

Online Discrepancy Minimization for Stochastic Arrivals

Nikhil Bansal* Haotian Jiang† Raghu Meka‡ Sahil Singla§ Makrand Sinha¶

Abstract

In the stochastic online vector balancing problem, vectors v_1, v_2, \dots, v_T chosen independently from an arbitrary distribution in \mathbb{R}^n arrive one-by-one and must be immediately given a \pm sign. The goal is to keep the norm of the discrepancy vector, *i.e.*, the signed prefix-sum, as small as possible for a given target norm.

We consider some of the most well-known problems in discrepancy theory in the above online stochastic setting, and give algorithms that match the known offline bounds up to $\text{polylog}(nT)$ factors. This substantially generalizes and improves upon the previous results of Bansal, Jiang, Singla, and Sinha (STOC' 20). In particular, for the Komlós problem where $\|v_t\|_2 \leq 1$ for each t , our algorithm achieves $\tilde{O}(1)$ discrepancy with high probability, improving upon the previous $\tilde{O}(n^{3/2})$ bound. For Tusnády's problem of minimizing the discrepancy of axis-aligned boxes, we obtain an $O(\log^{d+4} T)$ bound for arbitrary distribution over points. Previous techniques only worked for product distributions and gave a weaker $O(\log^{2d+1} T)$ bound. We also consider the Banaszczyk setting, where given a symmetric convex body K with Gaussian measure at least $1/2$, our algorithm achieves $\tilde{O}(1)$ discrepancy with respect to the norm given by K for input distributions with sub-exponential tails.

Our results are based on a new potential function approach. Previous techniques consider a potential that penalizes large discrepancy, and greedily chooses the next color to minimize the increase in potential. Our key idea is to introduce a potential that also enforces constraints on how the discrepancy vector evolves, allowing us to maintain certain anti-concentration properties. We believe that our techniques to control the evolution of states could find other applications in stochastic processes and online algorithms. For the Banaszczyk setting, we further enhance this potential by combining it with ideas from generic chaining. Finally, we also extend these results to the setting of online multi-color discrepancy.

*CWI Amsterdam and TU Eindhoven, N.Bansal@cwi.nl. Supported by the ERC Consolidator Grant 617951 and the NWO VICI grant 639.023.812.

†Paul G. Allen School of CSE, University of Washington, jhtdavid@cs.washington.edu. Supported in part by the National Science Foundation, Grant Number CCF-1749609, CCF-1740551, DMS-1839116.

‡Department of Computer Science, University of California, Los Angeles, raghum@cs.ucla.edu. Supported by NSF Grant CCF-1553605. Part of the work was done while visiting CWI Amsterdam.

§Institute for Advanced Study and Princeton University, singla@cs.princeton.edu. Supported in part by the Schmidt Foundation. Part of the work done while visiting CWI Amsterdam.

¶CWI Amsterdam, makrand@cwi.nl. Supported by the Netherlands Organization for Scientific Research, Grant Number 617.001.351 and the NWO VICI grant 639.023.812.

1 Introduction

We consider the following online vector balancing question, originally proposed by Spencer [Spe77]: vectors $v_1, v_2, \dots, v_T \in \mathbb{R}^n$ arrive online, and upon the arrival of v_t , a sign $\chi_t \in \{\pm 1\}$ must be chosen irrevocably, so that the ℓ_∞ -norm of the *discrepancy vector* (signed sum) $d_t := \chi_1 v_1 + \dots + \chi_t v_t$ remains as small as possible. That is, find the smallest B such that $\max_{t \in [T]} \|d_t\|_\infty \leq B$. More generally, one can consider the problem of minimizing $\max_{t \in T} \|d_t\|_K$ with respect to arbitrary norms given by a symmetric convex body K .

Offline setting. The offline version of the problem, where the vectors v_1, \dots, v_T are given in advance, has been extensively studied in discrepancy theory, and has various applications [Mat09, Cha01, CST⁺14]. Here we study three important problems in this vein:

Tusnády's problem. Given points $x_1, \dots, x_T \in [0, 1]^d$, we want to assign \pm signs to the points, so that for every axis-parallel box, the difference between the number of points inside the box that are assigned a plus sign and those assigned a minus sign is minimized.

Beck-Fiala and Komlós problem. Given $v_1, \dots, v_T \in \mathbb{R}^n$ with Euclidean norm at most one, we want to minimize $\max_{t \in T} \|d_t\|_\infty$. After scaling, a special case of the Komlós problem is the Beck-Fiala setting where $v_1, \dots, v_T \in [-1, 1]^n$ are s -sparse (with at most s non-zeros).

Banaszczyk's problem. Given $v_1, \dots, v_T \in \mathbb{R}^n$ with Euclidean norm at most one, and a convex body $K \in \mathbb{R}^n$ with Gaussian measure¹ $\gamma_n(K) \geq 1 - 1/(2T)$, find the smallest B so that there exist signs such that $d_t \in B \cdot K$ for all $t \in [T]$.

One of the most general and powerful results here is due to Banaszczyk [Ban12]: there exist signs such that $d_t \in O(1) \cdot K$ for all $t \in [T]$ for any convex body

¹The Gaussian measure $\gamma_n(S)$ of a set $S \subseteq \mathbb{R}^n$ is defined as $\mathbb{P}[G \in S]$ where G is standard Gaussian in \mathbb{R}^n .

$K \in \mathbb{R}^n$ with Gaussian measure² $\gamma_n(K) \geq 1 - 1/(2T)$. In particular, this gives the best known bounds of $O((\log T)^{1/2})$ for the Komlós problem; for the Beck-Fiala setting, when the vectors are s -sparse, the bound is $O((s \log T)^{1/2})$.

An extensively studied case, where sparsity plays a key role, is that of Tusnády's problem (see [Mat09] for a history), where the best known (non-algorithmic) results, building on a long line of work, are an $O(\log^{d-1/2} T)$ upper bound of [Nik17] and an almost matching $\Omega(\log^{d-1} T)$ lower bound of [MN15].

In general, several powerful techniques have been developed for offline discrepancy problems over the last several decades, starting with initial non-constructive approaches such as [Bec81, Spe85, Glu89, Gia97, Ban98, Ban12], and more recent algorithmic ones such as [Ban10, LM15, Rot14, MNT14, BDG16, LRR17, ES18, BDGL18, DNTT18]. However, none of them applies to the online setting that we consider here.

Online setting. A naïve algorithm is to pick each sign χ_t randomly and independently, which by standard tail bounds gives $B = \Theta((T \log n)^{1/2})$ with high probability. In typical interesting settings, we have $T \geq \text{poly}(n)$, and hence a natural question is whether the dependence on T can be improved from $T^{1/2}$ to say, poly-logarithmic in T , and ideally to even match the known offline bounds.

Unfortunately, the $\Omega(T^{1/2})$ dependence is necessary if the adversary is adaptive³: at each time t , the adversary can choose the next input vector v_t to be orthogonal to d_{t-1} , causing $\|d_t\|_2$ to grow as $\Omega(T^{1/2})$ (see [Spe87] for an even stronger lower bound). Even for very special cases, such as for vectors in $\{-1, 1\}^n$, strong $\Omega(2^n)$ lower bounds are known [Bár79]. Hence, we focus on a natural *stochastic* model where we relax the power of the adversary and assume that the arriving vectors are chosen in an i.i.d. manner from some—possibly adversarially chosen—distribution \mathbf{p} . In this case, one could hope to exploit that $\langle d_{t-1}, v_t \rangle$ is not always zero, *e.g.*, due to anti-concentration properties of the input distribution, and beat the $\Omega(T^{1/2})$ bound.

Recently, Bansal and Spencer [BS19], considered the special case where \mathbf{p} is the uniform distribution on all $\{-1, 1\}^n$ vectors, and gave an almost optimal

²We remark that if one only cares about the final discrepancy d_T , the condition in Banaszczyk's result can be improved to $\gamma_n(K) \geq 1/2$ (though, in all applications we are aware of, this makes no difference if $T = \text{poly}(n)$ and makes a difference of at most $\sqrt{\log T}$ for general T).

³In the sense that the adversary can choose the next vector v_t based on the current discrepancy vector d_{t-1} .

$O(n^{1/2} \log T)$ bound for the ℓ_∞ norm that holds with high probability for all $t \in [T]$. The setting of general distributions \mathbf{p} turns out to be harder and was considered recently by [JKS19] and [BJSS20], motivated by *envy minimization* problems and an online version of Tusnády's problem. The latter was also considered independently by Dwivedi, Feldheim, Gurel-Gurevich, and Ramadas [DFGGR19] motivated by the problem of placing points uniformly in a grid.

For an arbitrary distribution \mathbf{p} supported on vectors in $[-1, 1]^n$, [BJSS20] give an algorithm achieving an $O(n^2 \log T)$ bound for the ℓ_∞ -norm. In contrast, the best offline bound is $O((n \log T)^{1/2})$, and hence $\tilde{\Omega}(n^{3/2})$ factor worse, where $\tilde{\Omega}(\cdot)$ ignores poly-logarithmic factors in n and T .

More significantly, the existing bounds for the online version are much worse than those of the offline version for the case of s -sparse vectors (*Beck-Fiala* setting) — [BJSS20] obtain a much weaker bound of $O(sn \log T)$ for the online setting while the offline bound of $O((s \log T)^{1/2})$ is independent of the ambient dimension n . These technical limitations also carry over to the online Tusnády problem, where previous works [JKS19, DFGGR19, BJSS20] could only handle product distributions.

To this end, [BJSS20] propose two key problems in the i.i.d. setting. First, for a general distribution \mathbf{p} on vectors in $[-1, 1]^n$, can one get an optimal $\tilde{O}(n^{1/2})$ or even $\tilde{O}(n)$ dependence? Second, can one get poly($s, \log T$) bounds when the vectors are s -sparse. In particular, as a special case, can one get $(\log T)^{O(d)}$ bounds for the Tusnády problem, when points arrive from an *arbitrary* non-product distribution on $[0, 1]^d$.

1.1 Our Results In this paper we resolve both the above questions of [BJSS20], and prove much more general results that obtain bounds within poly-logarithmic factors of those achievable in the offline setting.

Online Komlós and Tusnády settings. We first consider Komlós' setting for online discrepancy minimization where the vectors have ℓ_2 -norm at most 1. Recall, the best known offline bound in this setting is $O((\log T)^{1/2})$ [Ban12]. We achieve the same result, up to poly-logarithmic factors, in the online setting.

THEOREM 1.1. (ONLINE KOMLÓS SETTING) *Let \mathbf{p} be a distribution in \mathbb{R}^n supported on vectors with Euclidean norm at most 1. Then, for vectors v_1, \dots, v_T sampled i.i.d. from \mathbf{p} , there is an online algorithm that with high probability maintains a discrepancy vector d_t such that $\|d_t\|_\infty = O(\log^4(nT))$ for all $t \in [T]$.*

In particular, for vectors in $[-1, 1]^n$ this gives an

$O(n^{1/2} \log^4(nT))$ bound, and for s -sparse vectors in $[-1, 1]^n$, this gives an $O(s^{1/2} \log^4(nT))$ bound, both of which are optimal up to poly-logarithmic factors.

The above result implies significant savings for the online Tuskány problem. Call a set $B \subseteq [0, 1]^n$ an axis-parallel box if $B = I_1 \times \dots \times I_n$ for intervals $I_i \subseteq [0, 1]$. In the online Tuskány problem, we see points $x_1, \dots, x_T \in [0, 1]^d$ and need to assign signs χ_1, \dots, χ_T in an online manner to minimize the discrepancy of every axis-parallel box at all times. More precisely, for an axis-parallel box B , define⁴

$$\text{disc}_t(B) := \left| \chi_1 \mathbf{1}_B(x_1) + \dots + \chi_t \mathbf{1}_B(x_t) \right|.$$

Our goal is to assign the signs χ_1, \dots, χ_t so as to minimize $\max_{t \leq T} \text{disc}_t(B)$ for every axis-parallel box B .

There is a standard reduction (see Section 6.2) from the online Tuskány problem to the case of s -sparse vectors in \mathbb{R}^N where $s = (\log T)^d$ but the ambient dimension N is $O_d(T^d)$. Using this reduction, along with Theorem 1.1, directly gives an $O(\log^{3d/2+4} T)$ bound for the online Tuskány's problem that works for any arbitrary distribution on points, instead of just product distributions as in [BJSS20]. In fact, we prove a more general result where we can choose arbitrary directions to test discrepancy and we use this flexibility (see Theorem 1.3 below) to improve the exponent of the bound further, and essentially match the best offline bound of $O((\log^{d-1/2} T))$ [Nik17].

THEOREM 1.2. (ONLINE TUSKÁNY'S PROBLEM) *Let \mathbf{p} be an arbitrary distribution on $[0, 1]^d$. For points x_1, \dots, x_T sampled i.i.d from \mathbf{p} , there is an algorithm which selects signs $\chi_t \in \{\pm 1\}$ such that with high probability for every axis-parallel box B , we have $\max_{t \in [T]} \text{disc}_t(B) = O_d(\log^{d+4} T)$.*

Theorem 1.1 and Theorem 1.2 follow from the more general result below about discrepancy in arbitrary test directions,

THEOREM 1.3. *Let $S \subseteq \mathbb{R}^n$ be a finite set of test vectors with Euclidean norm at most 1 and \mathbf{p} be a distribution in \mathbb{R}^n supported on vectors with Euclidean norm at most 1. Then, for vectors v_1, \dots, v_T sampled i.i.d. from \mathbf{p} , there is an online algorithm that with high probability maintains a discrepancy vector d_t satisfying*

$$\max_{z \in S} |d_t^\top z| = O((\log(|S|) + \log T) \cdot \log^3(nT))$$

for every $t \in [T]$.

⁴Here, and henceforth, for a set S , denote $\mathbf{1}_S(x)$ the indicator function that is 1 if $x \in S$ and 0 otherwise.

In fact, the proof of the above theorem also shows that given any arbitrary distribution on unit test vectors z , one can maintain a bound on the exponential moment $\mathbb{E}_z[\exp(|\langle d_t, z \rangle|)]$ at all times.

The key idea involved in proving Theorem 1.3 above, is a novel potential function approach. In addition to controlling the discrepancy d_t in the test directions, we also control how the distribution of d_t relates to the input vector distribution \mathbf{p} . This leads to better anti-concentration properties, which in turn gives better bounds on discrepancy in the test directions. We describe this idea in more detail in Section 1.2 and Section 2.

Online Banaszczyk setting. Next, we consider discrepancy with respect to general norms given by an arbitrary convex body K . To recall, in the offline setting, Banaszczyk's seminal result [Ban12] shows that if K is any convex body with Gaussian measure $1 - 1/(2T)$, then for any vectors v_1, \dots, v_T of ℓ_2 -norm at most 1, there exist signs χ_1, \dots, χ_T such that the discrepancy vectors $d_t \in O(1) \cdot K$ for all $t \in T$.

Here we study the online version when the input distribution $\mathbf{p} \in \mathbb{R}^n$ has sufficiently good tails. Specifically, we say a univariate random variable X has *sub-exponential tails* if for all $r > 0$, $\mathbb{P}[|X - \mathbb{E}[X]| > r\sigma(X)] \leq e^{-\Omega(r)}$, where $\sigma(X)$ denotes the standard-deviation of X . We say a multi-variate distribution $\mathbf{p} \in \mathbb{R}^n$ has sub-exponential tails if all its one-dimensional projections have sub-exponential tails. That is,

$$\mathbb{P}_{v \sim \mathbf{p}} \left[\left| \langle v, \theta \rangle - \mu_\theta \right| \geq \sigma_\theta \cdot r \right] \leq e^{-\Omega(r)}$$

for every $\theta \in \mathbb{S}^{n-1}$ and every $r > 0$,

where μ_θ and σ_θ are the mean and standard deviation⁵ of the scalar random variable $X_\theta = \langle v, \theta \rangle$.

Many natural distributions, such as when v is chosen uniform over the vertices of the $\{\pm 1\}^n$ hypercube (scaled to have Euclidean norm one), uniform from a convex body, Gaussian distribution (scaled to have bounded norm with high probability), or uniform on the unit sphere, have a sub-exponential tail and in these cases our bounds match the offline bounds up to poly-logarithmic factors.

