Balancing Vectors in Any Norm

Daniel Dadush *[1], Aleksandar Nikolov †[2], Kunal Talwar ‡[3], and Nicole Tomczak-Jaegermann §[4]

[1]Centrum Wiskunde & Informatica.
[2]University of Toronto.
[3]Google Brain.
[4]University of Alberta.

September 23, 2018

Abstract

In the vector balancing problem, we are given symmetric convex bodies C and K in $\mathbb{R}^n$, and our goal is to determine the minimum number $\beta \geq 0$, known as the vector balancing constant from C to K, such that for any sequence of vectors in C there always exists a signed combination of them lying inside βK. Many fundamental results in discrepancy theory, such as the Beck-Fiala theorem (Discrete Appl. Math '81), Spencer's "six standard deviations suffice" theorem (Trans. Amer. Math. Soc '85) and Banaszczyk's vector balancing theorem (Random Structures & Algorithms '98) correspond to bounds on vector balancing constants.

The above theorems have inspired much research in recent years within theoretical computer science, from the development of efficient polynomial time algorithms for matching existential vector balancing guarantees, to their applications in the context of approximation algorithms. In this work, we show that all vector balancing constants admit "good" approximate characterizations, with approximation factors depending only polylogarithmically on the dimension n. First, we show that a volumetric lower bound due to Banaszczyk is tight within a $O(\log n)$ factor. Our proof is algorithmic, and we show that Rothvoss's (FOCS '14) partial coloring algorithm can be analyzed to obtain these guarantees. Second, we present a novel convex program which encodes the "best possible way" to apply Banaszczyk's vector balancing theorem for bounding vector balancing constants from above, and show that it is tight within an $O(\log^{2.5} n)$ factor. This also directly yields a corresponding polynomial time approximation algorithm both for vector balancing constants, and for the hereditary discrepancy of any sequence of vectors with respect to an arbitrary norm.

Our results yield the first guarantees which depend only polylogarithmically on the dimension of the norm ball K. All prior works required the norm to be polyhedral and incurred a dependence of $O(\sqrt{\log m})$, where m is the number of facets. Our techniques rely on a novel combination of techniques from convex geometry and discrepancy theory. In particular, we give a new way to show lower bounds on Gaussian measures using only volumetric information, which may be of independent interest.

Keywords. Discrepancy, Convex Geometry, Gaussian Measure, M-ellipsoid, K-convexity.

*Email: dadush@cwi.nl. Supported by NWO Veni grant 639.071.510.
†Email: anikolov@cs.toronto.edu.
‡Email: kunal@google.com.
§Email: nicole.tomczak@ualberta.ca.

1 Introduction

The discrepancy of a set system is defined as the minimum, over the set of ± 1 colorings of the elements, of the imbalance between the number of $+1$ and -1 elements in the most imbalanced set. Classical combinatorial discrepancy theory studies bounds on the discrepancy of set systems, in terms of their structure. The tools developed for deriving bounds on the discrepancy of set systems have found many applications in mathematics and computer science [Mat99, Cha91], from the study of pseudorandomness, to communication complexity, and most recently, to approximation algorithms and privacy. Here we study a geometric generalization of combinatorial discrepancy, known as vector balancing, which captures some of the most powerful techniques in the area, and is of intrinsic interest.

Vector Balancing In many instances, the best known techniques for finding good bounds in combinatorial discrepancy were derived by working with more general vector balancing problems, where convex geometric techniques can be applied. Given symmetric convex bodies $C, K \subseteq \mathbb{R}^n$, the vector balancing constant of C into K is defined as

$$\mathrm{vb}(C,K) \triangleq \sup\left\{ \min_{x\in\{-1,1\}^N} \Big\| \sum_{i=1}^{N} x_i u_i \Big\|_K : N \in \mathbb{N}, u_1, \ldots, u_N \in C \right\},$$

where $\|x\|_K := \min\{s \geq 0 : x \in sK\}$ is the norm induced by K.

As an example, one may consider Spencer's "six standard deviations" theorem [Spe85], independently obtained by Gluskin [Glu89], which states that every set system on n points and n sets can be colored with discrepancy at most $O(\sqrt{n})$. In the vector balancing context, the more general statement is that $\mathrm{vb}(B_\infty^n, B_\infty^n) = O(\sqrt{n})$ (also proved in [Spe85, Glu89]), where we use the notation $B_p^n = \{x \in \mathbb{R}^n : \|x\|_p \leq 1\}$, $p \in [1, \infty]$, to denote the unit ball of the ℓ_p norm. To encode Spencer's theorem, we simply represent the set system using its incidence matrix $U \in \{0,1\}^{n\times n}$, where $U_{ji} = 1$ if element i is in set j and 0 otherwise. Here the columns of U have ℓ_∞ norm 1, and thus the sign vector $x \in \{-1,1\}^n$ satisfying $\|Ux\|_\infty = O(\sqrt{n})$ indeed yields the desired coloring.

In fact, vector balancing was studied earlier, and independently from combinatorial discrepancy. In 1963 Dvoretzky posed the general problem of determining $\mathrm{vb}(K,K)$ for a given symmetric convex body K. The more general version with two different bodies was introduced by Barany and Grinberg [BG81] who proved that for any symmetric convex body K in $\mathbb{R}^n$, $\mathrm{vb}(K,K) \leq n$. In addition to Spencer's theorem, as described above, many other fundamental discrepancy bounds, as well as conjectured bounds, can be stated in terms of vector balancing constants. The Beck-Fiala theorem, which bounds the discrepancy of any t-sparse set system by $2t-1$, i.e. where each element appears in at most t-sets, can be recovered from the bound $\mathrm{vb}(B_1^n, B_\infty^n) < 2$ [BF81]. The Beck-Fiala conjecture, which asks whether the bound for t-sparse set systems can be improved to $O(\sqrt{t})$, is generalized by the Kómlos conjecture [Spe94], which asks whether $\mathrm{vb}(B_2^n, B_\infty^n) = O(1)$. One of the most important vector balancing bounds is due to Banaszczyk [Ban98], who proved that for any convex body $K \subseteq \mathbb{R}^n$ of Gaussian measure $1/2$, one has the bound $\mathrm{vb}(B_2^n, K) \leq 5$. In particular, this implies the bound of $\mathrm{vb}(B_2^n, B_\infty^n) = O(\sqrt{\log n})$ for the Kómlos conjecture.

Hereditary Discrepancy. While vector balancing gives useful worst-case bounds, one is often interested in understanding the discrepancy guarantees one can get for instances derived from a fixed set of vectors, known as hereditary discrepancy. Given vectors $(u_i)_{i=1}^N$ in $\mathbb{R}^n$, the discrepancy and hereditary discrepancy with respect to a symmetric convex body $K \subseteq \mathbb{R}^n$ are defined as:

$$\operatorname{disc}((u_i)_{i=1}^N, K) \triangleq \min_{\varepsilon_1,\dots,\varepsilon_N \in \{-1,1\}} \Big\| \sum_{i=1}^N \varepsilon_i u_i \Big\|_K;$$
$$\operatorname{hd}((u_i)_{i=1}^N, K) \triangleq \max_{S \subseteq [N]} \operatorname{disc}((u_i)_{i\in S}, K).$$

When convenient, we will also use the notation $\operatorname{hd}(U, K) := \operatorname{hd}((u_i)_{i=1}^N, K)$, where $U := (u_1, \dots, u_N) \in \mathbb{R}^{n\times N}$, and $\operatorname{disc}(U_S, K) := \operatorname{disc}((u_i)_{i\in S}, K)$ for any subset $S \subseteq [N]$. In the context of set systems, ℓ_∞ hereditary discrepancy corresponds to the worst-case discrepancy of any element induced subsystem, which gives a robust notion of discrepancy, and can be seen as a measure of the complexity of the set system. As an interesting example, a set system has ℓ_∞ hereditary discrepancy 1 if and only if its incidence matrix is totally unimodular [GH62].

