1137/17M1126795](https://doi.org/10.1137/17M1126795)] 3, 5
- [6] NIKHIL BANSAL, DANIEL DADUSH, SHASHWAT GARG, AND SHACHAR LOVETT: The Gram–Schmidt walk: A cure for the Banaszczyk blues. In *Proc. 50th STOC*, pp. 587–597. ACM Press, 2018. [[doi:10.1145/3188745.3188850](https://doi.org/10.1145/3188745.3188850)]

- [7] NIKHIL BANSAL AND SHASHWAT GARG: Algorithmic discrepancy beyond partial coloring. In *Proc. 49th STOC*, pp. 914–926. ACM Press, 2017. [doi:10.1145/3055399.3055490, arXiv:1611.01805] 3, 5, 22
- [8] NIKHIL BANSAL AND VISWANATH NAGARAJAN: Approximation-friendly discrepancy rounding. In *Proc. 18th Ann. Conf. on Integer Programming and Combinatorial Optimization (IPCO '16)*, pp. 375–386. Springer, 2016. [doi:10.1007/978-3-319-33461-5_31, arXiv:1512.02254] 2
- [9] IMRE BÁRÁNY: On the power of linear dependencies. In *Building Bridges* (M. Grötschel and G. O. H. Katona, eds.), volume 19 of *Bolyai Soc. Math. Studies*, pp. 31–45. Springer, 2008. [doi:10.1007/978-3-540-85221-6_1] 2, 6
- [10] FRANCK BARTHE, OLIVIER GUÉDON, SHAHAR MENDELSON, AND ASSAF NAOR: A probabilistic approach to the geometry of the ℓ_p^n -ball. *Ann. Probab.*, 33(2):480–513, 2005. [doi:10.1214/009117904000000874] 3
- [11] JÓZSEF BECK: Roth’s estimate of the discrepancy of integer sequences is nearly sharp. *Combinatorica*, 1(4):319–325, 1981. [doi:10.1007/BF02579452] 2
- [12] JÓZSEF BECK AND TIBOR FIALA: “Integer-making” theorems. *Discr. Appl. Math.*, 3(1):1–8, 1981. [doi:10.1016/0166-218X(81)90022-6] 2, 3
- [13] MOSES CHARIKAR, ALANTHA NEWMAN, AND ALEKSANDAR NIKOLOV: Tight hardness results for minimizing discrepancy. In *Proc. 22nd Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'11)*, pp. 1607–1614. SIAM, 2011. [doi:10.1137/1.9781611973082.124] 6
- [14] BERNARD CHAZELLE: *The Discrepancy Method: Randomness and Complexity*. Cambridge Univ. Press, 2000. [doi:10.1017/CBO9780511626371] 2
- [15] DANIEL DADUSH, SHASHWAT GARG, SHACHAR LOVETT, AND ALEKSANDAR NIKOLOV: Towards a constructive version of Banaszczyk’s vector balancing theorem. *Theory of Computing*, 15(15):1–58, 2019. Preliminary version in *RANDOM'16*. [doi:10.4086/toc.2019.v015a015] 3, 4, 5, 19, 21
- [16] RONEN ELDAN AND MOHIT SINGH: Efficient algorithms for discrepancy minimization in convex sets. *Random Structures Algorithms*, 53(2):289–307, 2018. [doi:10.1002/rsa.20763] 2
- [17] DAVID A. FREEDMAN: On tail probabilities for martingales. *Ann. Probab.*, 3(1):100–118, 1975. [doi:10.1214/aop/1176996452] 8
- [18] RAJIV GANDHI, SAMIR KHULLER, SRINIVASAN PARTHASARATHY, AND ARAVIND SRINIVASAN: Dependent rounding and its applications to approximation algorithms. *J. ACM*, 53(3):324–360, 2006. Preliminary version in *FOCS'02*. [doi:10.1145/1147954.1147956] 8
- [19] EFIM DAVYDOVICH GLUSKIN: Extremal properties of orthogonal parallelepipeds and their applications to the geometry of Banach spaces. *Sbornik: Mathematics*, 64(1):85–96, 1989. [doi:10.1070/SM1989v064n01ABEH003295] 2