THEOREM 1.4. (ONLINE BANASZCZYK SETTING) *Let $K \subseteq \mathbb{R}^n$ be a symmetric convex body with $\gamma_n(K) \geq 1/2$ and \mathbf{p} be a distribution with sub-exponential tails that is supported over vectors of Euclidean norm at*

⁵Note that when the input distribution \mathbf{p} is α -isotropic, i.e. the covariance is αI_n , then $\sigma_\theta = \alpha$ for every direction θ , but the above definition is a natural generalization to handle an arbitrary covariance structure.

most 1. Then, for vectors v_1, \dots, v_T sampled i.i.d. from \mathbf{p} , there is an online algorithm that with high probability maintains a discrepancy vector d_t satisfying $d_t \in C \log^5(nT) \cdot K$ for all $t \in [T]$ and a universal constant C .

The proof of the above theorem, while similar in spirit to [Theorem 1.3](#), is much more delicate. In particular, we cannot use that theorem directly as capturing a general convex body as a polytope may require exponential number of constraints (the set \mathcal{S} of test vectors).

Online Weighted Multi-Color Discrepancy.

Finally we consider the setting of weighted multi-color discrepancy, where we are given vectors $v_1, \dots, v_T \in \mathbb{R}^n$ sampled i.i.d. from a distribution \mathbf{p} on vectors with ℓ_2 -norm at most one, an integer R which is the number of colors available, positive weights $w_c \in [1, \eta]$ for each color $c \in [R]$, and a norm $\|\cdot\|_*$. At each time t , the algorithm has to choose a color $c \in [R]$ for the arriving vector, so that the discrepancy disc_t with respect to $\|\cdot\|_*$, defined below, is minimized for every $t \in [T]$:

$$\text{disc}_t(\|\cdot\|_*) := \max_{c \neq c'} \text{disc}_t(c, c')$$

$$\text{where } \text{disc}_t(c, c') := \left\| \frac{d_c(t)/w_c - d_{c'}(t)/w_{c'}}{1/w_c + 1/w_{c'}} \right\|_*,$$

with $d_c(t)$ being the sum of all the vectors that have been given the color c till time t . We note that (up to a factor of two) the case of unit weights and $R = 2$ is the same as assigning \pm signs to the vectors $(v_i)_{i \leq T}$, and we will also refer to this setting as *signed discrepancy*.

We show that the bounds from the previous results also extend to the setting of multi-color discrepancy.

THEOREM 1.5. (MULTI-COLOR DISCREPANCY)

For any input distribution \mathbf{p} and any set \mathcal{S} of $\text{poly}(nT)$ test vectors with Euclidean norm at most one, there is an online algorithm for the weighted multi-color discrepancy problem that maintains discrepancy $O(\log^2(R\eta) \cdot \log^4(nT))$ with the norm $\|\cdot\|_* = \max_{z \in \mathcal{S}} |\langle \cdot, z \rangle|$.

Further, if the input distribution \mathbf{p} has sub-exponential tails then one can maintain multi-color discrepancy $O(\log^2(R\eta) \cdot \log^5(nT))$ for any norm $\|\cdot\|_*$ given by a symmetric convex body K satisfying $\gamma_n(K) \geq 1/2$.

As an application, the above theorem implies upper bounds for multi-player envy minimization in the online stochastic setting, as defined in [\[BKPP18\]](#), by reductions similar to those in [\[JKS19\]](#) and [\[BJSS20\]](#).

We remark that in the offline setting, such a statement with logarithmic dependence in R and η is easy

to prove by identifying the various colors with leaves of a binary tree and recursively using the offline algorithm for signed discrepancy. It is not clear how to generalize such a strategy to the online stochastic setting, since the algorithm for signed discrepancy might use the stochasticity of the inputs quite strongly.

By exploiting the idea of working with the Haar basis, we show how to implement such a strategy in the online stochastic setting: we prove that if there is a greedy strategy for the signed discrepancy setting that uses a potential satisfying certain requirements, then it can be converted to the weighted multi-color discrepancy setting in a black-box manner.

1.2 High-Level Approach

Before describing our ideas, it is useful to discuss the bottlenecks in the previous approach. In particular, the quantitative bounds for the online Komlós problem, as well as for the case of sparse vectors obtained in [\[BJSS20\]](#) are the best possible using their approach, and improving them further required new ideas. We describe these ideas at a high-level here, and refer to [Section 2](#) for a more technical overview.

Limitations of previous approach.

For intuition, let us first consider the simpler setting, where we care about minimizing the Euclidean norm of the discrepancy vector d_t — this will already highlight the main issues. As mentioned before, if the adversary is adaptive in the online setting, then they can always choose the next input vector v_t to be orthogonal to d_{t-1} (i.e., $\langle d_{t-1}, v_t \rangle = 0$) causing $\|d_t\|_2$ to grow as $T^{1/2}$. However, if $\langle d_{t-1}, v_t \rangle$ is typically large, then one can reduce $\|d_t\|_2$ by choosing $\chi_t = -\text{sign}(\langle d_{t-1}, v_t \rangle)$, as the following shows:

$$\begin{aligned} \|d_t\|_2^2 - \|d_{t-1}\|_2^2 &= 2\chi_t \cdot \langle d_{t-1}, v_t \rangle + \|v_t\|_2^2 \\ (1.1) \qquad \qquad \qquad &\leq -2|\langle d_{t-1}, v_t \rangle| + 1. \end{aligned}$$

The key idea in [\[BJSS20\]](#) was that if the vector v_t has uncorrelated coordinates (i.e. $\mathbb{E}_{v_t \sim \mathbf{p}}[v_t(i)v_t(j)] = 0$ for $i \neq j$), then one can exploit *anti-concentration* properties to essentially argue that $|\langle d_{t-1}, v_t \rangle|$ is typically large when $\|d_{t-1}\|_2$ is somewhat big, and the greedy choice above works, as it gives a *negative drift* for the ℓ_2 -norm. However, uncorrelated vectors satisfy provably weaker anti-concentration properties, by up to a $n^{1/2}$ factor ($s^{1/2}$ for s -sparse vectors), compared to those with independent coordinates. This leads up to an extra $n^{1/2}$ loss in general.

Moreover, to ensure uncorrelation one has to work in the eigenbasis of the covariance matrix of \mathbf{p} , which could destroy sparsity in the input vectors and give bounds that scale polynomially with n . [\[BJSS20\]](#) also

show that one can combine the above high-level uncorrelation idea with a potential function that tracks a soft version of maximum discrepancy in any coordinate,

$$(1.2) \quad \Phi_{t-1} = \sum_{i=1}^n \exp(\lambda d_{t-1}(i)),$$

to even get bounds on the ℓ_∞ -norm of d_t . However, this is also problematic as it might lead to another factor n loss, due to a change of basis (twice).

To achieve sparsity based bounds in the special case of online Tusnady’s problem, previous approaches use the above ideas and exploit the special problem structure. In particular, when the input distribution \mathbf{p} is a product distribution, [BJSS20] (and [DFGGR19]) observe that one can work with the natural Haar basis which also has a product structure in $[0, 1]^d$ — this makes the input vectors uncorrelated, while simultaneously preserving the sparsity due to the recursive structure of the Haar basis. However, this severely restricts \mathbf{p} to product distributions and previously, it was unclear how to even handle a mixture of two product distributions.

New potential: anti-concentration from exponential moments. Our results are based on a new potential. Typical potential analyses for online problems show that no matter what the current state is, the potential does not rise much when the next input arrives. As discussed above, this is typically exploited in the online discrepancy setting using *anti-concentration* properties of the incoming vector $v_t \sim \mathbf{p}$ — one argues that no matter the current discrepancy vector d_{t-1} , the inner product $\langle d_{t-1}, v_t \rangle$ is typically large so that a sign can be chosen to decrease the potential (recall (1.1)).

However, as in [BJSS20], such a worst-case analysis is restrictive as it requires \mathbf{p} to have additional desirable properties such as uncorrelated coordinates. A key conceptual idea in our work is that instead of just controlling a suitable proxy for the norm of the discrepancy vectors d_t , we also seek to control structural properties of the distribution d_t . Specifically, we also seek to evolve the distribution of d_t so that it has better anti-concentration properties with respect to the input distribution. In particular, one can get much better anti-concentration for a random variable if one also has control on the higher moments. For instance, if we can bound the fourth moment of the random variable $Y_t \equiv \langle d_{t-1}, v_t \rangle$, in terms of its variance, say $\mathbb{E}[Y_t^4] \ll \mathbb{E}[Y_t^2]^2$, then the Paley-Zygmund inequality implies that Y_t is far from zero. However, working with $\mathbb{E}[Y_t^4]$ itself is too weak as an invariant and necessitates looking at even higher moments.

A key idea is that these hurdles can be handled

cleanly by looking at another potential that controls the *exponential moment* of Y_t . Specifically, all our results are based on an aggregate potential function based on combining a potential of the form (1.2), which enforces *discrepancy constraints*, together with variants of the following potential, for a suitable parameter λ , which enforces *anti-concentration constraints*:

$$\Phi_t \sim \mathbb{E}_v[\exp(\lambda |\langle d_t, v \rangle|)].$$

This clearly allows us to control higher moments of $\langle d_t, v \rangle$, in turn allowing us to show strong anti-concentration properties without any assumptions on \mathbf{p} . We believe the above idea of controlling the space of possible states where the algorithm can be present in, could potentially be useful for other applications.

To illustrate the idea in the concrete setting of ℓ_2 -discrepancy, let us consider the case when the input distribution \mathbf{p} is mean-zero and $1/n$ -isotropic, meaning the covariance $\Sigma = \mathbb{E}_{v \sim \mathbf{p}}[vv^\top] = I_n/n$. Here, if we knew that the exponential moment $\mathbb{E}_{v \sim \mathbf{p}}[\exp(|\langle d_{t-1}, v \rangle|)] \leq T$, then it implies that with high probability $|\langle d_{t-1}, v \rangle| \leq \log T$ for $v \sim \mathbf{p}$. To avoid technicalities, let us assume that $|\langle d_{t-1}, v \rangle| \leq \log T$ holds with probability one. Therefore, when v_t sampled independently from \mathbf{p} arrives, then since $\mathbb{E}[|AB|] \geq \mathbb{E}[AB]/\|B\|_\infty$ for any coupled random variables A and B , taking $A = \langle d_{t-1}, v_t \rangle$ and $B = \langle d_{t-1}, v_t \rangle / \log T$, we get that

$$\begin{aligned} \mathbb{E}[|\langle d_{t-1}, v_t \rangle|] &\geq \frac{1}{\log T} \cdot \mathbb{E}_{v_t}[d_{t-1}^\top v_t v_t^\top d_{t-1}] \\ &= \frac{1}{\log T} \cdot d_{t-1}^\top \Sigma d_{t-1} = \frac{\|d_{t-1}\|_2^2}{n \log T}. \end{aligned}$$

Therefore, whenever $\|d_{t-1}\|_2 \gg (n \log T)^{1/2}$, then the drift in ℓ_2 -norm of the discrepancy vector d_t is negative. Thus, we can obtain the optimal ℓ_2 -discrepancy bound of $O((n \log T)^{1/2})$.

Banaszczyk setting. In the Banaszczyk setting, the algorithm uses a carefully chosen set of test vectors at different scales that come from *generic chaining*. In particular, we use a potential function based on test vectors derived from the generic chaining decomposition of the polar K° of the body K .

However, as there can now be exponentially many such test vectors, more care is needed. First, we use that the Gaussian measure of K is large to control the number of test vectors at each scale in the generic chaining decomposition of K° . Second, to be able to perform a union bound over the test vectors at each scale, one needs substantially stronger tail bounds than in Theorem 1.3. To do this, we scale the test vectors

to be quite large, but this becomes problematic with standard tools for potential analysis, such as Taylor approximation, as the update to each term in the potential can be much larger than potential itself, and hard to control. Nevertheless, we show that if the distribution has sub-exponential tails, then such an approximation holds “on average” and the growth in the potential can be bounded.

Concurrent and Independent Work. In a concurrent and independent work, Alweiss, Liu, and Sawhney [ALS20] obtained online algorithms achieving poly-logarithmic discrepancy bound for the Komlós and Tusnády’s problems in the more general setting where the adversary is oblivious. Their techniques, however, are completely different from the potential function based techniques of the present paper. In fact, as noted by the authors of [ALS20], a potential function analysis encounters significant difficulties here — the algorithm is required to control the evolution of the discrepancy vectors and such an invariant is difficult to maintain with a potential function, even for stochastic inputs. With the techniques and conceptual ideas we introduce, we can overcome this barrier in the stochastic setting. We believe that our potential-based approach to control the state space of the algorithm could prove useful for other stochastic problems.

2 Proof Overview

Recall the setting: the input vectors $(v_\tau)_{\tau \leq T}$ are sampled i.i.d. from \mathbf{p} and satisfy $\|v\|_2 \leq 1$, and we need to assign signs χ_1, \dots, χ_T in an online manner so as to minimize some target norm of the discrepancy vectors $d_t = \sum_{\tau \leq t} \chi_\tau v_\tau$. Moreover, we may also assume, without loss of generality that the distribution is mean-zero as the algorithm can toss a coin and work with either v or $-v$. This means that the covariance matrix $\Sigma = \mathbb{E}_v[vv^\top]$ satisfying $0 \preceq \Sigma \preceq I_n$.

2.1 Komlós Setting Here our goal is to minimize $\|d_t\|_\infty$. First, consider the potential function $\mathbb{E}_{v \sim \mathbf{p}}[\cosh(\lambda d_t^\top v)]$ where $\cosh(a) = \frac{1}{2} \cdot (e^a + e^{-a})$. This however only puts anti-concentration constraints on the discrepancy vector and does not track the discrepancy in the coordinate directions. It is natural to add a potential term to enforce discrepancy constraints. In particular, let $\mathbf{p}_x = \frac{1}{2}\mathbf{p} + \frac{1}{2}\mathbf{p}_y$, where \mathbf{p}_y is uniform over the standard basis vectors $(e_i)_{i \leq n}$, then the potential

$$(2.3) \quad \Phi_t = \mathbb{E}_{x \sim \mathbf{p}_x}[\cosh(\lambda d_t^\top x)],$$

allows us to control the exponential moments of $\langle d_{t-1}, v_t \rangle$ as well as the discrepancy in the target test directions. In particular, if the above potential $\Phi_t \leq$

$\text{poly}(T)$, then we get a bound of $O(\lambda^{-1} \log T)$ on $\|d_t\|_\infty$. Next we sketch a proof that for the greedy strategy using the above potential, one can take $\lambda = 1/\log T$, so that the potential remains bounded by $\text{poly}(T)$ at all times.

CLAIM 3. (INFORMAL: BOUNDED DRIFT) *If $\Phi_{t-1} \leq T^2$, then $\mathbb{E}_{v_t}[\Delta\Phi_t] := \mathbb{E}_{v_t}[\Phi_t - \Phi_{t-1}] \leq 2$.*

The above implies using standard martingale arguments, that the potential remain bounded by T^2 with high probability and hence $\|d_t\|_\infty = \text{polylog}(T)$ at all times $t \in [T]$.

To continue, let us first make a simplifying assumption that $\Sigma = I_n/n$ and that at time t , the condition $\lambda|d_{t-1}^\top v_t| \leq 2 \log T$ holds with probability 1. We give an almost complete proof below under these conditions. The first condition can be dealt with by an appropriate decomposition of the covariance matrix as sketched below. The second condition only holds with high probability $(1 - 1/\text{poly}(T))$, because we have a bound on the exponential moment, but the error event can be handled straightforwardly.