Beyond set systems, hereditary discrepancy can also usefully bound the worst-case "error" required for rounding a fractional LP solution to an integral one. More precisely, given any solution $y \in \mathbb{R}^n$ to a linear programming relaxation $Ax \le b$, $x \in [0,1]^n$, with $A \in \mathbb{R}^{m\times n}, b \in \mathbb{R}^m$, of a binary IP, and given any norm $\|\cdot\|$ on $\mathbb{R}^m$ measuring "constraint violation", one can ask what guarantees can be given on $\min_{x\in\{0,1\}^n} \|A(y-x)\|$? Using a well-known reduction of Lovász, Spencer and Vesztergombi [LSV86], this error can be bounded by $\operatorname{hd}(A, K)$ where K is the unit ball of $\|\cdot\|$. Furthermore, this reduction guarantees that x agrees with y on its integer coordinates. Note that we have the freedom to choose the norm $\|\cdot\|$ so that the error bounds meaningfully relate to the structure of the problem. Indeed, much work has been done on the achievable "error profiles" one can obtain algorithmically, e.g. for which $\Delta \in \mathbb{R}^m_{>0}$ we can always find $x \in \{0,1\}^m$ satisfying $|A(y-x)| \le \Delta$, $\forall y \in [0,1]^m$? Note that the feasibility of an error profile can be recovered from a bound of 1 on the hereditary discrepancy with respect to the weighted ℓ_∞ norm $\|y-x\|_\Delta = \max_{i\in[m]} |y_i - x_i|/\Delta_i$. Indeed, in many instances, this is (at least implicitly) how these bounds are proved. These error profile bounds have been fruitfully leveraged for problems where small "additive violations" to the constraints are either allowed or can be repaired. In particular, they were used in for the recent $O(\log n)$-additive approximation for bin packing [HR17], an additive approximation scheme for the train delivery problem [Rot12], and additive approximations of the degree bounded matroid basis problem [BN16].

Discrepancy Minimization. The original proofs of many of the aforementioned discrepancy upper bounds were existential, and did not come with efficient algorithms capable of constructing the requisite low discrepancy colorings. Starting with the breakthrough work of Bansal [Ban10], who gave a constructive version of Spencer's theorem using random walk and semidefinite programming techniques, nearly all known bounds have been made algorithmic in the last eight years.

One of the most important discrepancy minimization techniques is Beck's partial coloring method, which covers most of the above discrepancy results apart from Banaszczyk's vector balancing theorem. This method was first primarily applied to ℓ_∞ discrepancy minimization problems of the form

$$\min_{x\in\{-1,1\}^n} \Big\| \sum_{i=1}^n x_i v_i \Big\|_\infty, \text{ where } (v_i)_{i=1}^n \in \mathbb{R}^m.$$

As before, the goal is not to solve such problems near-optimally but instead to find solutions satisfying a guaranteed error bound. The partial coloring method solves this problem in phases, where at each phase it "colors" (i.e. sets to ± 1) at least a constant fraction of the remaining uncolored variables. This yields $O(\log n)$ partial coloring phases, where the discrepancy of the full coloring is generally bounded by the sum of discrepancies incurred in each phase. The existence of low discrepancy partial colorings, i.e. which color half the variables, was initially established via the pigeon hole principle and arguments based on the probabilistic and the entropy method. In particular, the entropy method gave a general sufficient condition for the feasibility of any error profile (as above) with respect to partial colorings. This method was made constructive by Lovett and Meka [LM12] using random walk techniques. These techniques were further generalized by Giannopoulos [Gia97b] to the general vector balancing setting using Gaussian measure. Precisely, he showed that if a symmetric convex body $K \subseteq \mathbb{R}^n$ has Gaussian measure at least 2^{-cn}, for c small enough, then for any sequence of vectors $v_1, \ldots, v_n \in B_2^n$, there exists a partial coloring $x \in \{-1, 0, 1\}^n$, having support at least $n/2$, such that $\sum_{i=1}^n x_i v_i \in O(1)K$. This method was made constructive by Rothvoss [Rot14], using a random projection algorithm, and later by Eldan and Singh [ES14] who used the solution of a random linear maximization problem. An important difference between the constructive and existential partial coloring methods, is that the constructive methods only guarantee that the "uncolored" coordinates of a partial coloring x are in $(-1, 1)$ instead of equal to 0. This relaxation seems to make the constructive methods more robust, i.e. the conditions needed for such "fractional" partial colorings are somewhat milder, without having noticeable drawbacks in most applications.

The main alternative to the partial coloring method comes from Banaszczyk's vector balancing theorem [Ban98]. Banaszczyk's method proves the existence of a full coloring when K has gaussian measure $1/2$, in contrast to Giannopoulos's result which gives a partial coloring but requires measure only 2^{-cn}. Banaszczyk's method was only very recently made constructive in the sequence of works [BDG16, DGLN16, BDGL18]. In particular, [DGLN16] showed an equivalence of Banaszczyk's theorem to the existence of certain subgaussian signing distributions, and [BDGL18] gave a random walk-based algorithm to build such distributions.

1.1 Approximating Vector Balancing and Hereditary Discrepancy

Given the powerful tools that have been developed above, a natural question is whether they can be extended to get nearly optimal bounds for any vector balancing or hereditary discrepancy problem. More precisely, we will be interested in the following computational and mathematical questions:

1. Given vectors $(u_i)_{i=1}^N$ and a symmetric convex body K in $\mathbb{R}^n$, can we (a) efficently compute a coloring whose K-discrepancy is approximately bounded by $\mathrm{hd}((u_i)_{i=1}^N, K)$? (b) efficiently approximate $\mathrm{hd}((u_i)_{i=1}^N, K)$?

2. Given two symmetric convex bodies $C, K \subseteq \mathbb{R}^n$, does $\mathrm{vb}(C, K)$ admit a "good" characterization? Namely, are there simple certificates which certify nearly tight upper and lower bounds on $\mathrm{vb}(C, K)$?

To begin, a few remarks are in order. Firstly, question 2 can be inefficiently encoded as question 1b, by letting $(u_i)_{i=1}^N$ denote a sufficiently fine net of C. Thus "good" characterizations for hereditary discrepancy transfer over to vector balancing, and thus we restrict for now the discussion to the former. For question 1a, one may be tempted to ask whether we can directly compute a coloring whose K-discrepancy is approximately $\mathrm{disc}((u_i)_{i=1}^N, K)$ instead of $\mathrm{hd}((u_i)_{i=1}^N, K)$. Unfortunately, even for $K = B_\infty^n$ and $(u_i)_{i=1}^n \in [-1, 1]^n$, it was shown in [CNN11] that it is NP-hard to distinguish whether $\mathrm{disc}((u_i)_{i=1}^n, B_\infty^n)$ is 0 or $\Omega(\sqrt{n})$ (note that $O(\sqrt{n})$ is guaranteed by Spencer's theorem), thus one cannot hope for any non-trivial approximation guarantee in this context.

We now discuss prior work on these questions and then continue with our main results.

Prior work. For both questions, prior work has mostly dealt with the case of ℓ_∞ or ℓ_2 discrepancy. Bounds on vector balancing constants for some combinations of ℓ_p to ℓ_q have also been studied, as described earlier, however without a unified approach. The question of obtaining near-optimal results for general vector balancing and hereditary discrepancy problems has on the other hand not been studied before.

In terms of coloring algorithms, Bansal [Ban10] gave a partial coloring based random walk algorithm which on $U \in \mathbb{R}^{m\times n}$, produces a full coloring of ℓ_∞ discrepancy $O(\sqrt{\log m \log \mathrm{rk}(U)}\,\mathrm{hd}(U, B_\infty^m))$, where $\mathrm{rk}(U)$ is the rank of U. Recently, Larsen [Lar17] gave an algorithm for the ℓ_2 norm achieving discrepancy $O(\sqrt{\log(\mathrm{rk}(U))}\,\mathrm{hd}(U, B_2^m))$.