- [20] NICHOLAS J. A. HARVEY, ROY SCHWARTZ, AND MOHIT SINGH: Discrepancy without partial colorings. In *Proc. 17th Internat. Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX'14)*, volume 28 of *LIPICs*, pp. 258–273. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2014. [doi:10.4230/LIPICs.APPROX-RANDOM.2014.258] 2
- [21] KASPER GREEN LARSEN: On range searching in the group model and combinatorial discrepancy. *SIAM J. Comput.*, 43(2):673–686, 2014. Preliminary version in *FOCS'11*. [doi:10.1137/120865240] 3
- [22] LAP-CHI LAU, R. RAVI, AND MOHIT SINGH: *Iterative Methods in Combinatorial Optimization*. Cambridge Univ. Press, 2011. [doi:10.1017/CBO9780511977152] 2
- [23] AVI LEVY, HARISHCHANDRA RAMADAS, AND THOMAS ROTHVOSS: Deterministic discrepancy minimization via the multiplicative weight update method. In *Proc. 19th Ann. Conf. on Integer Programming and Combinatorial Optimization (IPCO '17)*, pp. 380–391. Springer, 2017. [doi:10.1007/978-3-319-59250-3_31, arXiv:1611.08752] 3, 5
- [24] LÁSZLÓ LOVÁSZ, JOEL H. SPENCER, AND KATALIN VESZTERGOMBI: Discrepancy of set-systems and matrices. *European J. Combinatorics*, 7(2):151–160, 1986. [doi:10.1016/S0195-6698(86)80041-5] 19
- [25] SHACHAR LOVETT AND RAGHU MEKA: Constructive discrepancy minimization by walking on the edges. *SIAM J. Comput.*, 44(5):1573–1582, 2015. Preliminary version in *FOCS'12*. [doi:10.1137/130929400, arXiv:1203.5747] 2
- [26] JIŘÍ MATOUŠEK: An L_p version of the Beck-Fiala conjecture. *European J. Combinatorics*, 19(2):175–182, 1998. [doi:10.1006/eujc.1997.0162] 5
- [27] JIŘÍ MATOUŠEK: *Geometric Discrepancy: An Illustrated Guide*. Springer, 2009. [doi:10.1007/978-3-642-03942-3] 2
- [28] JIŘÍ MATOUŠEK, ALEKSANDAR NIKOLOV, AND KUNAL TALWAR: Factorization norms and hereditary discrepancy. *Internat. Mathematics Research Notices*, rny033, 2018. [doi:10.1093/imrn/rny033, arXiv:1408.1376] 3, 6, 19
- [29] ALEKSANDAR NIKOLOV: The Komlós conjecture holds for vector colorings, 2013. [arXiv:1301.4039] 5
- [30] ALEKSANDAR NIKOLOV: *New Computational Aspects of Discrepancy Theory*. Ph.D. thesis, Rutgers University, 2014. [doi:10.7282/T3RN3749] 2, 3
- [31] ALEKSANDAR NIKOLOV: Tighter bounds for the discrepancy of boxes and polytopes. *Mathematika*, 63(3):1091–1113, 2017. [doi:10.1112/S0025579317000250, arXiv:1701.05532] 3, 21
- [32] ALEKSANDAR NIKOLOV, KUNAL TALWAR, AND LI ZHANG: The geometry of differential privacy: The small database and approximate cases. *SIAM J. Comput.*, 45(2):575–616, 2016. Preliminary version in *STOC'13*. [doi:10.1137/130938943, arXiv:1212.0297] 2