By Taylor expansion, we have that for all a and for all $|\delta| \leq 1$,

$$(3.4) \quad \begin{aligned} \cosh(\lambda(a + \delta)) - \cosh(\lambda a) &\leq \lambda \sinh(\lambda a) \cdot \delta + \lambda^2 |\sinh(\lambda a)| \cdot \delta^2 + \lambda^2, \end{aligned}$$

where $\sinh(a) = \frac{1}{2} \cdot (e^a - e^{-a})$ and we used that $\cosh(a) \leq |\sinh(a)| + 1$. Therefore, since $d_t = d_{t-1} + \chi_t v_t$, by the above inequality we have

$$\begin{aligned} \Delta\Phi_t &\leq \chi_t \cdot \lambda \mathbb{E}_x [\sinh(\lambda d_{t-1}^\top x) \cdot x^\top v_t] \\ &\quad + \lambda^2 \mathbb{E}_x [|\sinh(\lambda d_{t-1}^\top x)| \cdot |x^\top v_t|^2] \\ &:= \chi_t \lambda L + \lambda^2 Q + \lambda^2. \end{aligned}$$

Since the algorithm chooses χ_t to minimize the potential, we have that $\mathbb{E}_{v_t}[\Delta\Phi_t] \leq -\lambda \mathbb{E}_{v_t}[L] + \lambda^2 \mathbb{E}_{v_t}[Q]$.

Upper bounding the quadratic term: Using that $\Sigma = \mathbb{E}_{v_t}[v_t v_t^\top] = I_n/n$, we have

$$\begin{aligned} \mathbb{E}_{v_t}[Q] &= \mathbb{E}_{v_t x} [|\sinh(\lambda d_{t-1}^\top x)| \cdot x^\top v_t v_t^\top x] \\ &= \mathbb{E}_x [|\sinh(\lambda d_{t-1}^\top x)| \cdot x^\top \Sigma x] \\ &= \frac{1}{n} \cdot \mathbb{E}_x [|\sinh(\lambda d_{t-1}^\top x)| \cdot \|x\|^2] \\ &\leq \frac{1}{n} \cdot \mathbb{E}_x [|\sinh(\lambda d_{t-1}^\top x)|], \end{aligned}$$

where the last inequality used that $\|x\|_2 \leq 1$.

Lower bounding the linear term: For this we use the aforementioned coupling trick: $\mathbb{E}_{v_t}[|L|] \geq \mathbb{E}_{v_t}[LY]/\|Y\|_\infty$ for any coupled random variable Y ⁶. Taking $Y = |d_{t-1}^\top v_t|$, we have that $\|Y\|_\infty \leq 2\lambda^{-1} \log T$. Therefore,

$$\begin{aligned} \mathbb{E}_{v_t}[|L|] &= \mathbb{E}_{v_t} \left| \mathbb{E}_x [\sinh(\lambda d_{t-1}^\top x) \cdot x^\top v_t] \right| \\ &\geq \frac{\lambda}{2 \log T} \cdot \mathbb{E}_{v_t} [\sinh(\lambda d_{t-1}^\top x) \cdot x^\top v_t v_t^\top d_{t-1}] \\ &= \frac{1}{2n \log T} \cdot \mathbb{E}_x [\sinh(\lambda d_{t-1}^\top x) \cdot \lambda d_{t-1}^\top x] \\ &\geq \frac{1}{2n \log T} \cdot \mathbb{E}_x [|\sinh(\lambda d_{t-1}^\top x)|] - 2, \end{aligned}$$

using that $\sinh(a)a \geq |\sinh(a)| - 2$ for all $a \in \mathbb{R}$.

Therefore, if $\lambda = 1/(2 \log T)$, we can bound the drift in the potential

$$\begin{aligned} \mathbb{E}_{v_t}[\Delta \Phi_t] &\leq -\frac{\lambda}{2n \log T} \cdot \mathbb{E}_x [|\sinh(\lambda d_{t-1}^\top x)|] \\ &\quad + \frac{\lambda^2}{n} \cdot \mathbb{E}_x [|\sinh(\lambda d_{t-1}^\top x)|] + 2 + \lambda^2 \leq 3. \end{aligned}$$

Non-Isotropic Covariance. To handle the general case when the covariance Σ is not isotropic, let us assume that all the non-zero eigenvalues are of the form 2^{-k} for integers $k \geq 0$. One can always rescale the input vectors and any potential set of test vectors, so that the covariance satisfies the above, while the discrepancy is affected only by a constant factor. See Section 5 for details.

With the above assumption $\Sigma = \sum_k 2^{-k} \Pi_k$ where Π_k is the orthogonal projection on to the subspace with eigenvalues 2^{-k} . Since, we only get T vectors, we can ignore the eigenvalues smaller than $(nT)^{-4}$ and only need to consider $O(\log(nT))$ different scales. Then, one can work with the following potential which imposes the alignment constraint in each such subspace:

$$\Phi_t = \sum_k \mathbb{E}_{x \sim p_x} [\cosh(\lambda d_t^\top \Pi_k x)].$$

As we have $O(\log(nT))$ pairwise orthogonal subspaces, we can still choose $\lambda = 1/\text{polylog}(nT)$ and with some care, the drift can be bounded using the aforementioned ideas. Once the potential is bounded, we can bound $\|d_t\|_\infty$ as before along with triangle inequality.

3.1 Banaszczyk Setting Recall that here we are given a convex body K with Gaussian volume at least $1/2$ and our goal is to bound K -norm of the discrepancy vector $\|d_t\|_K$. Here, $\|d\|_K$ intuitively is the minimum

scaling γ of K so that $d \in \gamma K$. To this end, we will use the dual characterization of K : Let $K^\circ = \{y : \sup_{x \in K} |\langle x, y \rangle| \leq 1\}$, then $\|d\|_K = \sup_{y \in K^\circ} |\langle d, y \rangle|$.

To approach this first note that the arguments from previous section allow us not only to bound $\|d_t\|_\infty$ but also $\max_{z \in \mathcal{S}} \langle d_t, z \rangle$ for an arbitrary set of *test directions* \mathcal{S} (of norm at most 1). As long as $|\mathcal{S}| \leq \text{poly}(nT)$, we can bound $\max_{z \in \mathcal{S}} \langle d_t, z \rangle = \text{poly}(\log(nT))$.

However, to handle a norm given by an arbitrary convex body K , one needs exponentially many test vectors, and the previous ideas are not enough. To design a suitable test distribution for an arbitrary convex body K , we use *generic chaining* to bound $\|d_t\|_K = \sup_{z \in K^\circ} \langle d_t, z \rangle$ by choosing epsilon-nets⁷ of K° at geometrically decreasing scales. Again let us assume that the $\Sigma = I_n/n$ for simplicity.

First, assuming Gaussian measure of K is at least $1/2$, it follows that $\text{diam}(K^\circ) = O(1)$ (see Section 4.3). So, one can choose the coarsest epsilon-net at $O(1)$ -scale while the finest epsilon-net can be taken at scale $\approx 1/\sqrt{n}$ since by adding the standard basis vectors to the test set, one can control $\|d_t\|_2 \leq \sqrt{n}$ (ignoring polylog factors) by using the previous ideas in the Komlós setting.

Now, one can use generic chaining as follows: define the directed layered graph \mathcal{G} (see Figure 1) where the vertices \mathcal{J}_ℓ in layer ℓ are the elements of an optimal ϵ_ℓ -net of K° with $\epsilon_\ell = 2^{-\ell}$. We add a directed edge from a vertex $u \in \mathcal{J}_\ell$ to vertex $v \in \mathcal{J}_{\ell+1}$ if $\|u - v\|_2 \leq \epsilon_\ell$ and identify the corresponding edge with the vector $v - u$. The length of any such edge $v - u$, defined as $\|v - u\|_2$, is at most ϵ_ℓ .

Let us denote the set of edges between layer ℓ and $\ell + 1$ by \mathcal{S}_ℓ . Now, one can express any $z \in K^\circ$ as $\sum_\ell w_\ell + w_{\text{err}}$ where $w_\ell \in \mathcal{S}_\ell$ and $\|w_{\text{err}}\|_2 \leq 1/\sqrt{n}$. Then, since we can control $\|d_t\|_2 \leq \sqrt{n}$, we have

$$\begin{aligned} \sup_{z \in K^\circ} \langle d_t, z \rangle &\leq \sum_\ell \max_{w \in \mathcal{S}_\ell} \langle d, w \rangle + \max_{\|w\|_2 \leq n^{-1/2}} \langle d, w_{\text{err}} \rangle \\ &= O(\log n) \cdot \max_\ell \max_{w \in \mathcal{S}_\ell} \langle d, w \rangle. \end{aligned}$$

Thus, it suffices to control $\max_{w \in \mathcal{S}_\ell} \langle d, w \rangle$ for each scale using a suitable test distribution in the potential.

For example, suppose we knew that $\mathbb{E}_{\tilde{w}}[\cosh(\lambda d^\top \tilde{w})] \leq T$ for \tilde{w} uniform in $r^2 \cdot \mathcal{S}_\ell$ for a scaling factor r^2 . Then, it would follow that $\max_{w \in \mathcal{S}_\ell} \langle d, w \rangle = O(\lambda^{-1} r^{-2} \log |\mathcal{S}_\ell| \cdot \log T)$. Standard results in convex geometry (see Section 4.3) imply that $|\mathcal{S}_\ell| \leq e^{O(1/\epsilon_\ell^2)}$, so to obtain a $\text{polylog}(nT)$ bound,

⁷We remark that one can also work with admissible nets that come from Talagrand's majorizing measures theorem and probably save a logarithmic factor, but for simplicity we work with epsilon-nets at different scales.

⁶Here $\|Y\|_\infty$ denotes the largest value of Y in its support.




Figure 1: The chaining graph \mathcal{G} showing epsilon-nets of the convex body at various scales. The edges connect near neighbors at two consecutive scales. Note that any point $z \in K^\circ$ can be expressed as the sum of the edge vectors w_ℓ where $w_\ell = v_\ell - v_{\ell-1}$, and $(v_{\ell-1}, v_\ell)$ is an edge between two points at scale $2^{-(\ell-1)}$ and $2^{-\ell}$.

one needs to scale the vectors $w \in \mathcal{S}_\ell$ by a factor of $r = 1/\epsilon_\ell$. This implies that the ℓ_2 -norm of scaled vector $r^2 \cdot w$ could be as large as \sqrt{n} .

This makes the drift analysis for the potential more challenging because now the Taylor expansion in (3.4) is not always valid as the update δ could be as large as \sqrt{n} . This is where the sub-exponential tail of the input distribution is useful for us. Since the input distribution is $1/n$ -isotropic and sub-exponential tailed, we know that if $\|w\|_2 \leq \sqrt{n}$, then for a typical choice of $v \sim \mathbf{p}$, the following holds

$$\langle v_t, w \rangle \approx \mathbb{E}_{v_t}[\langle v_t, w \rangle^2] = \mathbb{E}_{v_t}[w^\top v_t v_t^\top w] = \frac{\|w\|_2^2}{n} \leq 1.$$

Thus, with some work one can show that, the previous Taylor expansion essentially holds "on average" and the drift can be bounded. The case of general covariances can be handled by doing a decomposition as before. Although the full analysis becomes somewhat technical, all the main ideas are presented above.

3.2 Multi-color Discrepancy For the multi-color discrepancy setting, we show that if there is an online algorithm that uses a greedy strategy with respect to a certain kind of potential Φ , then one can adapt the same potential to the multi-color setting in a black-box manner.

In particular, let the number of colors $R = 2^h$ for an integer h and all weights be unit. Let us identify the leaves of a complete binary tree \mathcal{T} of height h with a color. Our goal is then to assign the incoming vector to one of the leaves. In the offline setting, this is easy to do with a logarithmic dependence of R — we start at the root and use the algorithm for the signed discrepancy setting to decide to which sub-tree the vector be assigned and then we recurse until the vector is assigned to one of the leaves. Such a strategy in the online stochastic setting is not obvious, as the

distribution of the incoming vector might change as one decides which sub-tree it belongs to.

By exploiting the idea used in [BJSS20] and [DFGGR19] of working with the Haar basis, we can implement such a strategy if the potential Φ satisfies certain requirements. Let us define $d_\ell(t)$ to be the sum of all the input vectors assigned to that leaf at time t . In the same way, for an internal node u of \mathcal{T} , we can define $d_u(t)$ to be the sum of the vectors $d_\ell(t)$ for all the leaves ℓ in the sub-tree rooted at u . The crucial insight is then, one can track the difference of the discrepancy vectors of the two children $d_u^-(t)$ for every internal node u of the tree \mathcal{T} . In particular, one can work with the potential

$$\Psi_t = \sum_{u \in \mathcal{T}} \Phi(\beta d_u^-(t)),$$

for some parameter β , and assign the incoming vector to the leaf that minimizes the increase in Ψ_t . Then, essentially we show that the analysis for the potential Φ translates to the setting of the potential Ψ_t if Φ satisfies certain requirements (see Section 12).

4 Preliminaries

4.1 Notation Throughout this paper, \log denotes the natural logarithm unless the base is explicitly mentioned. We use $[k]$ to denote the set $\{1, 2, \dots, k\}$. Sets will be denoted by script letters (e.g. \mathcal{T}).

Random variables are denoted by capital letters (e.g. A) and values they attain are denoted by lower-case letters possibly with subscripts and superscripts (e.g. a, a_1, a' , etc.). Events in a probability space will be denoted by calligraphic letters (e.g. \mathcal{E}). We also use $\mathbf{1}_\mathcal{E}$ to denote the indicator random variable for the event \mathcal{E} . We write $\lambda \mathbf{p} + (1-\lambda)\mathbf{p}'$ to denote the convex combination of the two distributions.

Given a distribution \mathbf{p} , we use the notation $x \sim \mathbf{p}$ to denote an element x sampled from the distribution \mathbf{p} .

For a real function f , we will write $\mathbb{E}_{x \sim \mathbf{p}}[f(x)]$ to denote the expected value of $f(x)$ under x sampled from \mathbf{p} . If the distribution is clear from the context, then we will abbreviate the above as $\mathbb{E}_x[f(x)]$.

For a symmetric matrix M , we use M^+ to denote the Moore-Penrose pseudo-inverse, $\|M\|_{\text{op}}$ for the operator norm of M and $\text{Tr}(M)$ for the trace of M .

4.2 Sub-exponential Tails Recall that a subexponential distribution \mathbf{p} on \mathbb{R} satisfies the following for every $r > 0$, $\mathbb{P}_{x \sim \mathbf{p}}[|x - \mu| \geq \sigma r] \leq e^{-\Omega(r)}$ where $\mu = \mathbb{E}_x[x]$ and $\sigma^2 = \mathbb{E}_x[(x - \mu)^2]$. A standard property of a distribution with a sub-exponential tail is *hypercontractivity* and a bound on the exponential moment (c.f. §2.7 in [Ver18]).

PROPOSITION 4.1. *Let \mathbf{p} be a distribution on \mathbb{R} that has a sub-exponential tail with mean zero and variance σ^2 . Then, for a constant $C > 0$, we have that $\mathbb{E}_{x \sim \mathbf{p}}[e^{|s|x|}] \leq C$ for all $|s| \leq 1/2\sigma$. Moreover, for every $k > 0$, we have $\mathbb{E}_{x \sim \mathbf{p}}[|x|^k]^{1/k} \leq C \cdot k\sigma$.*

4.3 Convex Geometry Given a convex body $K \subseteq \mathbb{R}^n$, its *polar* convex body is defined as $K^\circ = \{y \mid \sup_{x \in K} \langle x, y \rangle \leq 1\}$. If K is symmetric, then it defines a norm $\|\cdot\|_K$ which is defined as $\|\cdot\|_K = \sup_{y \in K^\circ} \langle \cdot, y \rangle$.

For a linear subspace $H \subseteq \mathbb{R}^n$, we have that $(K \cap H)^\circ = \Pi_H(K^\circ)$ where Π_H is the orthogonal projection on to the subspace H .

Gaussian Measure. We denote by γ_n the n -dimensional standard Gaussian measure on \mathbb{R}^n . More precisely, for any measurable set $\mathcal{A} \subseteq \mathbb{R}^n$, we have

$$\gamma_n(\mathcal{A}) = \frac{1}{(\sqrt{2\pi})^n} \int_{\mathcal{A}} e^{-\|x\|_2^2/2} dx.$$

For a k -dimensional linear subspace H of \mathbb{R}^n and a set $\mathcal{A} \subseteq H$, we denote by $\gamma_k(\mathcal{A})$ the Gaussian measure of the set \mathcal{A} where H is taken to be the whole space. For convenience, we will sometimes write $\gamma_H(\mathcal{A})$ to denote $\gamma_{\dim(H)}(\mathcal{A} \cap H)$.