In terms of certifying lower bounds on $\mathrm{hd}(U, B_\infty^m))$, the main tool has been the so-called determinant lower bound of [LSV86], where it was shown that

$$\mathrm{hd}(U, B_\infty^m) \geq \mathrm{detLB}(U) := \max_k \max_B \frac{1}{2} |\det(B)|^{1/k}$$

where the maximum is over $k \times k$ submatrices B of U. Matousek [Mat11], built upon the results of [Ban10] to show that

$$\mathrm{hd}(U, B_\infty^m) \leq O(\sqrt{\log m}\log^{3/2}(\mathrm{rk}(U))\,\mathrm{detLB}(U)).$$

For certifying tight upper bounds, [NT15, MNT18] showed that γ_2 norm of U, defined by

$$\gamma_2(U) := \min\left\{\|A\|_{2\to\infty}\|B\|_{1\to 2} : U = AB, A \in \mathbb{R}^{m\times k}, B \in \mathbb{R}^{k\times n}, k \in \mathbb{N}\right\}$$

where $\|A\|_{2\to\infty}$ is the maximum ℓ_2 norm of any row of A, and $\|B\|_{1\to 2}$ is the maximum ℓ_2 norm of any column of B, satisfies

$$\Omega(\gamma_2(U)/\log(\mathrm{rk}(U))) \leq \mathrm{detLB}(U) \leq \mathrm{hd}(U, B_\infty^m) \leq O(\sqrt{\log m}\,\gamma_2(U)) \tag{1}$$

which implies a $O(\sqrt{\log m}\log \mathrm{rk}(U))$ approximation to ℓ_∞ hereditary discrepancy. For the context of ℓ_2, it was shown in [NT15] that a relaxation of γ_2 yields an $O(\log \mathrm{rk}(U))$-approximation to $\mathrm{hd}(U, B_2^m)$. We note that part of the strategy of [NT15, MNT18] is to replace the ℓ_∞ norm via an averaged version of ℓ_2, where one optimizes over the averaging coefficients, which makes the ℓ_2 norm by itself an easier special case.

Moving to general norms. While at first glance it may seem that the above techniques for ℓ_∞ do not apply to more general norms, this is in some sense deceptive. Notwithstanding complexity considerations, every norm can be isometrically embedded into ℓ_∞, where in particular any polyhedral norm with m facets can be embedded into B_∞^m. Vice versa, starting from $U \in \mathbb{R}^{m\times N}$, with $\mathrm{rk}(U) = n$ and rank factorization $U = AB$, is it direct to verify $\mathrm{hd}(U, B_\infty^m) = \mathrm{hd}(B, K)$, where $K = \{x \in \mathbb{R}^n : |Ax| \leq 1\}$ is an n-dimensional symmetric polytope with m facets. Thus, for any $U \in \mathbb{R}^{n\times N}$, one can equivalently restate the guarantees of [Ban10] as yielding colorings of discrepancy $O(\sqrt{\log m \log n}\,\mathrm{hd}(U, K))$ and of [MNT18] as a $O(\sqrt{\log m}\log n)$ approximation to $\mathrm{hd}(U, K)$ for any n-dimensional symmetric polytope K with m facets. A natural question is therefore whether there exist corresponding coloring and approximation algorithms whose guarantees depend only polylogarithmically on the dimension of the norm and not on the complexity of its representation.

We note that polynomial bounds in n for general K can be achieved by simply approximating K by a sandwiching ellipsoid $E \subseteq K \subseteq \sqrt{n}E$ and applying the corresponding results for ℓ_2, which yield $O(\sqrt{n \log n})$ coloring and $O(\sqrt{n}\log n)$ approximations guarantees respectively. Interestingly, these guarantees are identical to what can be achieved by replacing K by a symmetric polytope with 3^n facets, which can achieve a sandwiching factor of 2, and applying the ℓ_∞ results.

1.2 Results

Our main results are that such polylogarithmic approximations are indeed possible. In particular, given $U \in \mathbb{R}^{n\times N}$ and a symmetric convex body $K \subseteq \mathbb{R}^n$ (by an appropriate oracle), we give randomized polynomial time algorithms for computing colorings of discrepancy $O(\log n \operatorname{hd}(U,K))$ and approximating $\operatorname{hd}(U,K)$ up to $O(\log^{2.5} n)$ factor. Furthermore, if K is a polyhedral norm with at most m facets, our approximation algorithm for $\operatorname{hd}(U,K)$ always achieves a tighter approximation factor than the γ_2 bound, and hence gives an $O(\min\left\{\log n\sqrt{\log m}, \log^{2.5} n\right\})$ approximation. To achieve these results, we first show that Rothvoss' partial coloring algorithm [Rot14] is nearly optimal for general hereditary discrepancy by showing near-tightness with respect to a volumetric lower bound of Banaszczyk [Ban93]. Second, we show that the "best possible way" to apply Banaszczyk's vector balancing theorem [Ban98] for the purpose of upper bounding $\operatorname{hd}(U,K)$ can be encoded as a convex program, and prove that this bound is tight to within an $O(\log^{2.5} n)$ factor. As a consequence, we show that Banaszczyk's theorem is essentially "universal" for vector balancing. To analyze these approaches we rely on a novel combination of tools from convex geometry and discrepancy. In particular, we give a new way to prove lower bounds on Gaussian measure using only volumetric information, which could be of independent interest. Furthermore, we make a natural geometric conjecture which would imply that Rothvoss' algorithm is (in a hereditary sense) optimal for finding partial colorings in any norm, and prove the conjecture for the special case of ℓ_2.

Comparing to prior work, our coloring and hereditary discrepancy approximation algorithms give uniformly better (or at at least no worse) guarantees in almost every setting which has been studied. Furthermore our methods provide a unified approach for studying discrepancy in arbitrary norms, which we expect to have further applications.

Interestingly, our results imply a tighter relationship between vector balancing and hereditary discrepancy than one might initially expect. That is, neither the volumetric lower bound we use nor our factorization based upper bound "see" the difference between them. More precisely, both bounds remain invariant when replacing $\operatorname{hd}(U,K)$ by $\operatorname{vb}(\operatorname{conv}\{\pm u_i : i \in [N]\}, K)$. This has the relatively non-obvious implication that

$$\operatorname{hd}(U,K) \leq \operatorname{vb}(\operatorname{conv}\{\pm u_i : i \in [N]\}, K) \leq O(\log n)\operatorname{hd}(U,K). \tag{2}$$

We believe it is an interesting question to understand whether a polylogarithmic separation indeed exists between the above quantities (we are currently unaware of any examples), as it would give a tangible geometric obstruction for tighter approximations.

1.3 Techniques

Starting with hereditary discrepancy, to push beyond the limitations of prior approaches the first two tasks at hand are: (1) find a stronger lower bound and (2) develop techniques to avoid the "union bound". Fortunately, a solution to the first problem was already given by Banaszczyk[Ban93], which we present in slightly adapted form below.

Lemma 1 (Volume Lower Bound). *Let $U = (u_1, \ldots, u_N) \in \mathbb{R}^{n\times N}$ and $K \subseteq \mathbb{R}^n$ be a symmetric convex body. For $S \subseteq [N]$, let U_S denote the columns of U in S. For $k \in [n]$, define*

$$\operatorname{volLB}_k^{\mathrm{h}}((u_i)_{i=1}^N, K) \triangleq \operatorname{volLB}_k^{\mathrm{h}}(U,K) \triangleq \max_{S\subseteq[N],|S|=k} \operatorname{vol}_k(\{x \in \mathbb{R}^k : U_S x \in K\})^{-1/k}. \tag{3}$$

Then, we have that

$$\operatorname{volLB}^{\mathrm{h}}((u_i)_{i=1}^N, K) \triangleq \operatorname{volLB}^{\mathrm{h}}(U,K) \triangleq \max_{k\in[n]} \operatorname{volLB}_k^{\mathrm{h}}(U,K) \leq \operatorname{hd}(U,K). \tag{4}$$

A formal proof of the above is given in the preliminaries (see section 2.1). At a high level, the proof is a simple covering argument, where it is argued that for any subset S, $|S| = k$, every point in $[0,1]^k$ is at distance at most $\mathrm{hd}(U, K)$ from $\{0,1\}^k$ under the norm induced by $C := \left\{x \in \mathbb{R}^k : U_S x \in K\right\}$. Equivalently an $\mathrm{hd}(U, K)$ scaling of C placed around the points of $\{0,1\}^k$ cover $[0,1]^k$, and hence by a standard lattice argument must have volume at least that of $[0,1]^k$, namely 1. This yields the desired lower bound after rearranging.