- [33] THOMAS ROTHVOSS: Better bin packing approximations via discrepancy theory. *SIAM J. Comput.*, 45(3):930–946, 2016. Preliminary version in FOCS’13. [doi:10.1137/140955367, arXiv:1301.4010] [2](#)
- [34] THOMAS ROTHVOSS: Constructive discrepancy minimization for convex sets. *SIAM J. Comput.*, 46(1):224–234, 2017. Preliminary version in FOCS’14. [doi:10.1137/141000282, arXiv:1404.0339] [2, 3](#)
- [35] JOEL H. SPENCER: Six standard deviations suffice. *Trans. Amer. Math. Soc.*, 289(2):679–706, 1985. [doi:10.1090/S0002-9947-1985-0784009-0] [2](#)
- [36] JOEL H. SPENCER: *Ten Lectures on the Probabilistic Method*. SIAM, 1987. [doi:10.1137/1.9781611970074] [5](#)
- [37] ARAVIND SRINIVASAN: Distributions on level-sets with applications to approximation algorithms. In *Proc. 42nd FOCS*, pp. 588–597. IEEE Comp. Soc. Press, 2001. [doi:10.1109/SFCS.2001.959935] [8](#)
- [38] MICHEL TALAGRAND: *The Generic Chaining: Upper and Lower Bounds of Stochastic Processes*. Springer, 2005. [doi:10.1007/3-540-27499-5] [4](#)

AUTHORS

Nikhil Bansal
 Researcher
 Centrum voor Wiskunde en Informatica
 Amsterdam, the Netherlands
 bansal@gmail.com
<http://www.win.tue.nl/~nikhil>

Daniel Dadush
 Researcher
 Centrum voor Wiskunde en Informatica
 Amsterdam, the Netherlands
 dndadush@gmail.com
<https://homepages.cwi.nl/~dadush/>

Shashwat Garg
Quantitative researcher
WorldQuant LLC
Budapest, Hungary
garg.shashwat@gmail.com
<https://www.win.tue.nl/~sgarg/>

Shachar Lovett
Associate professor
University of California, San Diego
slovett@cse.ucsd.edu
<http://cseweb.ucsd.edu/~slovett>

ABOUT THE AUTHORS

NIKHIL BANSAL is a researcher at the [Centrum voor Wiskunde en Informatica](#), Amsterdam. He attended the [Indian Institute of Technology, Mumbai](#) for his B. Tech. degree, and received his Ph. D. from [Carnegie Mellon University](#), where he was advised by [Avrim Blum](#).

He got fascinated by algorithms while taking an undergraduate course by [Prof. Ajit A. Diwan](#). Since then he has enjoyed thinking about various kinds of algorithmic questions.

DANIEL DADUSH is a tenured researcher at the [Centrum voor Wiskunde en Informatica](#), Amsterdam. He earned his Ph. D. in algorithms, combinatorics and optimization (ACO) from [Georgia Tech](#) in 2012, where his advisor was [Santosh Vempala](#). Before joining CWI, he spent two years as a Simons postdoctoral fellow in the Computer Science Department at [New York University](#). His research has focused on algorithms for lattice problems, integer programming, and high-dimensional convex geometry. He lives and works in Amsterdam, and is glad that, thus far, the city remains above water. He enjoys reading the New York Times, listening to NPR, and taking long bike rides along the canals when it is not raining.

SHASHWAT GARG is a Quantitative Researcher at WorldQuant LLC. He got his Ph. D. in algorithms from [Eindhoven University of Technology](#) in 2019 where his advisor was [Nikhil Bansal](#). His research focused on algorithms for combinatorial discrepancy and approximation algorithms for scheduling problems. He lives in Budapest, Hungary. He enjoys reading anything under the sun, exploring new cafés and painting.

THE GRAM–SCHMIDT WALK: A CURE FOR THE BANASZCZYK BLUES

SHACHAR LOVETT graduated from the [Weizmann Institute of Science](#) in 2010; his advisors were [Omer Reingold](#) and [Ran Raz](#). He was a member of the Institute for Advanced Study, School of Mathematics between 2010-2012. Since then, he has been a faculty member at the [University of California, San Diego](#). He is interested in the role that structure and randomness play in computation and mathematics, and in particular in computational complexity, coding theory, pseudorandomness, and algebraic constructions.