The following is a standard inequality for the Gaussian measure of slices of a convex body. For a proof, see Lemma 14 in [DGLN16].

PROPOSITION 4.2. *Let $K \subseteq \mathbb{R}^n$ with $\gamma_n(K) \geq 1/2$ and $H \subseteq \mathbb{R}^n$ be a linear subspace of dimension k . Then, $\gamma_k(K \cap H) \geq \gamma_n(K)$.*

Gaussian Width. For a set $\mathcal{T} \subseteq \mathbb{R}^n$, let $w(\mathcal{T}) = \mathbb{E}_g[\sup_{x \in \mathcal{T}} \langle g, x \rangle]$ denote the *Gaussian width* of \mathcal{T} where $g \in \mathbb{R}^n$ is sampled from the standard normal distribution. Let $\text{diam}(\mathcal{T}) = \sup_{x, y \in \mathcal{T}} \|x - y\|_2$ denote the diameter of the set \mathcal{T} .

The following lemma is standard up to the exact constants. For a proof, see Lemmas 26 and 27 in [DGLN16].

PROPOSITION 4.3. *Let $K \subseteq \mathbb{R}^n$ be a symmetric convex body with $\gamma_n(K) \geq 1/2$. Then, $w(K^\circ) \leq \frac{3}{2}$ and $\text{diam}(K^\circ) \leq 4$.*

To prevent confusion, we remark that the Gaussian width is $\Theta(\sqrt{n})$ factor larger than the *spherical width* defined as $\mathbb{E}_\theta[\sup_{x \in \mathcal{T}} \langle \theta, x \rangle]$ for a randomly chosen θ from the unit sphere \mathbb{S}^{n-1} . So the above proposition implies that the spherical width of K° is $O(1/\sqrt{n})$.

For a linear subspace $H \subseteq \mathbb{R}^n$ and a subset $\mathcal{T} \subseteq H$, we will use the notation $w_H(\mathcal{T}) = \mathbb{E}_g[\sup_{x \in \mathcal{T}} \langle g, x \rangle]$ to denote the Gaussian width of \mathcal{T} in the subspace H , where g is sampled from the standard normal distribution on the subspace H . Proposition 4.2 and Proposition 4.3 also imply that $w_H(\mathcal{T}) \leq 3/2$.

Covering Numbers. For a set $\mathcal{T} \subseteq \mathbb{R}^n$, let $N(\mathcal{T}, \epsilon)$ denote the size of the smallest ϵ -net of \mathcal{T} in the Euclidean metric, i.e., the smallest number of closed Euclidean balls of radius ϵ whose union covers \mathcal{T} . Then, we have the following inequality (c.f. [Wai19], §5.5).

PROPOSITION 4.4. (SUDAKOV MINORATION) *For any set $\mathcal{T} \subseteq \mathbb{R}^n$ and any $\epsilon > 0$*

$$w(\mathcal{T}) \geq \frac{\epsilon}{2} \sqrt{\log N(\mathcal{T}, \epsilon)}, \text{ or equivalently,}$$

$$N(\mathcal{T}, \epsilon) \leq e^{4w(\mathcal{T})^2/\epsilon^2}.$$

Analogously, for a linear subspace $H \subseteq \mathbb{R}^n$ and a subset $\mathcal{T} \subseteq H$, we also have $w_H(\mathcal{T}) \geq \frac{\epsilon}{2} \sqrt{\log N_H(\mathcal{T}, \epsilon)}$, where $N_H(\mathcal{T}, \epsilon)$ denote the covering numbering of \mathcal{T} when H is considered the whole space.

5 Reduction to κ -Dyadic Covariance

For all our problems, we may assume without loss of generality that the distribution \mathbf{p} has zero mean, i.e. $\mathbb{E}_{v \sim \mathbf{p}}[v] = 0$, since our algorithm can toss an unbiased random coin and work with either v or $-v$. Now the covariance matrix Σ of the input distribution \mathbf{p} is given by $\Sigma = \mathbb{E}_{v \sim \mathbf{p}}[vv^\top]$. Since $\|v\|_2 \leq 1$, we have that $0 \preceq \Sigma \preceq I$ and $\text{Tr}(\Sigma) \leq 1$.

However, it will be more convenient for the proof to assume that all the non-zero eigenvalues of the covariance matrix Σ are of the form 2^{-k} for an integer k . In this section, by slightly rescaling the input distribution and the test vectors, we show that one can assume this without any loss of generality.

Consider the spectral decomposition of $\Sigma = \sum_{i=1}^n \sigma_i u_i u_i^\top$, where $0 \leq \sigma_n \leq \dots \leq \sigma_1 \leq 1$ and

u_1, \dots, u_n form an orthonormal basis of \mathbb{R}^n . Moreover, since we only get T vectors, we can essentially ignore all eigenvalues smaller than, say $(nT)^{-8}$, as this error will not affect the discrepancy too much.

For a positive integer κ denoting the number of different scales, we say that Σ is κ -dyadic if every non-zero eigenvalue σ is 2^{-k} for some $k \in [\kappa]$.

LEMMA 5.1. *Let $\mathcal{S} \subseteq \mathbb{R}^n$ be an arbitrary set of test vectors with Euclidean norm at most nT and $v \sim \mathbf{p}$ with covariance $\Sigma = \sum_i \sigma_i u_i u_i^\top$. Then, there exists a positive-semi definite matrix M with $\|M\|_{\text{op}} \leq 1$ such that the covariance of Mv is κ -dyadic for $\kappa = \lceil 8 \log(nT) \rceil$. Moreover, there exists a test set \mathcal{S}' consisting of vectors with Euclidean norm at most $\max_{y \in \mathcal{S}} \|y\|$, such that for any signs $(\chi_t)_{t \in T}$, the discrepancy vector $d_t = \sum_{\tau=1}^t \chi_\tau v_\tau$ satisfies the following with probability $1 - (nT)^{-4}$,*

$$\max_{y \in \mathcal{S}} |d_t^\top y| = 2 \cdot \max_{z \in \mathcal{S}'} |(Md_t)^\top z| + O(1).$$

Proof. For notational simplicity, we use d to denote d_t . We construct matrix M to be positive semi-definite with eigenvectors u_1, \dots, u_n . For any $i \in [n]$ such that $\sigma_i \in (2^{-k}, 2^{-k+1}]$ for some $k \in [\kappa]$, we set $Mu_i = (2^k \sigma_i)^{-1/2} \cdot u_i$, and for every $i \in [n]$ such that $\sigma_i \leq 2^{-\kappa}$, we set $Mu_i = 0$. It is easy to check that the covariance of Mv for $v \sim \mathbf{p}$ is κ -dyadic.

We define the new test set to be $\mathcal{S}' = \{\frac{1}{2}M^+y \mid y \in \mathcal{S}\}$ where M^+ is the pseudo-inverse of M . Note that $\|M^+\|_{\text{op}} \leq 2$, so every $z \in \mathcal{S}'$ satisfies $\|z\|_2 \leq \max_{y \in \mathcal{S}} \|y\| \leq nT$. To upper bound the discrepancy with respect to the test set, let Π_{err} be the projector onto the span of eigenvectors u_i with $\sigma_i \leq 2^{-\kappa}$ and let Π be the projector onto its orthogonal subspace. Then, for any $y \in \mathcal{S}$, we have

$$\begin{aligned} |d^\top y| &\leq |d^\top \Pi y| + |d^\top \Pi_{\text{err}} y| \\ &\leq |(Md)^\top (M^+y)| + nT \cdot \|\Pi_{\text{err}} d\|_2. \end{aligned}$$

Note that $\mathbb{E}\|\Pi_{\text{err}} d\|_2^2 \leq (nT)^{-8} \cdot (nT)^2$, so by Markov's inequality, with probability at least $1 - (nT)^{-4}$, we have that $\|\Pi_{\text{err}} d\|_2 \leq (nT)^{-1}$ and hence, $|d^\top \Pi_{\text{err}} y| = O(1)$ for every $y \in \mathcal{S}$. It follows that

$$\max_{y \in \mathcal{S}} |d^\top y| \leq 2 \cdot \max_{z \in \mathcal{S}'} |(Md)^\top z| + O(1).$$

This finishes the proof. \square

For all applications in this paper, the test vectors will always have Euclidean norm at most nT , so we can always assume without loss of generality that the input distribution \mathbf{p} , which is supported over vectors with Euclidean norm at most one, has mean $\mathbb{E}_{v \sim \mathbf{p}}[v] = 0$,

and its covariance $\Sigma = \mathbb{E}_v[vv^\top]$ is κ -dyadic for $\kappa = 8 \lceil \log(nT) \rceil$. We will make this assumption in the rest of this paper without stating it explicitly sometimes.

6 Discrepancy for Arbitrary Test Vectors

In this section, we consider discrepancy minimization with respect to an arbitrary set of test vectors with Euclidean length at most 1.

THEOREM 1.3. *Let $\mathcal{S} \subseteq \mathbb{R}^n$ be a finite set of test vectors with Euclidean norm at most 1 and \mathbf{p} be a distribution in \mathbb{R}^n supported on vectors with Euclidean norm at most 1. Then, for vectors v_1, \dots, v_T sampled i.i.d. from \mathbf{p} , there is an online algorithm that with high probability maintains a discrepancy vector d_t satisfying*

$$\max_{z \in \mathcal{S}} |d_t^\top z| = O((\log(|\mathcal{S}|) + \log T) \cdot \log^3(nT))$$

for every $t \in [T]$.

Before getting into the details of the proof, we first give two important applications of [Theorem 1.3](#) to the Komlós problem in [Section 6.1](#) and to the Tusnady's problem in [Section 6.2](#). The proof of [Theorem 1.3](#) will be discussed in [Section 6.3](#).

6.1 Discrepancy for Online Komlós Setting

THEOREM 1.1. (ONLINE KOMLÓS SETTING) *Let \mathbf{p} be a distribution in \mathbb{R}^n supported on vectors with Euclidean norm at most 1. Then, for vectors v_1, \dots, v_T sampled i.i.d. from \mathbf{p} , there is an online algorithm that with high probability maintains a discrepancy vector d_t such that $\|d_t\|_\infty = O(\log^4(nT))$ for all $t \in [T]$.*

Proof. [[Theorem 1.1](#)] Taking the set of test vectors $\mathcal{S} = \{e_1, \dots, e_n\}$ where e_i 's are the standard basis vectors in \mathbb{R}^n , [Theorem 1.3](#) implies an algorithm that w.h.p. maintains a discrepancy vector d_t such that $\|d_t\|_\infty = O(\log^4(nT))$ for all $t \in [T]$. \square

6.2 An Application to Online Tusnady's Problem

THEOREM 1.2. (ONLINE TUSNÁDY'S PROBLEM) *Let \mathbf{p} be an arbitrary distribution on $[0, 1]^d$. For points x_1, \dots, x_T sampled i.i.d. from \mathbf{p} , there is an algorithm which selects signs $\chi_t \in \{\pm 1\}$ such that with high probability for every axis-parallel box B , we have $\max_{t \in [T]} \text{disc}_t(B) = O_d(\log^{d+4} T)$.*

Firstly, using the probability integral transformation along each dimension, we may assume without loss of generality that the marginal of \mathbf{p} along each dimension $i \in [d]$, denoted as \mathbf{p}_i , is the uniform distribution on

$[0, 1]$. More specifically, we replace each incoming point $x \in [0, 1]^d$ by $(F_1(x_1), \dots, F_d(x_d))$, where F_i is the cumulative density function for \mathbf{p}_i . Note that $F_i(x_i)$ is uniform on $[0, 1]$ when $x_i \sim \mathbf{p}_i$. We make such an assumption throughout this subsection.

A standard approach in tackling Tuskány's problem is to decompose the unit cube $[0, 1]^d$ into a canonical set of boxes known as dyadic boxes (see [Mat09]).

Define dyadic intervals $I_{j,k} = [k2^{-j}, (k+1)2^{-j}]$ for $j \in \mathbb{Z}_{\geq 0}$ and $0 \leq k < 2^j$. A dyadic box is one of the form

$$B_{\mathbf{j}, \mathbf{k}} := I_{j(1), k(1)} \times \dots \times I_{j(d), k(d)},$$

with $\mathbf{j}, \mathbf{k} \in \mathbb{Z}^d$ such that $0 \leq j$ and $0 \leq k < 2^j$, and each side has length at least $1/T$. One can handle the error from the smaller dyadic boxes separately since few points will land in each such box. Denoting the set of dyadic boxes as $\mathcal{D} = \{B_{\mathbf{j}, \mathbf{k}} \mid 0 \leq \mathbf{j} \leq (\log T)\mathbf{1}, 0 \leq \mathbf{k} < 2^{\mathbf{j}}\}$, where $\mathbf{1} \in \mathbb{R}^d$ is the all ones vector, we note that $|\mathcal{D}| = O_d(T^d)$.

Usually, one proves a discrepancy upper bound on the set of dyadic boxes, which implies a discrepancy upper bound on all axis-parallel boxes since each axis-parallel box can be expressed roughly as the disjoint union of $O_d(\log^d T)$ dyadic boxes. This was precisely the approach used for the online Tuskány's problem in [BJSS20]. However, such an argument has a fundamental barrier. Since each arrival lands in approximately $O_d(\log^d T)$ boxes in \mathcal{D} , one can at best obtain a discrepancy upper bound of $O_d(\log^{d/2} T)$ for the set of dyadic boxes, which leads to $O_d(\log^{3d/2} T)$ discrepancy for all boxes.

Using the idea of test vectors in Theorem 1.3, we can save a factor of $O_d(\log^{d/2} T)$ over the approach above. Roughly, this saving comes from the discrepancy of dyadic boxes accumulates in an ℓ_2 manner as opposed to directly adding up. A similar idea was previously exploited by [BG17] for the offline Tuskány's problem.

Proof. [Theorem 1.2] We view Tuskány's problem as a vector balancing problem in $|\mathcal{D}|$ -dimensions with coordinates indexed by dyadic boxes, where we define $v_t(B) = \mathbf{1}_B(x_t)$ for each arrival $t \in [T]$ and every dyadic box $B \in \mathcal{D}$. Each coordinate B of the discrepancy vector $d_t = \sum_{i=1}^t \chi_i v_i$ is exactly $\text{disc}_t(B)$. Notice that $\|v_t\|_2 \leq O_d(\log^{d/2} T)$ since v_t is $O_d(\log^d T)$ -sparse. Note that v_t 's are the input vectors for the vector balancing problem.

Now we define the set of test vectors \mathcal{S} that will allow us to bound the discrepancy of any axis-parallel box. For every box B that can be exactly expressed as the disjoint union of several dyadic boxes, i.e. $B =$

$\cup_{B' \in \mathcal{D}'} B'$ for some subset $\mathcal{D}' \subseteq \mathcal{D}$ of disjoint dyadic boxes, we create a test vector $z_B \in \{0, 1\}^{|\mathcal{D}|}$ with $z_B(B') = 1$ if and only if $B' \in \mathcal{D}'$. We call such box B a *dyadic-generated* box. Since there are multiple choices of \mathcal{D}' that give the same dyadic-generated box B , we only take \mathcal{D}' to be the one that contains the smallest number of dyadic boxes. \mathcal{S} will be the set of all such dyadic-generated boxes.

Recalling that $|\mathcal{D}| = O_d(T^d)$, it follows that $|\mathcal{S}| \leq 2^{|\mathcal{D}|} = O_d(T^d)$, as each coordinate of a box in \mathcal{S} corresponds to an endpoint of one of the dyadic intervals in \mathcal{D} . Moreover, every test vector $z_B \in \mathcal{S}$ is $O_d(\log^d T)$ -sparse and thus $\|z_B\|_2 \leq O_d(\log^{d/2} T)$. Using Theorem 1.3 with both the input and test vectors scaled down by $O_d(\log^{d/2} T)$, we obtain an algorithm that w.h.p. maintains discrepancy vector d_t such that for all $t \in [T]$,

$$\max_{z_B \in \mathcal{S}} |d_t^\top z_B| \leq O_d(\log^{d+4} T).$$

Since $d_t^\top z_B = \text{disc}_t(B)$ which follows from B being a disjoint union of dyadic boxes, we have $\text{disc}_t(B) \leq O_d(\log^{d+4} T)$ for any dyadic-generated box B .