We note that the volume lower bound extends in the obvious way to vector balancing. In particular, for two symmetric convex bodies $C, K \subseteq \mathbb{R}^n$,

$$\mathrm{volLB}^{\mathrm{h}}(C, K) \triangleq \sup\left\{\mathrm{volLB}((u_i)_{i=1}^k, K) : k \in [n], u_1, \ldots, u_k \in C\right\} \geq \mathrm{vb}(C, K). \tag{5}$$

The above lower bound can be substantially stronger than the determinant lower bound for ℓ_∞ discrepancy. As a simple example, let $U \in \mathbb{R}^{2^n \times n}$ be the matrix having a row for each vector in $\{-1,1\}^n$. Since U has rank n, the determinant lower bound is restricted to $k \times k$ matrices for $k \in [n]$. Hadamard's inequality implies for any $k \times k$ matrix B with ± 1 entries that $|\det(B)|^{1/k} \leq \sqrt{k} \leq \sqrt{n}$. A moment's thought however, reveals that for $x \in \mathbb{R}^n$, $\|Ux\|_\infty = \|x\|_1$ and hence any coloring $x \in \{-1,1\}$ must have discrepancy $\|x\|_1 = n$. Using the previous logic, the volume lower bound to the full system yields by standard estimates

$$\mathrm{volLB}(U, B_\infty^m) \geq \mathrm{vol}_n(\{x \in \mathbb{R}^n : \|x\|_1 \leq 1\})^{-1/n} = \mathrm{vol}_n(B_1^n)^{-1/n} = (n!/2^n)^{1/n} \geq n/(2e),$$

which is essentially tight.

From Volume to Coloring. The above example gives hope that the volume lower bound can circumvent a dependency on the facet complexity of the norm. Our first main result, shows that indeed this is the case:

Theorem 2 (Tightness of the Volume Lower Bound)**.** *For any $U \in \mathbb{R}^{n \times N}$ and symmetric convex body K in $\mathbb{R}^n$, we have that*

$$\mathrm{volLB}^{\mathrm{h}}(U, K) \leq \mathrm{hd}(U, K) \leq O(\log n)\, \mathrm{volLB}^{\mathrm{h}}(U, K)\ , \tag{6}$$

Furthermore, there exists a randomized polynomial time algorithm that computes a coloring of U with K-discrepancy $O(\log n\, \mathrm{volLB}^{\mathrm{h}}(U, K))$, given a membership oracle for K.

We note that the above immediately implies the corresponding approximate tightness of the volume lower bound for vector balancing. The above bound can also be shown to be tight. In particular, the counterexample to the 3-permutations conjecture from [NNN12], which has ℓ_∞ discrepancy $\Omega(\log n)$, can be shown to have volume lower bound $O(1)$. The computations for this are somewhat technical, so we defer a detailed discussion to the full version. As mentioned previously, an interesting property of the volume lower bound is its invariance under taking convex hulls, namely $\mathrm{volLB}^{\mathrm{h}}(\mathrm{conv}\{\pm U\}, K) = \mathrm{volLB}^{\mathrm{h}}(U, K)$. In combination with Theorem 2, this establishes the claimed inequality 2. This invariance is proved in section 6.1, where we use a theorem of Ball [Bal88] to show that the volume lower bound is essentially convex, and hence maximized at extreme points.

Our proof of Theorem 2 is algorithmic, and relies on iterated applications of Rothvoss's partial coloring algorithm. We now explain our high level strategy as well as the differences with respect to prior approaches.

For simplicity of the presentation, we shall assume that $U = (e_1, \ldots, e_n) \in \mathbb{R}^{n \times n}$ and that the volume lower bound $\mathrm{volLB}^{\mathrm{h}}((e_i)_{i=1}^n, K) = 1$. This can be (approximately) achieved by applying a standard reduction to the case where U is non-singular, so $N \leq n$, "folding" U into K, and appropriately guessing the volume lower bound (see section 3 for full details).

For any subset $S \subseteq [n]$, let $K_S := \{x \in K : x_i = 0, i \in [n] \setminus S\}$ denote the coordinate section of K induced by S. Since the vectors of U now correspond to the coordinate basis, it is direct to verify that

$$\mathrm{volLB}^{\mathrm{h}}((e_i)_{i=1}^n, K) = \max_{S \subseteq [n], k := |S|} \mathrm{vol}_k(K_S)^{-1/k}.$$

In particular, the assumption $\mathrm{volLB}^{\mathrm{h}}((e_i)_{i=1}^n, K) = 1$ implies that

$$\mathrm{vol}_{|S|}(K_S) \geq 1, \ \forall S \subseteq [n]. \tag{7}$$

Under this condition, our goal can now be stated as finding a coloring $x \in \{-1,1\}^n \in O(\log n)K$.

When K is a symmetric polytope $|Ax| \leq 1$, with m facets, Bansal [Ban10] uses a "sticky" random walk on the coordinates, where the increments are computed via an SDP to guarantee that their variance along any facet is at most $\mathrm{hd}((e_i)_{i=1}^n, K)^2$, while the variance along all (active) coordinate directions is at least 1 (i.e. we want to hit cube constraints faster). As this only gives probabilistic error guarantees for each constraint in isolation, a union bound is used to get a global guarantee, incurring the $O(\sqrt{\log m})$ dependence.

To avoid the "union bound", we instead use Rothvoss's partial coloring algorithm, which simply samples a random Gaussian vector $X \in \mathbb{R}^n$ and computes the closest point in Euclidean distance x to X in $K \cap [-1,1]^n$ as the candidate partial coloring. As long as K has "large enough" Gaussian measure, Rothvoss shows that x has at least a constant fraction of its components at ± 1. While this method can in essence better leverage the geometry of K than Bansal's method (in particular, it does not need an explicit description of K), it is apriori unclear why Gaussian measure should be large enough in the present context.

Our main technical result is that if all the coordinate sections of K have volume at least 1 (i.e. condition 7), then there indeed exists a section of K of dimension close to n, whose Gaussian measure is "large" after appropriately scaling. Specifically, we show that for any $\delta \in (0,1)$, there exists a subspace H of dimension $(1-\delta)n$ such that the Gaussian measure of $2^{O(1/\delta)}(K \cap H)$ is at least $2^{-\delta n}$ (see Theorem 10 for the exact statement). We sketch the ideas in the next subsection.

The existence of a large section of K with not too small Gaussian measure in fact suffices to run Rothvoss's partial coloring algorithm (see Theorem 9). Conveniently, one does not need to know the section explicitly, as its existence is only used in the analysis of the algorithm. Since condition 7 is hereditary, we can now find partial colorings of K-discrepancy $O(1)$ on any subset of coordinates. Thus, applying $O(\log n)$ partial coloring phases in the standard way yields the desired full coloring.

A useful restatement of the above is that Rothvoss's algorithm can always find partial colorings with discrepancy $O(1)$ times the volume lower bound. We note that this guarantee is a natural by-product of algorithm (once one has guessed the appropriate scaling), which does not need to be explicitly enforced as in Bansal's algorithm.

Finding a section with large Gaussian measure. We now sketch how to find a section of K of large Gaussian measure the assumption that $\mathrm{vol}_{|S|}(K_S) \geq 1, \forall S \subseteq [n]$. The main tool we require is the M-ellipsoid from convex geometry [Mil86]. The M-ellipsoid E of K is an ellipsoid which approximates K well from the perspective of covering, that is $2^{O(n)}$ translates of E suffice to cover K and vice versa.

The main idea is to use the volumetric assumption to show that the largest $(1-\delta)n$ axes of E, for $\delta \in (0,1)$ of our choice, have length at least $\sqrt{n}2^{-O(1/\delta)}$, and then use the subspace generated by these axes for the section of K we use. On this subspace H, we have that a $2^{O(1/\delta)}$ scaling of $E \cap H$ contains the $\sqrt{n}$ ball, and thus by the covering estimate $2^{O(n)}$ translates of $2^{O(1/\delta)}(K \cap H)$ covers the $\sqrt{n}$ ball. Since the $\sqrt{n}$ ball on H has Gaussian measure at least $1/2$, the prior covering estimate indeed implies that $2^{O(1/\delta)}(K \cap H)$ has Gaussian measure $2^{-O(n)}$, noting that shifting $2^{O(1/\delta)}(K \cap H)$ away from the origin only reduces Gaussian measure. Using an M-ellipsoid with appropriate regularity properties (see Theorem 11), one can scale $K \cap H$ by another $2^{O(1/\delta)}$ factor, so that the preceding argument yields Gaussian measure at least $2^{-\delta n}$.

We now explain why the axes of E are indeed long. By the covering estimates, for any $S \subseteq [n]$, $|S| = \delta n$, the sections E_S and K_S satisfy

$$\mathrm{vol}_{\delta n}(E_S)^{1/\delta n} \geq 2^{-O(1/\delta)} \, \mathrm{vol}_{\delta n}(K_S)^{1/\delta n} \geq 2^{-O(1/\delta)},$$

where the last inequality is by assumption. Using a form of the restricted invertibility principle for determinants (see Lemma 7), one can show that if all coordinate sections of E of dimension δn have large volume, then so does every section of E of the same dimension. Precisely, one gets that

$$\min_{\dim(W)=\delta n} \mathrm{vol}_{\delta n}(E \cap W)^{1/\delta n} \geq \binom{n}{\delta n}^{-1/\delta n} \min_{|S|=\delta n} \mathrm{vol}_{\delta n}(E_S)^{1/\delta n} \geq 2^{O(-1/\delta)}.$$

In particular, the above implies that the geometric average of the *shortest* δn axes of E (corresponding to the minimum volume section above), must have length $\sqrt{n}2^{-O(1/\delta)}$ since the ball of volume 1 in dimension δn has radius $\Omega(\sqrt{\delta n})$. But then, the longest $(1-\delta)n$ axes all have have length $\sqrt{n}2^{-O(1/\delta)}$. This completes the proof sketch.

The Discrepancy of Partial Colorings. Our analysis of Rothvoss's algorithm opens up the tantalizing possibility that it may indeed be optimal for finding partial colorings in a hereditary sense. More precisely, we conjecture that if when run on an instance U with norm ball K, the algorithm almost always produces partial colorings with K-discrepancy at least D, then there exists a subset of S of the columns of U such that every partial coloring of U_S has discrepancy $\Omega(D)$. The starting point for this conjecture is our upper bound of $O(1)\,\mathrm{volLB}^{\mathrm{h}}(U, K)$, on the discrepancy of the partial colorings the algorithm computes. We now provide a purely geometric conjecture, which would imply the above "hereditary optimality" for Rothvoss's algorithm.

As in the last subsection, we may assume that $U = (e_1, \dots, e_n)$ is the standard basis of $\mathbb{R}^n$ and that $\mathrm{volLB}((e_i)_{i=1}^n, K) = 1$. To prove the conjecture, it suffices to show that exists some subset $S \subseteq [n]$ of coordinates, such that all partial colorings have K-discrepancy $\Omega(1)$. For concreteness, let us ask for partial colorings which color at least $|S|/2$ coordinates (the precise constant will not matter). For $x \in [-1,1]^n$, define $\mathrm{bounds}(x) = \{i \in [n] : x_i \in \{-1, 1\}\}$. With this notation, our goal is to find $S \subseteq [n]$, such that $\forall x \in [-1,1]^S$, $|\mathrm{bounds}(x)| \geq |S|/2$, $\|\sum_{i \in S} x_i e_i\|_K \geq \Omega(1)$.

We explain the candidate geometric obstruction to low discrepancy partial colorings, which is a natural generalization of the so-called spectral lower bound for ℓ_2 discrepancy. Assume now that for some subset $S \subseteq [n]$, we have that

$$K_S \subseteq c\sqrt{|S|}B_2^S, \tag{8}$$

where $B_2^S := (B_2^n)_S$, for some constant $c > 0$. Since any partial coloring $x \in [-1,1]^S$, $|\mathrm{bounds}(x)| \geq |S|/2$, clearly has $\|x\|_2 \geq \sqrt{|S|/2}$, we must have that

$$\frac{1}{c\sqrt{2}} \leq \Big\|\sum_{i \in S} x_i e_i\Big\|_{c\sqrt{|S|}B_2^S} \leq \Big\|\sum_{i \in S} x_i e_i\Big\|_{K_S}. \tag{9}$$

In particular, every partial coloring on S has discrepancy at least $\frac{1}{c\sqrt{2}} = \Omega(1)$, as desired.

Given the above, we may now reduce the conjecture to the following natural geometric question:

Conjecture 3 (Restricted Invertibility for Convex Bodies)**.** *There exists an absolute constant $c \geq 1$, such that for any $n \in \mathbb{N}$ and symmetric convex body $K \subseteq \mathbb{R}^n$ of volume at most 1, there exists $S \subseteq [n]$, $S \neq \emptyset$, such that $K_S \subseteq c\sqrt{|S|}B_2^S$.*

To see that this indeed implies the required statement, note that if $\mathrm{volLB}((e_i)_{i=1}^n, K) = 1$, then by definition there exists $A \subseteq [n]$, $|A| \geq 1$, such that $\mathrm{vol}_{|A|}(K_A) \leq 1$. Now applying the above conjecture to K_A yields the desired result.

Two natural relaxations of the conjectures are to ask (1) does it hold for ellipsoids and (2) does it hold for general sections instead of coordinate sections? Our main evidence for this conjecture is that indeed both these statements are true. We note that (1) indeed implies the optimality of Rothvoss's partial coloring algorithm for ℓ_2 discrepancy. Our results here are slightly stronger than (1)+(2), as we in some sense manage to get "halfway there" with coordinates sections, by working with the M-ellipsoid, and only for the last step do we need to resort general sections. We note that the above conjecture is closely related to the Bourgain-Tzafriri restricted invertibility principle, and indeed our proof for ellipsoids reduces to it. We refer the reader to section 3.1 for further details and proofs.

A Factorization Approach for Vector Balancing. While Theorem 2 gives an efficient and approximately optimal method of balancing a given set of vectors, it does not give an efficiently computable tight upper bound on the vector balancing constant or on hereditary discrepancy. Even though we proved that, after an appropriate scaling, the volume lower bound also gives an upper bound on the vector balancing constant, we are not aware of an efficient algorithm for computing the volume lower bound, which is itself a maximum over an exponential number of terms. To address this shortcoming, we study a different approach to vector balancing which relies on applying Banaszczyk's theorem in an optimal way in order to get an efficiently computable, and nearly tight, upper bound on both vector balancing constants and hereditary discrepancy.

Recall that Banaszczyk's vector balancing theorem states that if a body K has Gaussian measure at least $1/2$, then $\mathrm{vb}(B_2^n, K) \leq 5$. In order to apply the theorem to bodies K of small Gaussian measure, we can use rescaling. In particular, if r is the smallest number such that the Gaussian measure of rK is $\frac{1}{2}$, then the theorem tells us that $\mathrm{vb}(B_2^n, K) \leq 5r$. A natural way to use this upper bound for bodies C different from B_2^n is to find a mapping of C into B_2^n, and then use the theorem as above. As an illustration of this idea, let us see how we can get nearly tight bounds on $\mathrm{vb}(B_p^n, B_q^n)$ (the ℓ_p and ℓ_q balls) by applying Banaszczyk's theorem. Let us take an arbitrary sequence of points $u_1, \ldots, u_N \in B_p^n$, and rescale them to define new points $v_i \triangleq u_i/\max\{1, n^{1/2-1/p}\}$. The rescaled points $v_1, \ldots, v_N$ lie in B_2^n and we can apply Banaszczyk's theorem to them and the convex body $K \triangleq L\sqrt{q}n^{1/q}B_q^n$, which has Gaussian measure at least $\frac{1}{2}$ as long as we choose L to be a large enough constant. We get that there exist signs $\varepsilon_1, \ldots, \varepsilon_N \in \{-1, 1\}$ such that

$$\left\|\sum_{i=1}^N \varepsilon_i v_i\right\|_K \leq 5 \iff \left\|\sum_{i=1}^N \varepsilon_i u_i\right\|_q \leq 5L\sqrt{q}\max\{n^{1/q}, n^{1/q+1/2-1/p}\}.$$

In other words, we have that

$$\mathrm{vb}(B_p^n, B_q^n) \leq 5L\sqrt{q}\max\{n^{1/q}, n^{1/q+1/2-1/p}\}.$$

The volume lower bound (Lemmas 1) can be used to show that this bound is tight up to the $O(\sqrt{q})$ factor. Indeed one can show that B_p^n contains n vectors $u_1, \ldots u_n$ such that the matrix

$U \triangleq (u_1, \ldots, u_n)$ has determinant $\det(U) \geq e^{-1} \max\{1, n^{1/2-1/p}\}$ (see [Bal89] or [Nik15]). By standard estimates, $\mathrm{vol}(B_q^n)^{1/n} \geq cn^{1/q}$ for an absolute constant $c > 0$. Plugging these estimates into Lemma 1 shows $\mathrm{vb}(B_p, B_q) \geq c' \max\{n^{1/q}, n^{1/q+1/2-1/p}\}$ for a constant $c' > 0$.