To upper bound the discrepancy of arbitrary axis-parallel boxes, we first introduce the notion of *stripes*. A stripe in $[0, 1]^d$ is an axis-parallel box that is of the form $I_1 \times \dots \times I_d$ where exactly one of the intervals I_i is allowed to be a proper sub-interval $[a, b] \subseteq [0, 1]$. The width of such a stripe is defined to be $b - a$. Stripes whose projection is $[a, b]$ in dimension i satisfying $b - a = 1/T$ correspond to the smallest dyadic interval in dimension i . We call such stripes *minimum dyadic stripes*. There are exactly T minimum dyadic stripes for each dimension $i \in [d]$. Since minimum dyadic stripes have width $1/T$ and the marginal of \mathbf{p} along any dimension is the uniform distribution over $[0, 1]$, a standard application of Chernoff bound implies that w.h.p. the total number of points in all the minimum dyadic stripes is at most $O_d(\log(T))$ points.

For a general axis-parallel box \tilde{B} , it is well-known that \tilde{B} can be expressed as the disjoint union of a dyadic-generated box B together with at most $k \leq 2d$ boxes B_1, \dots, B_k where each $B_i \subseteq S_i$ is a subset of a minimum dyadic stripe. We can thus upper bound

$$\text{disc}_t(\tilde{B}) \leq \text{disc}_t(B) + \sum_{i=1}^k \text{disc}_t(B_i) \leq \text{disc}_t(B) + \sum_{i=1}^k r_i.$$

where r_i is the total number of points in the stripe S_i . As mentioned, w.h.p. we can upper bound $\sum_{i=1}^k r_i = O_d(\log(T))$ and thus one obtains $\text{disc}_t(\tilde{B}) = O_d(\log^{d+4} T)$ for any axis-parallel box \tilde{B} . This proves the theorem. \square

6.3 Proof of Theorem 1.3

Potential Function and Algorithm. By Lemma 5.1, it is without loss of generality to assume that \mathbf{p} is κ -dyadic, where $\kappa = 8\lceil \log(nT) \rceil$. For any $k \in [\kappa]$, we use Π_k to denote the projection matrix onto the eigenspace of Σ corresponding to the eigenvalue 2^{-k} and define $\Pi = \sum_{k=1}^{\kappa} \Pi_k$ to be the sum of these projection matrices.

The algorithm for Theorem 1.3 will use a greedy strategy that chooses the next sign so that a certain potential function is minimized. To define the potential, we first define a distribution where some noise is added to the input distribution \mathbf{p} to account for the test vectors. Let \mathbf{p}_z be the uniform distribution over the set of test vectors \mathcal{S} . We define the noisy distribution \mathbf{p}_x to be $\mathbf{p}_x := \mathbf{p}/2 + \mathbf{p}_z/2$, i.e., a random sample from \mathbf{p}_x is drawn with probability 1/2 each from \mathbf{p} or \mathbf{p}_z . Note that any vector x in the support of \mathbf{p}_x satisfies $\|x\|_2 \leq 1$ since both the input distribution \mathbf{p} and the set of test vectors \mathcal{S} lie inside the unit Euclidean ball.

At any time step t , let $d_t = \chi_1 v_1 + \dots + \chi_t v_t$ denote the current discrepancy vector after the signs $\chi_1, \dots, \chi_t \in \{\pm 1\}$ have been chosen. Set $\lambda^{-1} = 100\kappa \log(nT)$ and define the potential

$$\Phi_t = \Phi(d_t) := \sum_{k=1}^{\kappa} \mathbb{E}_{x \sim \mathbf{p}_x} [\cosh(\lambda d_t^\top \Pi_k x)].$$

When the vector v_t arrives, the algorithm greedily chooses the sign χ_t that minimizes the increase $\Phi_t - \Phi_{t-1}$.

Analysis. The above potential is useful because it allows us to give tail bounds on the length of the discrepancy vectors in most directions given by the distribution \mathbf{p} while simultaneously controlling the discrepancy in the test directions. In particular, let \mathcal{G}_t denote the set of *good* vectors v in the support of \mathbf{p} that satisfy $\lambda |d_t^\top \Pi v| \leq \kappa \cdot \log(4\Phi_t/\delta)$. Then, we have the following lemma.

LEMMA 6.1. *For any $\delta > 0$ and any time t , we have*

- (a) $\mathbb{P}_{v \sim \mathbf{p}}(v \notin \mathcal{G}_t) \leq \delta$.
- (b) $|d_t^\top \Pi_k z| \leq \lambda^{-1} \log(4|\mathcal{S}|\Phi_t)$ for all $z \in \mathcal{S}$ and $k \in [\kappa]$.

Proof. (a) Recall that with probability 1/2 a sample from \mathbf{p}_x is drawn from the input distribution \mathbf{p} . Using this and the fact that $0 \leq \exp(x) \leq 2 \cosh(x)$ for any $x \in \mathbb{R}$, we have $\sum_{k \in [\kappa]} \mathbb{E}_{v \sim \mathbf{p}} [\exp(\lambda |d_t^\top \Pi_k v|)] \leq 4\Phi_t$. Note that for any $v \notin \mathcal{G}_t$, we have $\lambda |d_t^\top \Pi v| > \kappa \cdot \log(4\Phi_t/\delta)$ by

definition, so it follows that $\lambda |d_t^\top \Pi_k v| > \log(4\Phi_t/\delta)$ for at least one $k \in [\kappa]$. Thus, applying Markov's inequality we get that $\mathbb{P}_{v \sim \mathbf{p}}(v \notin \mathcal{G}_t) \leq \delta$.

- (b) Similarly, a random sample from \mathbf{p}_x is drawn from the uniform distribution over \mathcal{S} with probability 1/2, so $\exp(\lambda |d_t^\top \Pi_k z|) \leq 4|\mathcal{S}|\Phi_t$ for every $z \in \mathcal{S}$ and $k \in [\kappa]$.

This implies that $|d_t^\top \Pi_k z| \leq \lambda^{-1} \log(4|\mathcal{S}|\Phi_t)$.

□

The next lemma shows that the expected increase in the potential is small on average.

LEMMA 6.2. (BOUNDED POSITIVE DRIFT) *For any intermediate time step $t \in [T]$, if $\Phi_{t-1} \leq 3T^5$, then $\mathbb{E}_{v_t}[\Phi_t] - \Phi_{t-1} \leq 2$.*

Using Lemma 6.2, we first finish the proof of Theorem 1.3.

Proof. [Theorem 1.3] We first use Lemma 6.2 to prove that with probability at least $1 - T^{-4}$, the potential $\Phi_t \leq 3T^5$ for every $t \in [T]$. Such an argument is standard and has previously appeared in [JKS19, BJSS20]. In particular, we consider a truncated random process $\tilde{\Phi}_t$ which is the same as Φ_t until $\Phi_{t_0} > 3T^5$ for some time step t_0 ; for any t from time t_0 to T , we define $\tilde{\Phi}_t = 3T^5$. It follows that $\mathbb{P}[\tilde{\Phi}_t \geq 3T^5] = \mathbb{P}[\Phi_t \geq 3T^5]$. Lemma 6.2 implies that for any time $t \in [T]$, the expected value of the truncated process $\tilde{\Phi}_t$ over the input sequence v_1, \dots, v_T is at most $3T$. By Markov's inequality, with probability at least $1 - T^{-4}$, the potential $\Phi_t \leq 3T^5$ for every $t \in [T]$.

When the potential $\Phi_t \leq 3T^5$, part (b) of Lemma 6.1 implies that $|d_t^\top \Pi_k z| = O(\lambda^{-1} \cdot (\log(|\mathcal{S}|) + \log T))$ for any $z \in \mathcal{S}$ and $k \in [\kappa]$. Thus, it follows that for every $z \in \mathcal{S}$,

$$\begin{aligned} |d_t^\top z| &\leq \sum_{k \in [\kappa]} |d_t^\top \Pi_k z| \\ &= O(\kappa \lambda^{-1} (\log(|\mathcal{S}|) + \log T)) \\ &= O((\log(|\mathcal{S}|) + \log T) \cdot \log^3(nT)), \end{aligned}$$

which completes the proof of the theorem.

□

To finish the proof, we prove the remaining Lemma 6.2 next.

Proof. [Lemma 6.2]

Let us fix a time t . To simplify the notation, let $\Phi = \Phi_{t-1}$ and $\Delta\Phi = \Phi_t - \Phi$, and let $d = d_{t-1}$ and $v =$

v_t . To bound the change $\Delta\Phi$, we use Taylor expansion. Since $\cosh'(a) = \sinh(a)$ and $\sinh'(a) = \cosh(a)$, for any $a, b \in \mathbb{R}$ satisfying $|a - b| \leq 1$, we have

$$\begin{aligned} & \cosh(\lambda a) - \cosh(\lambda b) \\ &= \lambda \sinh(\lambda b) \cdot (a - b) + \frac{\lambda^2}{2!} \cosh(\lambda b) \cdot (a - b)^2 \\ & \quad + \frac{\lambda^3}{3!} \sinh(\lambda b) \cdot (a - b)^3 + \dots, \\ & \leq \lambda \sinh(\lambda b) \cdot (a - b) + \lambda^2 \cosh(\lambda b) \cdot (a - b)^2, \\ & \leq \lambda \sinh(\lambda b) \cdot (a - b) + \lambda^2 |\sinh(\lambda b)| \cdot (a - b)^2 \\ & \quad + \lambda^2 (a - b)^2, \end{aligned}$$

where the first inequality follows since $|\sinh(a)| \leq \cosh(a)$ for all $a \in \mathbb{R}$, and since $|a - b| \leq 1$ and $\lambda < 1$, so the higher order terms in the Taylor expansion are dominated by the first and second order terms. The second inequality uses that $\cosh(a) \leq |\sinh(a)| + 1$ for $a \in \mathbb{R}$.

After choosing the sign χ_t , the discrepancy vector $d_t = d + \chi_t v$. Defining $s_k(x) = \sinh(\lambda \cdot d^\top \Pi_k x)$ and noting that $|v^\top \Pi_k x| \leq 1$, the above upper bound on the Taylor expansion gives us that

$$\begin{aligned} \Delta\Phi &= \sum_{k \in [\kappa]} \mathbb{E}_x [\cosh(\lambda(d + \chi_t v)^\top \Pi_k x)] \\ & \quad - \sum_{k \in [\kappa]} \mathbb{E}_x [\cosh(\lambda d^\top \Pi_k x)] \\ & \leq \underbrace{\chi_t \left(\sum_{k \in [\kappa]} \lambda \mathbb{E}_x [s_k(x) v^\top \Pi_k x] \right)}_{:= \chi_t L} \\ & \quad + \underbrace{\sum_{k \in [\kappa]} \lambda^2 \mathbb{E}_x [s_k(x) \cdot x^\top \Pi_k v v^\top \Pi_k x]}_{:= Q} \\ & \quad + \underbrace{\sum_{k \in [\kappa]} \lambda^2 \mathbb{E}_x [x^\top \Pi_k v v^\top \Pi_k x]}_{:= Q_*}, \end{aligned}$$

where $\chi_t L$, Q , and Q_* denote the first, second, and third terms respectively. Recall that our algorithm uses the greedy strategy by choosing χ_t to be the sign that minimizes the potential. Taking expectation over the random incoming vector $v \sim \mathbf{p}$, we get

$$\mathbb{E}_v[\Delta\Phi] \leq -\mathbb{E}_v[|L|] + \mathbb{E}_v[Q] + \mathbb{E}_v[Q_*].$$

We will prove the following upper bounds on the quadratic (in λ) terms Q and Q_* .

CLAIM 7. $\mathbb{E}_v[Q] \leq 2\lambda^2 \sum_{k \in [\kappa]} 2^{-k} \mathbb{E}_x[|s_k(x)|]$ and $\mathbb{E}_v[Q_*] \leq 4\lambda^2$.

On the other hand, we will show that the linear (in λ) term L is also large in expectation.

CLAIM 8. $\mathbb{E}_v[|L|] \geq \lambda B^{-1} \sum_{k \in [\kappa]} 2^{-k} \mathbb{E}_x[|s_k(x)|] - 1$ for some value $B \leq 2\kappa \cdot \log(\Phi^2 \kappa n)$.

By our assumption that $\Phi \leq 3T^5$, we have that $2\lambda \leq B^{-1}$. Therefore, combining the above two claims, we get that

$$\begin{aligned} \mathbb{E}_v[\Delta\Phi] &\leq (2\lambda^2 - \lambda B^{-1}) \left(\sum_{k \in [\kappa]} 2^{-k} \mathbb{E}_x[|s_k(x)|] \right) \\ & \quad + 1 + 4\lambda^2 \leq 2. \end{aligned}$$

This finishes the proof of Lemma 6.2 assuming the claims which we prove next. \square

Proof. [Claim 7] Recall that $\mathbb{E}_v[vv^\top] = \Sigma$ and that $\Pi_k \Sigma \Pi_k = 2^{-k} \Pi_k$. Using linearity of expectation,

$$\begin{aligned} \mathbb{E}_v[Q] &= \sum_{k \in [\kappa]} \lambda^2 \mathbb{E}_x[|s_k(x)| \cdot x^\top \Pi_k \Sigma \Pi_k x] \\ &= \lambda^2 \sum_{k \in [\kappa]} 2^{-k} \mathbb{E}_x[|s_k(x)| \cdot x^\top \Pi_k x] \\ &\leq 2\lambda^2 \sum_{k \in [\kappa]} 2^{-k} \mathbb{E}_x[|s_k(x)|], \end{aligned}$$

where the last inequality uses that $\|x\|_2 \leq 1$. Similarly,

$$\begin{aligned} \mathbb{E}_v[Q_*] &= \sum_{k \in [\kappa]} \lambda^2 \mathbb{E}_x[x^\top \Pi_k \Sigma \Pi_k x] \\ &\leq 2\lambda^2 \sum_{k \in [\kappa]} 2^{-k} \leq 4\lambda^2. \end{aligned}$$

This completes the proof of the claim. \square

Proof. [Claim 8]

To lower bound the linear term, we use the fact that $|L(v)| \geq \|f\|_\infty^{-1} \cdot f(v) \cdot L(v)$ for any real-valued non-zero function f . We will choose the function $f(v) = d^\top \Pi v \cdot \mathbf{1}_{\mathcal{G}}(v)$ where \mathcal{G} will be the event that $|d^\top \Pi v|$ is small, which we know is true because of Lemma 6.1.

In particular, set $\delta^{-1} = \lambda \Phi T$ and let \mathcal{G} denote the set of vectors v in the support of \mathbf{p} such that $\lambda |d^\top \Pi v| \leq \kappa \cdot \log(4\Phi/\delta) := B$.