It is easy to see that, unlike the example above, in general simply rescaling C and K and applying Banaszczyk's theorem to the rescaled bodies may not give a tight bound on $\mathrm{vb}(C, K)$. However, we will show that we can get such tight bounds if we expand the class of transformations we allow on C and K from simple rescaling to arbitrary linear transformations. It turns out that the most convenient language for this approach is that of linear operators between normed spaces. We can generalize the notion of a vector balancing constant between a pair of convex bodies to arbitrary linear operators $U : X \to Y$ between two n-dimensional normed spaces X, with norm $\| \cdot \|_X$), and Y, with norm $\| \cdot \|_Y$), as follows

$$\mathrm{vb}(U) = \sup \left\{ \min_{\varepsilon_1, \ldots, \varepsilon_N \in \{-1,1\}} \left\| \sum_{i=1}^{N} \varepsilon_i U(x_i) \right\|_Y : N \in \mathbb{N}, x_1, \ldots, x_N \in B_X \right\} \tag{10}$$

where $B_X = \{x : \|x\|_X \leq 1\}$ is the unit ball of X. This definition is indeed a generalization of the geometric one. If C and K are two centrally symmetric convex bodies in $\mathbb{R}^n$, and we define the corresponding normed spaces $X_C = (\mathbb{R}^n, \| \cdot \|_C)$ and $X_K = (\mathbb{R}^n, \| \cdot \|_K)$, then the vector balancing constant $\mathrm{vb}(I)$ of the formal identity operator $I : X_C \to X_K$ recovers $\mathrm{vb}(C, K)$. However, the more abstract setting makes it plain that a simple rescaling is not the right approach to applying Banaszczyk's theorem to arbitrary norms: if X is an arbitrary norm, then X and B_2^n may not be defined on the same vector space, and rescaling B_X so that it is a subset of B_2^n does not even make sense. Instead, when dealing with general norms, it becomes very natural to embed B_X into B_2^n via a linear map $T : X \to \ell_2^n$ so that $T(B_X) \subseteq B_2^n$. Our approach is based on this idea, and, in particular, on choosing such a map T optimally.

To formalize the above, we use the ℓ-norm, which has been extensively studied in the theory of operator ideals, and in asymptotic convex geometry (see e.g. [TJ89, Pis89, AAGM15]). For a linear operator $S : \ell_2^n \to Y$ into a normed space Y with norm $\| \cdot \|_Y$, the ℓ-norm of S is defined as

$$\ell(S) \triangleq \left(\int \|S(x)\|_Y^2 d\gamma_n(x) \right)^{1/2},$$

where γ_n is the standard Gaussian measure on $\mathbb{R}^n$. I.e., if Z is a standard Gaussian random variable in $\mathbb{R}^n$, then $\ell(S) = (\mathbb{E}\|S(Z)\|_Y^2)^{1/2}$. It is easy to verify that $\ell(\cdot)$ is a norm on the space of linear operators from ℓ_2^n to Y, for any normed space Y as above. The reason the ℓ-norm is useful to us is the fact that the smallest r for which the set $K = \{x \in \mathbb{R}^n : \|Sx\|_Y \leq r\}$ has Gaussian measure at least $1/2$ is approximately $\ell(S)$, due to the concentration of measure phenomenon.

We now define our main tool: a factorization constant λ, which, for any two n-dimensional normed spaces X and Y and an operator $U : X \to Y$ is defined by

$$\lambda(U) \triangleq \inf\{\ell(S)\|T\| : T : X \to \ell_2^n,\ S : \ell_2^n \to Y, U = ST\}.$$

In other words, $\lambda(U)$ is the minimum of $\ell(S)\|T\|$ over all ways to factor U through ℓ_2^n as $U = ST$. Here $\|T\|$ is the operator norm, equal to $\max\{\|Tx\|_2/\|x\|_X\}$. This definition captures an optimal application of Banaszczyk's theorem. Using the theorem, it is not hard to show that $\mathrm{vb}(U) \leq C\lambda(U)$ for an absolute constant C. Our main result is showing $\mathrm{vb}(U)$ and $\lambda(U)$ are in fact equal up to a factor which is polynomial in $\log n$. To prove this, we formulate $\lambda(U)$ as a convex minimization problem. Such a formulation is important both for our structural results, which rely on Lagrange duality, and also for giving an algorithm to compute $\lambda(U)$ efficiently, and, therefore, approximate

$\mathrm{vb}(U)$ efficiently, which turns out to be sufficient to approximate hereditary discrepancy in arbitrary norms.

The most immediate way to formulate $\lambda(U)$ as an optimization problem is to minimize $\ell(UT^{-1})$ over operators $T : X \to \ell_2^n$ and subject to the constraint $\|T\| \le 1$. Unfortunately, this optimization problem is not convex in T: the value of the objective function is finite for any nonzero T, but infinite for $0 = \frac{1}{2}(T + (-T))$, for example. The key observation that allows us to circumvent this issue is that the objective function is completely determined by the operator $A \triangleq T^*T$, and is in fact convex in A. Here T^* is the dual operator of T (see Section 4.1 for more details). We use $f(A)$ to denote this objective function, i.e. to denote $\ell(UT^{-1})$ where T is an operator such that $T^*T = A$. We give more justification why this function is well-defined and convex in Section 4.3. Then, our convex formulation of $\lambda(U)$ is

$$\begin{aligned}
&\inf f(A)\\
&\text{s.t.}\\
&A : X \to X^*, \|A\| \le 1\\
&A \succ 0.
\end{aligned}$$

Above, X^* is the dual space of X, and $\|A\|$ is the operator norm. The first constraint is equivalent to the constraint $\|T\| \le 1$ where $U = ST$ is the factorization in the definition of $\lambda(U)$. The last constraint says that A should be positive definite, which is important so that A can be written as T^*T and $f(A)$ is well-defined.

We utilize this convex formulation and Lagrange duality to derive a dual formulation of $\lambda(U)$ as a supremum over "dual certificates". Such a formulation is useful in approximately characterizing $\mathrm{vb}(U)$ in terms of $\lambda(U)$ because it reduces our task to relating the dual certificates to the terms in the volume lower bound (3). If we can show that every dual certificate bounds from below one of the terms of the volume lower bound (up to factors polynomial in $\log n$), then we can conclude that $\lambda(U)$ also bounds the volume lower bound from below, and therefore $\mathrm{vb}(U)$ as well.

Before we can give the dual formulation, we need to introduce the dual norm ℓ^* of the ℓ-norm, defined via trace duality: for any linear opartor $R : Y \to \ell_2^n$, let

$$\ell^*(R) \triangleq \sup\{\mathrm{tr}(RS) : S : \ell_2^n \to Y, \ell(S) \le 1\}.$$

The norms ℓ and ℓ^* form a dual pair, and in particular we have

$$\ell(S) = \sup\{\mathrm{tr}(RS) : R : Y \to \ell_2^n, \ell^*(R) \le 1\}.$$

For a finite dimensional space Y, both suprema above are achieved.

The derivation of our dual formulation uses standard tools, but is quite technical due to the complicated nature of the function $f(A)$. We give the formulation for norms X such that $B_X = \mathrm{conv}\{\pm x_1, \dots, \pm x_m\}$. This is without loss of generality since every symmetric convex body can be approximated by a symmetric polytope. The dual formulation is as follows:

$$\begin{aligned}
&\sup \mathrm{tr}((RU(\sum_{i=1}^m p_i x_i \otimes x_i)U^* R^*)^{1/3})^{3/2}\\
\text{s.t.}\quad &R : Y \to \ell_2^n, \ell^*(R) \le 1\\
&\sum_{i=1}^m p_i = 1, \quad p_1, \dots, p_m \ge 0.
\end{aligned}$$

Above $x_i \otimes x_i$ is the rank-1 operator from the dual space X^* to X, given by $(x_i \otimes x_i)(x^*) = \langle x^*, x_i \rangle x_i$.

We relate the volume lower bound to this dual via deep inequalities between the ℓ^* and the ℓ norms (K-convexity), and between the ℓ norm and packing and covering numbers (Sudakov's minoration). Our main result is the theorem below.

Theorem 4. *There exists a constant C such that for any two n-dimensional normed spaces X and Y, and any linear operator $U : X \to Y$ between them, we have*

$$\frac{1}{C} \le \frac{\lambda(U)}{\mathrm{vb}(U)} \le C(1 + \log n)^{5/2}.$$

Moreover, for any vectors $u_1, \dots, u_N$ and convex body K in $\mathbb{R}^n$ we can define a norm X on $\mathbb{R}^n$ so that for the space Y with unit ball K and the identity map $I : X \to Y$,

$$\frac{\lambda(I)}{C(1 + \log n)^{5/2}} \le \mathrm{hd}((u_i)_{i=1}^N, K) \le \mathrm{vb}(I) \le C\lambda(I).$$

Finally, $\lambda(U)$ is computable in polynomial time given appropriate access to X and Y. [1]

1.4 Organization

In section 2 we present basic definitions and preliminary material. In section 3, we present our proof of Theorem 2. In subsection 3.1, we present our partial progress on the restricted invertibility conjecture for convex bodies. In section 4, we present the proof of tightness for our factorization approach to vector balancing. In section 5, we give a polynomial time algorithm to compute the factorization constant up to a constant factor. In section 6.1, we show that the volume lower bound is invariant under taking convex hulls.

2 Preliminaries

We use the notation $[n] = \{1, \dots, n\}$. For vectors $x, y \in \mathbb{R}^n$, we define $\langle x, y \rangle = \sum_{i=1}^n x_i y_i$ to be the standard inner product in $\mathbb{R}^n$. For a square matrix $T \in \mathbb{R}^{n \times n}$, we define $\mathrm{tr}(T) = \sum_{i=1}^n T_{ii}$, and a matrix $M \in \mathbb{R}^{n \times m}$, we define its transpose $M_{ij}^{\mathsf{T}} := M_{ji}$. For two sets $A, B \in \mathbb{R}^n$, we define their Minkowski sum $A + B = \{a + b : a \in A, b \in B\}$.

For a linear subspace $W \subseteq \mathbb{R}^n$, we denote the orthogonal projection onto W by π_W. For $S \subseteq [n]$, we write π_S to denote the projection onto the coordinate subspace $\mathrm{span}\{e_i : i \in S)\}$.

Convexity. A convex body $K \subseteq \mathbb{R}^n$ is a compact convex set with non-empty interior. K is symmetric if $K = -K$. A symmetric convex body induces a norm $\|x\|_K = \min\{s \ge 0 : x \in sK\}$. If K contains the origin is its interior, the polar of K is defined by $K^\circ = \{x \in \mathbb{R}^n : \langle x, y \rangle \le 1, \forall y \in K\}$. Furthermore, by convex duality, we have that relation $(K^\circ)^\circ = K$.

For a subset $S \subseteq [n]$, we denote the coordinate section of K on S by $K_S \triangleq \{x \in K : x_i = 0, \ \forall i \notin S\}$.

For a vector $x \in \mathbb{R}^n$, for $p \in [1, \infty)$, we let $\|x\|_p = (\sum_{i=1}^n |x_i|^p)^{1/p}$ denote the ℓ_p norm, and $\|x\|_\infty = \max_{i=1}^n |x_i|$ denote the ℓ_∞ norm. We use B_p^n, $p \in [1, \infty]$, to denote the unit ℓ_p ball in dimension n, $B_p^S := (B_p^n)_S$ and $B_p^W := (B_p^n) \cap W$ for corresponding coordinate and general sections, where $S \subseteq [n]$ is a subset and $W \subseteq \mathbb{R}^n$ is a linear subspace.

[1] See Theorem 33 for the necessary assumptions.

Probability and Measure. We denote the n-dimensional Lebesgue measure by $\mathrm{vol}_n(\cdot)$. Let $\kappa_n := \mathrm{vol}_n(B_2^n)$ denote the volume of the Euclidean ball, which can be estimated by $\kappa_n^{1/n} \approx \sqrt{\frac{2\pi e}{n}}$. For a matrix $A \in \mathbb{R}^{n\times k}$, for any measurable set $S \subseteq \mathbb{R}^k$, we have $\mathrm{vol}_k(AS) = \det(A^\mathsf{T} A)^{1/2} \mathrm{vol}_k(S)$.

We define γ_n to be the standard Gaussian measure on $\mathbb{R}^n$, that is $\gamma_n(A) = \frac{1}{\sqrt{2\pi}^n} \int_A e^{-\|x\|^2/2}$. We will often use the k-dimensional Gaussian measure restricted to k-dimensional linear subspace H of $\mathbb{R}^n$, for which we use the notation γ_H.

Positive Definite Matrices and Ellipsoids. A matrix $A \in \mathbb{R}^{n\times n}$ is symmetric if $A = A^\mathsf{T}$. A symmetric matrix $A \in \mathbb{R}^{n\times n}$ is positive semidefinite (PSD), written $A \succ 0$, if $x^\mathsf{T} Ax \geq 0$ for all $x \in \mathbb{R}^n$. Equivalently, it is PSD if A if it is symmetric and all its eigenvalues are non-negative. A is positive definite, $A \succ 0$, it is eigenvalues are all strictly positive. We write $A \succeq B$ to mean $A - B \succeq 0$ and similarly for $A \succ B$. Every positive semidefinite matrix A has a unique positive semidefinite square root, which we denote $A^{1/2}$.

For an $n \times n$ positive definite matrix Q, we define the ellipsoid $E(Q) = \{x \in \mathbb{R}^n : x^\mathsf{T} Qx \leq 1\} = A^{-1/2}B_2^n$. The polar ellipsoid is $E(Q)^\circ = E(Q^{-1})$ that $\mathrm{vol}_n(E(Q)) = \kappa_n \det(Q)^{-1/2}$. The length of the principal axes of Q, which are aligned with the eigen vectors of Q, have length $1/\sqrt{\lambda_n} \geq \cdots \geq 1/\sqrt{\lambda_1}$, where $\lambda_1 \geq \cdots \lambda_n > 0$ are the eigenvalues of Q.

Membership Oracles. To interact with a convex body $K \subseteq \mathbb{R}^n$, we will assume that it is given by a well-guaranteed membership oracle O_K, where $O_K(x) = 1$ if $x \in K$ and 0 otherwise. It comes with guarantees (a_0, r, R), $a_0 \in \mathbb{R}^n$ a center, $0 < r < R$, for which $a_0 + rB_2^n \subseteq K \subseteq a_0 + RB_2^n$. With access to such oracle, one can perform many standard tasks in convex optimization, such as approximately maximize a linear function over K, or compute the closest point in K to an input point y, compute the norm $\|x\|_K$ (when K is symmetric), using a polynomial number of queries to the oracle and arithmetic operations. See for example [GLS88] for a reference. All our algorithms will rely upon the real model of computation.

Inequalities for Convex Bodies. We will need the following inequalities to relate the volume of a symmetric convex body to that of its polar.