Then, $f(v) = d^\top \Pi v \cdot \mathbf{1}_{\mathcal{G}}(v)$ satisfies $\|f\|_\infty \leq \lambda^{-1} B$, and we can lower bound,

$$\begin{aligned} \mathbb{E}_v[|L|] &\geq \frac{\lambda}{\lambda^{-1} B} \sum_{k \in [\kappa]} \mathbb{E}_{v,x}[s_k(x) \cdot d^\top \Pi v \cdot v^\top \Pi_k x \cdot \mathbf{1}_{\mathcal{G}}(v)] \\ &= \frac{\lambda^2}{B} \sum_{k \in [\kappa]} \mathbb{E}_x[s_k(x) \cdot d^\top \Pi \Sigma \Pi_k x] \\ (8.5) \quad &\quad - \frac{\lambda^2}{B} \sum_{k \in [\kappa]} \mathbb{E}_x[s_k(x) \cdot d^\top \Pi \Sigma_{\text{err}} \Pi_k x], \end{aligned}$$

where $\Sigma_{\text{err}} = \mathbb{E}_v[vv^\top (1 - \mathbf{1}_{\mathcal{G}}(v))]$ satisfies $\|\Sigma_{\text{err}}\|_{\text{op}} \leq \mathbb{P}_{v \sim \mathbf{p}}(v \notin \mathcal{G}) \leq \delta$ using Lemma 6.1. To bound the first term in (8.5), recall that $s_k(x) = \sinh(\lambda d^\top \Pi_k x)$. Using $\Pi \Sigma \Pi_k = 2^{-k} \Pi_k$ and the fact that $\sinh(a)a \geq |\sinh(a)| - 2$ for any $a \in \mathbb{R}$, we have

$$\begin{aligned} \lambda \mathbb{E}_x[s_k(x) \cdot d^\top \Pi \Sigma \Pi_k x] &= 2^{-k} \mathbb{E}_x[s_k(x) \cdot \lambda d^\top \Pi_k x] \\ &\geq 2^{-k} (\mathbb{E}_x[|s_k(x)|] - 2). \end{aligned}$$

For the second term, we use the bound $\|\Sigma_{\text{err}}\|_{\text{op}} \leq \delta$ to obtain

$$|d^\top \Pi \Sigma_{\text{err}} \Pi_k x| \leq \|\Sigma_{\text{err}}\|_{\text{op}} \cdot \|d\|_2 \cdot \|x\|_2 \leq \delta \|d\|_2.$$

Since $\|d\|_2 \leq T$ always holds, by our choice of δ ,

$$\lambda |d^\top \Pi \Sigma_{\text{err}} \Pi_k x| \leq \Phi^{-1}.$$

Plugging the above bounds in (8.5),

$$\begin{aligned} \mathbb{E}_v[|L|] &\geq \frac{\lambda}{B} \sum_{k \in [\kappa]} 2^{-k} (\mathbb{E}_x[|s_k(x)|] - 2) \\ &\quad - \frac{\lambda}{B} \cdot \Phi^{-1} \left(\sum_{k \in [\kappa]} \mathbb{E}_x[|s_k(x)|] \right) \\ &\geq \frac{\lambda}{B} \sum_{k \in [\kappa]} 2^{-k} \mathbb{E}_x[|s_k(x)|] - \frac{\lambda}{B} \sum_{k \in [\kappa]} 2^{-k+1} - \frac{\lambda}{B} \\ &\geq \frac{\lambda}{B} \sum_{k \in [\kappa]} 2^{-k} \mathbb{E}_x[|s_k(x)|] - 1, \end{aligned}$$

where the second inequality follows since $\sum_{k \in [\kappa]} \mathbb{E}_x[|s_k(x)|] \leq \Phi$. \square

9 Discrepancy with respect to Arbitrary Convex Bodies

Our main result of this section is the following theorem.

THEOREM 1.4. (ONLINE BANASZCZYK SETTING) *Let $K \subseteq \mathbb{R}^n$ be a symmetric convex body with $\gamma_n(K) \geq 1/2$ and \mathbf{p} be a distribution with sub-exponential tails*

that is supported over vectors of Euclidean norm at most 1. Then, for vectors v_1, \dots, v_T sampled i.i.d. from \mathbf{p} , there is an online algorithm that with high probability maintains a discrepancy vector d_t satisfying $d_t \in C \log^5(nT) \cdot K$ for all $t \in [T]$ and a universal constant C .

9.1 Potential Function and Algorithm As in the previous section, it is without loss of generality to assume that \mathbf{p} is κ -dyadic, where $\kappa = 8 \lceil \log(nT) \rceil$. For any $k \in [\kappa]$, recall that Π_k denotes the projection matrix onto the eigenspace of Σ corresponding to the eigenvalue 2^{-k} and $\Pi = \sum_{k=1}^{\kappa} \Pi_k$. Further, let us also recall that Π_{err} is the projection matrix onto the subspace spanned by eigenvectors corresponding to eigenvalues of Σ that are at most $2^{-\kappa}$. We also note that $\dim(\text{im}(\Pi_k)) \leq \min\{2^k, n\}$ since $\text{Tr}(\Sigma) \leq 1$.

Our algorithm to bound the discrepancy with respect to an arbitrary symmetric convex body $K \subseteq \mathbb{R}^n$ with $\gamma_n(K) \geq 1/2$ will use a greedy strategy with a similar potential function as in §6.

Let \mathbf{p}_z be a distribution on *test vectors* in \mathbb{R}^n that will be specified later. Define the noisy distribution $\mathbf{p}_x = \mathbf{p}/2 + \mathbf{p}_z/2$, i.e., a random sample from \mathbf{p}_x is drawn from \mathbf{p} or \mathbf{p}_z with probability 1/2 each.

At any time step t , let $d_t = \chi_1 v_1 + \dots + \chi_t v_t$ denote the current discrepancy vector after the signs $\chi_1, \dots, \chi_t \in \{\pm 1\}$ have been chosen. Set $\lambda^{-1} = 100\kappa \log(nT)$, and define the potential

$$\Phi_t = \Phi(d_t) := \sum_{k \in [\kappa]} \mathbb{E}_{x \sim \mathbf{p}_x} [\exp(\lambda d_t^\top \Pi_k x)].$$

When the vector v_t arrives, the algorithm chooses the sign χ_t that minimizes the increase $\Phi_t - \Phi_{t-1}$.

Test Distribution. To complete the description of the algorithm, we need to choose a suitable distribution \mathbf{p}_z on test vectors to give us control on the norm $\|\cdot\|_K = \sup_{y \in K^\circ} \langle \cdot, y \rangle$. For this, we will use generic chaining.

First let us denote by $H_k = \text{im}(\Pi_k)$ the linear subspace that is the image of the projection matrix Π_k where the subspaces $\{H_k\}_{k \in [\kappa]}$ are orthogonal and span \mathbb{R}^n . Moreover, recall that $\dim(H_k) \leq \min\{2^k, n\}$.

Let us denote by $K_k = K \cap H_k$ the slice of the convex body K with the subspace H_k . Proposition 4.2 implies that $\gamma_{H_k}(K) \geq 1/2$ for each $k \in [\kappa]$ and combined with Proposition 4.3 this implies that $K_k^\circ := (K_k)^\circ = \Pi_k(K^\circ)$ satisfies $\text{diam}(K_k^\circ) \leq 4$ and $w_{H_k}(K_k^\circ) \leq 3/2$ for every k .

Consider ϵ -nets of the polar bodies K_k° at geomet-

rically decreasing dyadic scales. Let

$$\epsilon_{\min}(k) = 2^{-\lceil \log_2(\frac{1}{10\lambda} \sqrt{\dim(H_k)}) \rceil} \text{ and ,}$$

$$\epsilon_{\max}(k) = 2^{-\log_2 \lceil 1/\text{diam}(K_k^\circ) \rceil},$$

be the finest and the coarsest scales for a fixed k , and for integers $\ell \in [\log_2(1/\epsilon_{\max}(k)), \log_2(1/\epsilon_{\min}(k))]$, define the scale $\epsilon(\ell, k) = 2^{-\ell}$. We call these *admissible* scales for any fixed k .

Note that for a fixed $k \in [\kappa]$, the number of admissible scales is at most $2 \log_2(nT)$ since $\text{diam}(K_k^\circ) \leq 4$. The smallest scale is chosen because with high probability we can always control the Euclidean norm of the discrepancy vector in the subspace H_k to be $\lambda^{-1} \log(nT) \sqrt{\dim(H_k)}$ using a test distribution as used in Komlos's setting.

Let $\mathcal{J}(\ell, k)$ be an optimal $\epsilon(\ell, k)$ -net of K_k° . For each k , define the following directed layered graph \mathcal{G}_k (recall Figure 1) where the vertices in layer ℓ are the elements of $\mathcal{J}(\ell, k)$. Note that the first layer indexed by $\log_2(1/\epsilon_{\max}(k))$ consists of a single vertex, the origin. We add a directed edge from $u \in \mathcal{J}(\ell, k)$ to $v \in \mathcal{J}(\ell+1, k)$ if $\|v - u\|_2 \leq \epsilon(\ell, k)$. We identify an edge (u, v) with the vector $v - u$ and define its length as $\|v - u\|_2$. Let $\mathcal{S}(\ell, k)$ denote the set of edges between layer ℓ and $\ell + 1$. Note that any edge $(u, v) \in \mathcal{S}(\ell, k)$ has length at most $\epsilon(\ell, k)$ and since $w_{H_k}(K_k^\circ) \leq 3/2$, Proposition 4.4 implies that,

$$(9.6) \quad |\mathcal{S}(\ell, k)| \leq |\mathcal{J}(\ell + 1, k)|^2 \leq 2^{16/\epsilon(\ell, k)^2}.$$

Pick the final test distribution as $\mathbf{p}_z = \mathbf{p}_\Sigma/2 + \mathbf{p}_y/2$ where \mathbf{p}_Σ and \mathbf{p}_y denote the distributions given in Figure 2.

The above test distribution completes the description of the algorithm. Note that adding the eigenvectors will allow us to control the Euclidean length of the discrepancy vectors in the subspaces H_k as they form an orthonormal basis for these subspaces. Also observe that, as opposed to the previous section, the test vectors chosen above may have large Euclidean length as we scaled them. For future reference, we note that the entire probability mass assigned to length r vectors in the support of \mathbf{p}_y is at most 2^{-2r^2} where $r \geq 1/4$.

9.2 Potential Implies Low Discrepancy The test distribution \mathbf{p}_z is useful because of the following lemma. In particular, a $\text{poly}(n, T)$ upper bound on the potential function implies a polylogarithmic discrepancy upper bound on $\|d_t\|_K$.

LEMMA 9.1. *At any time t , we have that*

$$\|\Pi_k d_t\|_2 \leq \lambda^{-1} \log(4n\Phi_t) \sqrt{\dim(H_k)} \text{ and,}$$

$$\|d_t\|_K \leq O(\kappa \cdot \lambda^{-1} \cdot \log(nT) \cdot \log(\Phi_t)).$$

Proof. To derive a bound on the Euclidean length of $\Pi_k d_t$, we note that a random sample from \mathbf{p}_x is drawn from the uniform distribution over $\{u_i\}_{i \leq n}$ with probability $1/4$, so $\exp(\lambda |d_t^\top \Pi_k u_i|) \leq 4n\Phi_t$ for every $k \in [\kappa]$ and every $i \in [n]$. Since $\{u_i\}_{i \leq n}$ also form an eigenbasis for Π , we get that $|d_t^\top \Pi_k u_i| \leq \lambda^{-1} \log(4n\Phi_t)$ which implies that $\|\Pi_k d_t\|_2 \leq \lambda^{-1} \log(4n\Phi_t) \sqrt{\dim(H_k)}$.

To see the bound on $\|d_t\|_K$, we note that

$$(9.7) \quad \|d_t\|_K = \sup_{y \in K^\circ} \langle d_t, y \rangle \leq \sum_{k \in [\kappa]} \sup_{y \in K_k^\circ} \langle \Pi_k d_t, y \rangle$$

$$\leq \sum_{k \in [\kappa]} \left(\sup_{z \in \mathcal{J}(\ell, k)} |d_t^\top \Pi_k z| + \epsilon_{\min}(k) \|\Pi_k d_t\|_2 \right),$$

where the last inequality holds since $\mathcal{J}(\ell, k)$ is an $\epsilon_{\min}(k)$ -net of K_k° . By our choice of $\epsilon_{\min}(k)$ and the bound on $\|\Pi_k d_t\|_2$ from the first part of the Lemma, we have that $\epsilon_{\min}(k) \|\Pi_k d_t\|_2 \leq 10 \log(4n\Phi_t)$.

To upper bound $\sup_{z \in \mathcal{J}(\ell, k)} \langle \Pi_k d_t, z \rangle$, we pick any arbitrary $z \in \mathcal{J}(\ell, k)$ and consider any path from the origin to z in the graph \mathcal{G}_k . Let $(u_\ell, u_{\ell+1})$ be the edges of this path for $\ell \in [\log_2(1/\epsilon_{\min}), \log_2(1/\epsilon_{\max})]$ where $u_\ell = 0$ for $\ell = \log_2(1/\epsilon_{\max})$ and $u_\ell = z$ for $\ell = \log_2(1/\epsilon_{\min})$. Then $z = \sum_\ell w_\ell$ where $w_\ell = (u_{\ell+1} - u_\ell)$. By our choice of the test distribution, the bound on the potential implies the following for any edge $w \in \mathcal{S}(\ell, k)$,

$$\exp(\lambda \cdot r(\ell, k)^2 \cdot |d_t^\top \Pi_k w|)$$

$$\leq 2^{2/\epsilon(\ell, k)^2} \cdot |\mathcal{S}(\ell, k)| \cdot 4\Phi_t \leq 2^{18/\epsilon(\ell, k)^2} \cdot 4\Phi_t,$$

where the second inequality follows from $|\mathcal{S}(\ell, k)| \leq 2^{16/\epsilon(\ell, k)^2}$ in (9.6). This implies that for any edge $w \in \mathcal{S}(\ell, k)$,

$$|d_t^\top \Pi_k w| \leq \lambda^{-1} \log(4\Phi_t).$$

Since $z = \sum_\ell w_\ell$ and there are at most $\log(n)$ different scales ℓ , we get that $|d_t^\top \Pi_k z| \leq \lambda^{-1} \cdot \log(n) \cdot \log(4\Phi_t)$. Since z was arbitrary in $\mathcal{J}(\ell, k)$, plugging the above bound in (9.7) completes the proof. \square

The next lemma shows that the expected increase (or drift) in the potential is small on average.

LEMMA 9.2. (BOUNDED POSITIVE DRIFT) *Let \mathbf{p} be supported on the unit Euclidean ball in \mathbb{R}^n and has a sub-exponential tail. There exist an absolute constant $C > 0$ such that if $\Phi_{t-1} \leq T^5$ for any t , then $\mathbb{E}_{v_t \sim \mathbf{p}}[\Phi_t] - \Phi_{t-1} \leq C$.*

Analogous to the proof of Theorem 1.3, Lemma 9.2 implies that w.h.p. the potential $\Phi_t \leq T^5$ for every $t \in [T]$. Combined with Lemma 9.1, and recalling that $\kappa = O(\log nT)$ and $\lambda^{-1} = O(\kappa \log(nT))$, this proves Theorem 1.4. To finish the proof, we prove Lemma 9.2 in the next section.

- (a) \mathbf{p}_Σ is uniform over the eigenvectors u_1, \dots, u_n of the covariance matrix Σ .
- (b) \mathbf{p}_y samples a random vector as follows: pick an integer k uniformly from $[\kappa]$ and an admissible scale $\epsilon(\ell, k)$ with probability $\frac{2^{-2/\epsilon(\ell, k)^2}}{\sum_{\ell} 2^{-2/\epsilon(\ell, k)^2}}$. Choose a uniform vector from $r(\ell, k)^2 \cdot \mathcal{S}(\ell, k)$, where the scaling factor $r(\ell, k) := 1/\epsilon(\ell, k)$.

Figure 2: Test distributions \mathbf{p}_Σ and \mathbf{p}_y

9.3 Drift Analysis: Proof of Lemma 9.2 The proof is quite similar to the analysis for Komlos’s setting. In particular, we have the following tail bound analogous to Lemma 6.1. Let \mathcal{G}_t denote the set of good vectors v in the support of \mathbf{p} that satisfy $\lambda|d_t^\top \Pi v| \leq \kappa \cdot \log(4\Phi_t/\delta)$.

LEMMA 9.3. For any $\delta > 0$ and any time t , we have $\mathbb{P}_{v \sim \mathbf{p}}(v \notin \mathcal{G}_t) \leq \delta$.

We omit the proof of the above lemma as it is the same as that of Lemma 6.1.

Proof. [Proof of Lemma 9.2] Recall that our potential function is defined to be

$$\Phi_t := \sum_{k \in [\kappa]} \mathbb{E}_{x \sim \mathbf{p}_x} [\exp(\lambda d_t^\top \Pi_k x)],$$

where $\mathbf{p}_x = \mathbf{p}/2 + \mathbf{p}_\Sigma/4 + \mathbf{p}_y/4$ is a combination of the input distribution \mathbf{p} and test distributions \mathbf{p}_Σ and \mathbf{p}_y , each constituting a constant mass.

Let us fix a time t . To simplify the notation, we denote $\Phi = \Phi_{t-1}$ and $\Delta\Phi = \Phi_t - \Phi$, and denote $d = d_{t-1}$ and $v = v_t$. To bound the potential change $\Delta\Phi$, we use the following inequality, which follows from a modification of the Taylor series expansion of $\cosh(r)$ and holds for any $a, b \in \mathbb{R}$,

$$(9.8) \quad \begin{aligned} & \cosh(\lambda a) - \cosh(\lambda b) \\ & \leq \lambda \sinh(\lambda b) \cdot (a - b) + \frac{\lambda^2}{2} \cosh(\lambda b) \cdot e^{|a-b|} (a - b)^2. \end{aligned}$$

Note that when $|a - b| \ll 1$, then $e^{|a-b|} \leq 2$, so one gets the first two terms of the Taylor expansion as an upper bound, but here we will also need it when $|a - b| \gg 1$.

Note that every vector in the support of \mathbf{p} and \mathbf{p}_Σ has Euclidean length at most 1, while $y \sim \mathbf{p}_y$ may have large Euclidean length due to the scaling factor of $r(\ell, k)^2$. Therefore, we decompose the distribution \mathbf{p}_x appearing in the potential as $\mathbf{p}_x = \frac{3}{4}\mathbf{p}_w + \frac{1}{4}\mathbf{p}_y$, where the distribution $\mathbf{p}_w = \frac{2}{3}\mathbf{p} + \frac{1}{3}\mathbf{p}_\Sigma$ is supported on vectors with Euclidean length at most 1.

After choosing the sign χ_t for v , the discrepancy vector d_t becomes $d + \chi_t v$. For ease of notation, define $s_k(x) = \sinh(\lambda \cdot d^\top \Pi_k x)$ and $c_k(x) = \cosh(\lambda \cdot d^\top \Pi_k x)$ for any $x \in \mathbb{R}^n$. Now (9.8) implies that $\Delta\Phi := \Delta\Phi_1 + \Delta\Phi_2$ where

$$\begin{aligned} \Delta\Phi_1 & \leq \chi_t \cdot \frac{3}{4} \left(\sum_{k \in [\kappa]} \lambda \mathbb{E}_w [s_k(w) v^\top \Pi_k w] \right) \\ & \quad + \frac{3}{4} \sum_{k \in [\kappa]} \lambda^2 \mathbb{E}_w [c_k(w) \cdot w^\top \Pi_k v v^\top \Pi_k w] \\ & := \chi_t L_1 + Q_1, \text{ and,} \\ \Delta\Phi_2 & \leq \chi_t \cdot \frac{1}{4} \left(\sum_{k \in [\kappa]} \lambda \mathbb{E}_y [s_k(y) v^\top \Pi_k y] \right) \\ & \quad + \frac{1}{4} \sum_{k \in [\kappa]} \lambda^2 \mathbb{E}_y [c_k(y) \cdot e^{\lambda |v^\top \Pi_k y|} y^\top \Pi_k v v^\top \Pi_k y] \\ & := \chi_t L_2 + Q_2. \end{aligned}$$

Since our algorithm chooses sign χ_t to minimize the potential increase, taking expectation over the incoming vector v , we get

$$\mathbb{E}_v[\Delta\Phi] \leq -\mathbb{E}_v[|L_1 + L_2|] + \mathbb{E}_v[Q_1 + Q_2].$$

We will prove the following upper bounds on the quadratic terms (in λ) Q_1 and Q_2 .

CLAIM 10. For an absolute constant $C > 0$,

$$\mathbb{E}_v[Q_1 + Q_2] \leq C \cdot \lambda^2 \sum_{k \in [\kappa]} 2^{-k} \mathbb{E}_x [c_k(x) \|x\|_2^2].$$

On the other hand, we will show that the linear (in λ) terms $L_1 + L_2$ is also large in expectation.

CLAIM 11. For some $B \leq 4\kappa \log(\Phi^2 n \kappa)$,

$$\mathbb{E}_v[|L_1 + L_2|] \geq \lambda B^{-1} \sum_{k \in [\kappa]} 2^{-k} \mathbb{E}_x [c_k(x) \|x\|_2^2] - O(1).$$

By our assumption of $\Phi \leq T^5$, so it follows that

$2\lambda \leq B^{-1}$. Therefore, combining the above two claims,

$$\mathbb{E}_v[\Delta\Phi] \leq (2\lambda^2 - \lambda B^{-1}) \left(\sum_{k \in [\kappa]} 2^{-k} \mathbb{E}_x [c_k(x) \|x\|_2^2] \right) + C \leq C,$$

which finishes the proof of [Lemma 9.2](#) assuming the claims. \square

To prove the missing claims, we need the following property that follows from the sub-exponential tail of the input distribution \mathfrak{p} .

LEMMA 11.1. *There exists a constant $C > 0$, such that for every integer $k \in [\kappa]$, and any $y \in \text{im}(\Pi_k)$ satisfying $\|y\|_2 \leq \frac{1}{4} \sqrt{\min\{2^k, n\}}$, the following holds*

$$\mathbb{E}_{v \sim \mathfrak{p}} [e^{\lambda |v^\top y|} \cdot |v^\top y|^2] \leq C \cdot 2^{-k} \cdot \|y\|_2^2 \text{ for all } \lambda \leq 1.$$

We remark that this is the only step in the proof which requires the sub-exponential tail, as otherwise the exponential term above may be quite large. It may however be possible to exploit some more structure from the test vectors y and the discrepancy vector to prove the above lemma without any sub-exponential tail requirements from the input distribution.

Proof. As $y \in \text{im}(\Pi_k)$, we have that $v^\top y = v^\top \Pi_k y$ which is a scalar sub-exponential random variable with zero mean and variance at most

$$\begin{aligned} \sigma_y^2 &:= \mathbb{E}_v [v^\top \Pi_k y]^2 \leq \|\Pi_k \Sigma \Pi_k\|_{\text{op}} \|y\|_2^2 \\ &\leq 2^{-k} \|y\|_2^2 \leq 1/16. \end{aligned}$$

Using Cauchy-Schwarz and [Proposition 4.1](#), we get that

$$\begin{aligned} \mathbb{E}_v [e^{\lambda |v^\top y|} \cdot |v^\top y|^2] &\leq \sqrt{\mathbb{E}_v [e^{2\lambda |v^\top y|}]} \cdot \sqrt{\mathbb{E}_v [|v^\top y|^4]} \\ &\leq C \cdot \mathbb{E}_v [|v^\top \Pi_k y|^2] \\ &\leq C \cdot 2^{-k} \|y\|_2^2, \end{aligned}$$

where the exponential term is bounded as $\sigma_y \leq \frac{1}{4}$. \square

Proof. [[Claim 10](#)] Recall that $\mathbb{E}_v [vv^\top] = \Sigma$ which satisfies $\Pi_k \Sigma \Pi_k = 2^{-k} \Pi_k$. Therefore, using linearity of expectation,

$$\begin{aligned} \mathbb{E}_v [Q_1] &= \frac{3}{4} \sum_{k \in [\kappa]} \lambda^2 \mathbb{E}_w [c_k(w) \cdot w^\top \Pi_k \Sigma \Pi_k w] \\ &= \lambda^2 \cdot \frac{3}{4} \sum_{k \in [\kappa]} 2^{-k} \mathbb{E}_w [c_k(w) \cdot w^\top \Pi_k w] \\ (11.9) \quad &\leq 2\lambda^2 \cdot \frac{3}{4} \sum_{k \in [\kappa]} 2^{-k} \mathbb{E}_w [c_k(w) \|w\|_2^2]. \end{aligned}$$

We next use [Lemma 11.1](#) to bound the second quadratic term

$$\mathbb{E}_v [Q_2] = \frac{1}{4} \sum_{k \in [\kappa]} \lambda^2 \mathbb{E}_y [c_k(y) \cdot e^{\lambda |v^\top \Pi_k y|} y^\top \Pi_k v v^\top \Pi_k y].$$

For any $k \in [\kappa]$ and any $y \in \text{im}(\Pi_k)$ that is in the support of \mathfrak{p}_y , we have that

$$\begin{aligned} \lambda \|\Pi_k y\|_2 &\leq \lambda \cdot \|y\|_2 \leq \lambda / \epsilon_{\min}(k) \\ &\leq \lambda \cdot \frac{1}{10\lambda} \cdot \sqrt{\dim(H_k)} \leq \frac{1}{4} \sqrt{\min\{n, 2^k\}}. \end{aligned}$$

On the other hand, if $y \in \text{im}(\Pi_{k'})$ for $k' \neq k$, then the above quantity is zero. [Lemma 11.1](#) then implies that for any y in the support of \mathfrak{p}_y ,

$$\begin{aligned} \mathbb{E}_v [e^{\lambda v^\top \Pi_k y} \cdot |\lambda v^\top \Pi_k y|^2] &\leq C_1 \cdot 2^{-k} \|\lambda \Pi_k y\|_2^2 \\ &\leq C_1 \lambda^2 \cdot 2^{-k} \|y\|_2^2, \end{aligned}$$

where C_1 is some absolute constant. Therefore, we obtain the following bound

$$(11.10) \quad \mathbb{E}_v [Q_2] \leq C_1 \cdot \lambda^2 \cdot \sum_{k \in [\kappa]} 2^{-k} \mathbb{E}_y [c_k(y) \|y\|_2^2].$$

Summing up [\(11.9\)](#) and [\(11.10\)](#) finishes the proof of the claim. \square

Proof. [[Proof of Claim 11](#)] Let $L = L_1 + L_2$. To lower bound the linear term, we proceed similarly as in the proof of [Claim 8](#) and use the fact that $|L(v)| \geq \|f\|_\infty^{-1} \cdot f(v) \cdot L(v)$ for any real-valued non-zero function f . We will choose the function $f(v) = d^\top \Pi v \cdot \mathbf{1}_{\mathcal{G}}(v)$ where \mathcal{G} will be the event that $|d^\top \Pi v|$ is small which we know is true because of [Lemma 9.3](#).

In particular, set $\delta^{-1} = \lambda^{-2} n \cdot \Phi \cdot \log(4n\Phi)$ and let \mathcal{G} denote the set of vectors v in the support of \mathfrak{p} such that $\lambda |d^\top \Pi v| \leq \kappa \cdot \log(4\Phi/\delta) := B$.

Then, $f(v) = d^\top \Pi v \cdot \mathbf{1}_{\mathcal{G}}(v)$ satisfies $\|f\|_\infty \leq \lambda^{-1} B$,

and we can lower bound,

$$\begin{aligned}
 (11.11) \quad & \mathbb{E}_v[|L|] \\
 & \geq \frac{\lambda}{\lambda^{-1}B} \cdot \frac{3}{4} \sum_{k \in [\kappa]} \mathbb{E}_{vw}[s_k(w) \cdot d^\top \Pi v \cdot v^\top \Pi_k w \cdot \mathbf{1}_{\mathcal{G}}(v)] \\
 & \quad + \frac{\lambda}{\lambda^{-1}B} \cdot \frac{1}{4} \sum_{k \in [\kappa]} \mathbb{E}_{vy}[s_k(y) \cdot d^\top \Pi v \cdot v^\top \Pi_k y \cdot \mathbf{1}_{\mathcal{G}}(v)] \\
 & = \frac{\lambda^2}{B} \cdot \frac{3}{4} \sum_{k \in [\kappa]} \mathbb{E}_w[s_k(w) \cdot d^\top \Pi \Sigma \Pi_k w] \\
 & \quad - \frac{\lambda^2}{B} \cdot \frac{3}{4} \sum_{k \in [\kappa]} \mathbb{E}_w[s_k(w) \cdot d^\top \Pi \Sigma_{\text{err}} \Pi_k w] \\
 & \quad + \frac{\lambda^2}{B} \cdot \frac{1}{4} \sum_{k \in [\kappa]} \mathbb{E}_y[s_k(y) \cdot d^\top \Pi \Sigma \Pi_k y] \\
 & \quad - \frac{\lambda^2}{B} \cdot \frac{1}{4} \sum_{k \in [\kappa]} \mathbb{E}_y[s_k(y) \cdot d^\top \Pi \Sigma_{\text{err}} \Pi_k y],
 \end{aligned}$$

where $\Sigma_{\text{err}} = \mathbb{E}_v[vv^\top(1 - \mathbf{1}_{\mathcal{G}}(v))]$ satisfies $\|\Sigma_{\text{err}}\|_{\text{op}} \leq \mathbb{P}_{v \sim \mathfrak{p}}(v \notin \mathcal{G}) \leq \delta$ using [Lemma 6.1](#).

To bound the terms involving Σ in (11.11), we recall that $s_k(x) = \sinh(\lambda d^\top \Pi_k x)$ and $c_k(x) = \cosh(\lambda d^\top \Pi_k x)$. Using $\Pi \Sigma \Pi_k = 2^{-k} \Pi_k$ and the fact that $\sinh(a)a \geq \cosh(a)|a| - 2$ for any $a \in \mathbb{R}$, we have

$$\begin{aligned}
 \lambda \mathbb{E}_w[s_k(w) \cdot d^\top \Pi \Sigma \Pi_k w] &= 2^{-k} \mathbb{E}_w[s_k(w) \cdot \lambda d^\top \Pi_k w] \\
 &\geq 2^{-k} (\mathbb{E}_w[c_k(w)|\lambda d^\top \Pi_k w|] - 2),
 \end{aligned}$$

and similarly for y .

The terms with Σ_{err} can be upper bounded using $\|\Sigma_{\text{err}}\|_{\text{op}} \leq \delta$. In particular, we have

$$|d^\top \Pi \Sigma_{\text{err}} \Pi_k x| \leq \|\Pi d\|_2 \|\Sigma_{\text{err}}\|_{\text{op}} \|x\|_2 \leq \delta \|\Pi d\|_2 \|x\|_2.$$

Since $\Pi = \sum_{k \in [\kappa]} \Pi_k$ and $(\Pi_k)_{k \in [\kappa]}$ are orthogonal projectors, [Lemma 9.1](#) implies that $\|\Pi d\|_2 \leq \lambda^{-1} \log(4n\Phi) \sqrt{n}$. Moreover, we have $\|w\|_2 \leq 1$ and $\|y\|_2 \leq \min_k \{1/\epsilon_{\min}(k)\} \leq \frac{1}{10\lambda} \cdot \sqrt{n}$. Then, by our choice of $\delta^{-1} = \lambda^{-2} n \Phi \cdot \log(4n\Phi)$, we have

$$\lambda |d^\top \Pi \Sigma_{\text{err}} \Pi_k x| \leq \delta \lambda^{-1} n \log(4n\Phi) = \Phi^{-1}.$$

Plugging the above bounds in (11.11), we obtain

$$\begin{aligned}
 (11.12) \quad \mathbb{E}_v[|L|] &\geq \frac{\lambda}{B} \cdot \frac{3}{4} \sum_{k \in [\kappa]} 2^{-k} \mathbb{E}_w[c_k(w)|\lambda d^\top \Pi_k w|] \\
 &\quad + \frac{\lambda}{B} \cdot \frac{1}{4} \sum_{k \in [\kappa]} 2^{-k} \mathbb{E}_y[c_k(y)|\lambda d^\top \Pi_k y|] - 4,
 \end{aligned}$$

where we used the upper bound $\sum_{k \in [\kappa]} \mathbb{E}_x[|s_k(x)|] \leq \Phi$ to control the error term involving Σ_{err} .

To finish the proof, we bound the two terms in (11.12) separately. We first use the inequality that $\cosh(a)a \geq \cosh(a) - 2$ for all $a \in \mathbb{R}$ and the fact that $\|w\|_2 \leq 1$ for every w in the support of \mathfrak{p}_w to get that

$$\begin{aligned}
 (11.13) \quad \mathbb{E}_w[c_k(w)|\lambda d^\top \Pi_k w|] &\geq \mathbb{E}_w[c_k(w)] - 2 \\
 &\geq \mathbb{E}_w[c_k(w)\|w\|_2^2] - 2.
 \end{aligned}$$

To bound the second term in (11.12), we recall that the entire probability mass assigned to length r vectors (i.e. $\epsilon(\ell, k) = 1/r$) in the support of \mathfrak{p}_y is at most 2^{-2r^2} , where $r \geq 1/4$. Let \mathcal{E} be the event that $|\lambda d^\top \Pi_k y| \leq \|y\|_2^2$.

Note that $c_k(y)\|y\|_2^2 \leq 2^{r^2} r^2$ if $\|y\|_2 = r$. This implies that

$$\begin{aligned}
 (11.14) \quad \mathbb{E}_y[c_k(y)|\lambda d^\top \Pi_k y|] &\geq \mathbb{E}_y[c_k(y)\|y\|_2^2] - \mathbb{E}_y[c_k(y)\|y\|_2^2 \cdot \mathbf{1}_{\mathcal{E}}(y)] \\
 &\geq \mathbb{E}_y[c_k(y)\|y\|_2^2] - \int_{1/4}^{\infty} 2^{-2r^2} 2^{r^2} r^2 \\
 &\geq \mathbb{E}_y[c_k(y)\|y\|_2^2] - 1.
 \end{aligned}$$

Since $\mathfrak{p}_x = \frac{3}{4}\mathfrak{p}_w + \frac{1}{4}\mathfrak{p}_y$, plugging (11.13) and (11.14) into (11.12) give that $\mathbb{E}_v[|L|] \geq \lambda B^{-1} \sum_{k \in [\kappa]} 2^{-k} \mathbb{E}_x[c_k(x)\|x\|_2^2] - C$, for some constant $C > 0$, which completes the proof of the claim. \square

12 Generalization to Weighted Multi-Color Discrepancy

In this section, we prove [Theorem 1.5](#) which follows from a black-box way of converting an algorithm for the signed discrepancy setting to the multi-color setting.

In particular, for a parameter $0 \leq \lambda \leq 1$, let $\Phi : \mathbb{R}^n \rightarrow \mathbb{R}_+$ be a potential function satisfying

$$\begin{aligned}
 (12.15) \quad & \Phi(d + \alpha v) \leq \Phi(d) + \lambda \alpha L_d(v) + \lambda^2 \alpha^2 Q_d(v) \\
 & \quad \text{for every } d, v \in \mathbb{R}^n \text{ and } |\alpha| \leq 1, \text{ and} \\
 & -\lambda \cdot \mathbb{E}_{v \sim \mathfrak{p}}[|L_d(v)|] + \lambda^2 \cdot \mathbb{E}_{v \sim \mathfrak{p}}[Q_d(v)] = O(1) \\
 & \quad \text{for any } d \text{ such that } \Phi(d) \leq 3T^5,
 \end{aligned}$$

where $L_d : \mathbb{R}^n \rightarrow \mathbb{R}$ and $Q_d : \mathbb{R}^n \rightarrow \mathbb{R}_+$ are arbitrary functions of v that depend on d .

One can verify that the first condition is always satisfied for the potential functions used for proving [Theorem 1.3](#) and [Theorem 1.4](#), while the second condition holds for $\lambda = O(1/\log^2(nT))$ because of [Lemma 6.2](#) and [Lemma 9.2](#).

Moreover, for parameters n and T , let $B_{\|\cdot\|_*}$ be such that if the potential $\Phi(d) = \Phi$, then the corresponding norm $\|d\|_* \leq B_{\|\cdot\|_*} \log(nT\Phi)$. Part (b) of [Lemma 6.1](#) implies that for any test set \mathcal{S} of $\text{poly}(nT)$ vectors

contained in the unit Euclidean ball, if the norm $\|\cdot\|_* = \max_{z \in S} |\langle \cdot, z \rangle|$, then $B_{\|\cdot\|_*} = O(\log^3(nT))$. Similarly, if $\|\cdot\|_*$ is given by a symmetric convex body with Gaussian measure at least $1/2$, then Lemma 9.1 implies that $B_{\|\cdot\|_*} = O(\log^4(nT))$.

We use the above properties of the potential Φ to give a greedy algorithm for the multi-color discrepancy setting.

12.1 Weighted Binary Tree Embedding We first show how to embed the weighted multi-color discrepancy problem into a binary tree \mathcal{T} of height $O(\log(R\eta))$. For each color c , we create $\lfloor w_c \rfloor$ nodes with weight $w_c/\lfloor w_c \rfloor \in [1, 2]$ each. The total number of nodes is thus $M_\ell = \sum_{c \in [R]} \lfloor w_c \rfloor = O(R\eta)$. In the following, we place these nodes as the leaves of an (incomplete) binary tree.

Take the height $h = O(\log(R\eta))$ to be the smallest exponent of 2 such that $2^h \geq M_\ell$. We first remove $2^h - M_\ell < 2^{h-1}$ leaves from the complete binary tree of height h such that none of the removed leaves are siblings. Denote the set of remaining leaves as $\mathcal{L}(\mathcal{T})$. Then from left to right, assign the leaves in $\mathcal{L}(\mathcal{T})$ to the R colors so that leaves corresponding to the same color are consecutive. For each leaf node $\ell \in \mathcal{L}(\mathcal{T})$ that is assigned the color $c \in [R]$, we assign it the weight $w_\ell = w_c/\lfloor w_c \rfloor$.

We index the internal nodes of the tree as follows: for integers $0 \leq j \leq h-1$ and $0 \leq k \leq 2^j$, we use (j, k) to denote the 2^k -th node at depth j . Note that the left and right children of a node (j, k) are the nodes $(j+1, 2k)$ and $(j+1, 2k+1)$. The weight $w_{j,k}$ of an internal node (j, k) is defined to be sum of weights of all the leaves in the sub-tree rooted at (j, k) . This way of embedding satisfies certain desirable properties which we give in the following lemma.

LEMMA 12.1. (BALANCED TREE EMBEDDING) *For the weighted (incomplete) binary tree \mathcal{T} defined above, for any two nodes (j, k) and (j, k') in the same level,*

$$1/4 \leq w_{j,k}/w_{j,k'} \leq 4.$$

Proof. Observe that each leaf node $\ell \in \mathcal{L}(\mathcal{T})$ has weight $w_\ell \in [1, 2]$. Moreover, for each internal node $(h-1, k)$ in the level just above the leaves, at least one of its children is not removed in the construction of \mathcal{T} . Therefore, it follows that $w_{j,k} = a_{j,k}2^{h-j}$ for some $a_{j,k} \in [1/2, 2]$ and similarly for (j, k') . The lemma now immediately follows from these observations. \square

Induced random walk on the weighted tree. Randomly choosing a leaf with probability proportional to its weight induces a natural random walk on the

tree \mathcal{T} : the walk starts from the root and moves down the tree until it reaches one of the leaves. Conditioned on the event that the walk is at some node (j, k) in the j -th level, it goes to left child $(j+1, 2k)$ with probability $q_{j,k}^l = w_{j+1,2k}/w_{j,k}$ and to the right child $(j+1, 2k+1)$ with probability $q_{j,k}^r = w_{j+1,2k+1}/w_{j,k}$. Note that by Lemma 12.1 above, we have that both $q_{j,k}^l, q_{j,k}^r \in [1/5, 4/5]$ for each internal node (j, k) in the tree. Note that $w_{j,k}/w_{0,0}$ denotes the probability that the random walk passes through the vertex j, k .

12.2 Algorithm and Analysis Recall that each leaf $\ell \in \mathcal{L}(\mathcal{T})$ of the tree \mathcal{T} is associated with a color. Our online algorithm will assign each arriving vector v_t to one of the leaves $\ell \in \mathcal{L}(\mathcal{T})$ and its color will then be the color of the corresponding leaf.

For a leaf $\ell \in \mathcal{L}(\mathcal{T})$, let $d_\ell(t)$ denote the sum of all the input vectors that are associated with the leaf ℓ at time t . For an internal node (j, k) , we define $d_{j,k}(t)$ to be the sum $\sum_{\ell \in \mathcal{L}(\mathcal{T}_{j,k})} d_\ell(t)$ where $\mathcal{L}(\mathcal{T}_{j,k})$ is the set of all the leaves in the sub-tree rooted at (j, k) . Also, let $d_{j,k}^l(t) = d_{j+1,2k}(t)$ and $d_{j,k}^r(t) = d_{j+1,2k+1}(t)$ be the vectors associated with the left and right child of the node (j, k) .

Finally let,

$$d_{j,k}^-(t) = \frac{d_{j,k}^l(t)/q_{j,k}^l - d_{j,k}^r(t)/q_{j,k}^r}{1/q_{j,k}^l + 1/q_{j,k}^r} = q_{j,k}^r d_{j,k}^l(t) - q_{j,k}^l d_{j,k}^r(t),$$

denote the weighted difference between the two children vectors for the (j, k) -th node of the tree.

Algorithm. For $\beta = 1/(400h)$, consider the following potential function

$$\Psi_t = \sum_{j,k \in \mathcal{T}} \Phi(\beta d_{j,k}^-(t)),$$

where the sum is over all the internal nodes (j, k) of \mathcal{T} .

The algorithm assigns the incoming vector v_t to the leaf $\ell \in \mathcal{L}(\mathcal{T})$, so that the increase in the potential $\Psi_t - \Psi_{t-1}$ is minimized. The color assigned to the vector v_t is then the color of the corresponding leaf ℓ .

We show that if the potential Φ satisfies (12.15), then the drift for the potential Ψ can be bounded.

LEMMA 12.2. *If at any time t , if $\Psi_{t-1} \leq T^5$, then the following holds*

$$\mathbb{E}_{v_t \sim p}[\Delta \Psi_t] := \mathbb{E}_{v_t \sim p}[\Psi_t - \Psi_{t-1}] = O(1).$$

Using standard arguments as used in the proof of Theorem 1.3, this implies that with high probability $\Psi_t \leq T^5$ at all times t .

Moreover, the above potential also gives a bound on the discrepancy because of the following lemma.

LEMMA 12.3. If $\Psi_t \leq T^5$, then

$$\begin{aligned} \text{disc}_t &= O(\beta^{-1}h \cdot B_{\|\cdot\|_*} \cdot \log(nT\Psi_t)) \\ &= O(h^2 \cdot B_{\|\cdot\|_*} \cdot \log(nT)). \end{aligned}$$

Combined with part (b) of Lemma 6.1 and Lemma 9.1, the above implies Theorem 1.5. For proofs of Lemma 12.3 and Lemma 12.2, we refer to the full version of this paper [BJM⁺20].

References

- [ALS20] Ryan Alweiss, Yang P. Liu, and Mehtaab Sawhney. Discrepancy minimization via a self-balancing walk. *CoRR*, abs/2006.14009, 2020.
- [Ban98] Wojciech Banaszczyk. Balancing vectors and Gaussian measures of n -dimensional convex bodies. *Random Struct. Algorithms*, 12(4):351–360, 1998.
- [Ban10] Nikhil Bansal. Constructive Algorithms for Discrepancy Minimization. In *Proceedings of FOCS 2010*, pages 3–10, 2010.
- [Ban12] Wojciech Banaszczyk. On series of signed vectors and their rearrangements. *Random Struct. Algorithms*, 40(3):301–316, 2012.
- [Bár79] Imre Bárány. On a Class of Balancing Games. *J. Comb. Theory, Ser. A*, 26(2):115–126, 1979.
- [BDG16] Nikhil Bansal, Daniel Dadush, and Shashwat Garg. An Algorithm for Komlós Conjecture Matching Banaszczyk’s Bound. In *Proceedings of FOCS 2016*, pages 788–799, 2016.
- [BDGL18] Nikhil Bansal, Daniel Dadush, Shashwat Garg, and Shachar Lovett. The Gram-Schmidt Walk: A Cure for the Banaszczyk Blues. In *Proceedings of STOC 2018*, pages 587–597, 2018.
- [Bec81] József Beck. Balanced two-colorings of finite sets in the square I. *Combinatorica*, 1(4):327–335, 1981.
- [BG17] Nikhil Bansal and Shashwat Garg. Algorithmic discrepancy beyond partial coloring. In *Proceedings of STOC 2017*, pages 914–926, 2017.
- [BJM⁺20] Nikhil Bansal, Haotian Jiang, Raghu Meka, Sahil Singla, and Makrand Sinha. Online Discrepancy Minimization for Stochastic Arrivals. *CoRR*, abs/2007.10622, 2020.
- [BJSS20] Nikhil Bansal, Haotian Jiang, Sahil Singla, and Makrand Sinha. Online vector balancing and geometric discrepancy. In *Proceedings of STOC*, 2020.
- [BKPP18] Gerdus Benade, Aleksandr M. Kazachkov, Ariel D. Procaccia, and Christos-Alexandros Psomas. How to Make Envy Vanish Over Time. In *Proceedings of EC 2018*, pages 593–610, 2018.
- [BS19] Nikhil Bansal and Joel H. Spencer. On-Line Balancing of Random Inputs. *CoRR*, abs/1903.06898, 2019.
- [Cha01] Bernard Chazelle. *The discrepancy method: randomness and complexity*. Cambridge University Press, 2001.
- [CST⁺14] William Chen, Anand Srivastav, Giancarlo Travaglino, et al. *A panorama of discrepancy theory*, volume 2107. Springer, 2014.
- [DFGGR19] Raaz Dwivedi, Ohad N. Feldheim, Ori Gurel-Gurevich, and Aaditya Ramdas. The power of online thinning in reducing discrepancy. *Probability Theory and Related Fields*, 2019.
- [DGLN16] Daniel Dadush, Shashwat Garg, Shachar Lovett, and Aleksandar Nikolov. Towards a constructive version of banaszczyk’s vector balancing theorem. In *Proceedings of APPROX/RANDOM*, 2016.
- [DNTT18] Daniel Dadush, Aleksandar Nikolov, Kunal Talwar, and Nicole Tomczak-Jaegermann. Balancing Vectors in Any Norm. In *Proceedings of FOCS*, 2018.
- [ES18] Ronen Eldan and Mohit Singh. Efficient algorithms for discrepancy minimization in convex sets. *Random Struct. Algorithms*, 53(2):289–307, 2018.
- [Gia97] Apostolos A Giannopoulos. On some vector balancing problems. *Stud. Math.*, 122(3):225–234, 1997.
- [Glu89] 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.
- [JKS19] Haotian Jiang, Janardhan Kulkarni, and Sahil Singla. Online Geometric Discrepancy for Stochastic Arrivals with Applications to Envy Minimization. *arXiv:1910.01073*, 2019.
- [LM15] Shachar Lovett and Raghu Meka. Constructive Discrepancy Minimization by Walking on the Edges. *SIAM J. Comput.*, 44(5):1573–1582, 2015.
- [LRR17] Avi Levy, Harishchandra Ramadas, and Thomas Rothvoss. Deterministic Discrepancy Minimization via the Multiplicative Weight Update Method. In *Proceedings of IPCO 2017*, pages 380–391, 2017.
- [Mat09] Jiri Matousek. *Geometric discrepancy: An illustrated guide*, volume 18. SSBM, 2009.
- [MN15] Jiří Matoušek and Aleksandar Nikolov. Combinatorial discrepancy for boxes via the γ_2 norm. In *Proceedings of SoCG 2015*, pages 1–15, 2015.
- [MNT14] Jiří Matoušek, Aleksandar Nikolov, and Kunal Talwar. Factorization norms and hereditary discrepancy. *CoRR*, abs/1408.1376, 2014.
- [Nik17] Aleksandar Nikolov. Tighter bounds for the discrepancy of boxes and polytopes. *CoRR*, abs/1701.05532, 2017.
- [Rot14] Thomas Rothvoß. Constructive Discrepancy Minimization for Convex Sets. In *Proceedings of FOCS 2014*, pages 140–145, 2014.
- [Spe77] Joel Spencer. Balancing games. *J. Comb. Theory, Ser. B*, 23(1):68–74, 1977.
- [Spe85] Joel Spencer. Six standard deviations suffice. *Trans. Am. Math. Soc.*, 289(2):679–706, 1985.
- [Spe87] Joel H. Spencer. *Ten lectures on the probabilistic method*, volume 52. Society for Industrial and Applied Mathematics Philadelphia, 1987.
- [Ver18] Roman Vershynin. *High-Dimensional Probability*, volume 1. Cambridge University Press, 2018.
- [Wai19] Martin J. Wainwright. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2019.