Theorem 5 (Blaschke-Santaló). *Let $K \subseteq \mathbb{R}^n$ be a symmetric convex body. Then $\mathrm{vol}_n(K) \cdot \mathrm{vol}_n(K^\circ) \leq \kappa_n^2$, where equality holds if and only if K is an origin centered ellipsoid. Here $\kappa_n = \mathrm{vol}_n(B_2^n)$.*

Restricted Invertibility. We will need a refinement of the restricted invertibility theorem of Bourgain and Tzafriri [BT87] due to Spielman and Srivastava.

Theorem 6 ([SS10]). *Let $Q \in \mathbb{R}^{n\times n}$ be positive definite quadratic form and $\varepsilon \in (0, 1)$. Let $\lambda_1 := \lambda_1(Q) > 0$ denote the maximum eigenvalue of Q. For $k = \lfloor \varepsilon^2 \operatorname{tr}(Q)/\lambda_1 \rfloor$, there exists $S \subseteq [n]$, $|S| = k$, such that $\lambda_{\min}(Q_{S,S}) > \frac{(1-\varepsilon)^2 \operatorname{tr}(Q)}{n}$, where $\lambda_{\min}(Q_{S,S})$ is the minimum eigenvalue of $Q_{S,S}$.*

We will also need a couple of simple determinantal analogues of the restricted invertibility principle.

Lemma 7. *Let Q be an $n \times n$ real positive semi-definite matrix with eigenvalues $\lambda_1 \geq \ldots \geq \lambda_n$. For any integer k, $1 \leq k \leq n$, there exists a set $S \subseteq [n]$ of size k such that*

$$\prod_{i=1}^{k} \lambda_i \leq \binom{n}{k} \det(Q_{S,S}).$$

Proof. To prove the lemma, we will rely on the classical identity for applying the elementary symmetric polynomials to the eigenvalues of Q:

$$\sum_{S\in[n],|S|=k}\prod_{i\in S}\lambda_i \triangleq p_k(\lambda) = \sum_{S\subset[n]:|S|=k}\det(Q_{S,S}).$$

To verify this equation, consider the coefficient of t^{n-k} in the polynomial $\det(Q+tI)$. Calculating the coefficient using the Leibniz formula for the determinant gives the right hand side; calculating it using $\det(Q+tI) = (\lambda_1+t)\dots(\lambda_n+t)$ gives the left hand side. Since the eigenvalues are all non-negative, we get that

$$\prod_{i=1}^k \lambda_i \le p_k(\lambda) = \sum_{S\subseteq[n]:|S|=k}\det(Q_{S,S}) \le \binom{n}{k}\max_{S\subseteq[n]:|S|=k}\det(Q_{S,S}),$$

as needed. □

2.1 Proof of Volume Lower Bound

Lemma 1 (Volume Lower Bound)**.** *Let $U = (u_1,\dots,u_N) \in \mathbb{R}^{n\times N}$ and $K \subseteq \mathbb{R}^n$ be a symmetric convex body. For $S\subseteq[N]$, let U_S denote the columns of U in S. For $k\in[n]$, define*

$$\mathrm{volLB}^{\mathrm{h}}_k((u_i)_{i=1}^N, K) \triangleq \mathrm{volLB}^{\mathrm{h}}_k(U,K) \triangleq \max_{S\subseteq[N],|S|=k}\mathrm{vol}_k(\{x\in\mathbb{R}^k : U_S x\in K\})^{-1/k}. \tag{3}$$

Then, we have that

$$\mathrm{volLB}^{\mathrm{h}}((u_i)_{i=1}^N, K) \triangleq \mathrm{volLB}^{\mathrm{h}}(U,K) \triangleq \max_{k\in[n]}\mathrm{volLB}^{\mathrm{h}}_k(U,K) \le \mathrm{hd}(U,K). \tag{4}$$

Proof. For $S\subseteq[N]$, $|S| = k\in[n]$, let $C = \left\{x\in\mathbb{R}^k : U_S x\in K\right\}$. It is direct to verify $\mathrm{hd}((e_i)_{i=1}^k, C) = \mathrm{hd}(U_S, K) \le \mathrm{hd}(U,K)$, where $(e_i)_{i=1}^k$ is the standard basis of $\mathbb{R}^k$. Thus, it suffices to show that $\mathrm{hd}((e_i)_{i=1}^k, C) \ge \mathrm{vol}_k(C)^{-1/k}$. For $x\in\mathbb{R}^k$, $A\subseteq\mathbb{R}^k$ finite, let $d(x,A) := \min_{a\in A}\|a-x\|_C$ denote the minimum distance between x and A under the
(semi-)norm induced by C. From here, we apply the standard reduction from linear discrepancy to hereditary discrepancy [LSV86], to get

$$\begin{aligned}
\max_{x\in[0,1]^k} d(x,\{0,1\}^k) &\le \max_{x\in[0,1]^n} d(x,\{0,\tfrac12,1\}^k) + \max_{x'\in\{0,\frac12,1\}^k} d(x',\{0,1\}^n)\\
&\le \frac12\max_{x\in[0,1]^k} d(x,\{0,1\}^k) + \frac12\mathrm{hd}((e_i)_{i=1}^k, C)\\
&\Rightarrow\\
\max_{x\in[0,1]^k} d(x,\{0,1\}^k) &\le \mathrm{hd}((e_i)_{i=1}^k, C).
\end{aligned}$$

Let $r = \mathrm{hd}((e_i)_{i=1}^k, C)$, we in particular have that $[0,1]^k \subseteq \{0,1\}^k + rC$. Thus

$$\begin{aligned}
\mathrm{vol}_k(rC) &\ge \mathrm{vol}_k(\cup_{x\in\{0,1\}^k} rC\cap(-x+[0,1)^k)) = \mathrm{vol}_k(\cup_{x\in\{0,1\}^k}(rC+x)\cap[0,1)^k)\\
&= \mathrm{vol}_k((\{0,1\}^k+rC)\cap[0,1)^k) \ge \mathrm{vol}_k([0,1)^k) = 1.
\end{aligned}$$

In particular, $r \ge \mathrm{vol}_k(C)^{-1/k}$ as needed. □

3 Tightness of the Volume Lower Bound

In this section, we will show that the volume lower bound (3) is tight within a logarithmic factor.

Theorem 2 (Tightness of the Volume Lower Bound). *For any $U \in \mathbb{R}^{n\times N}$ and symmetric convex body K in $\mathbb{R}^n$, we have that*

$$\mathrm{volLB}^{\mathrm{h}}(U,K) \le \mathrm{hd}(U,K) \le O(\log n)\,\mathrm{volLB}^{\mathrm{h}}(U,K)\ , \tag{6}$$

Furthermore, there exists a randomized polynomial time algorithm that computes a coloring of U with K-discrepancy $O(\log n\,\mathrm{volLB}^{\mathrm{h}}(U,K))$, given a membership oracle for K.

The main technical result of this section is that the volume lower bound, restricted to subsets of size at least $\Omega(n)$, is in fact an upper bound on the discrepancy of so-called partial colorings. This allows us to easily recover Theorem 2 using $O(\log n)$ partial coloring phases in the standard way. We state our technical result below, restricted to the case where the vectors are aligned with the standard basis. We note that since the output norm is general, this is essentially without loss of generality. For a vector $x \in [-1,1]^n$, we use the notation $\mathrm{bounds}(x) = \{i \in [n] : x_i \in \{-1,1\}\}$, to denote the coordinates in x set to ± 1.

Lemma 8 (Partial Colorings via Volume). *There exists a universal constants $C \ge 1, \varepsilon_0 \in (0,1), \delta \in (0,1)$, such that for any $y \in (-1,1)^n$ and symmetric convex body $K \subseteq \mathbb{R}^n$ satisfying $\forall S \subseteq [n]$, $|S| = \lceil \delta n \rceil$, $\mathrm{vol}_{|S|}(K_S) \ge 1$, there exists a polynomial time algorithm which with high probability finds $x \in [-1,1]^n$ with $|\mathrm{bounds}(x)| \ge \lceil \varepsilon_0 n \rceil$ and $x - y \in CK$.*

We now give the straightforward reduction from Theorem 2 to Lemma 8.

Proof of Theorem 2. By Lemma 1, we may restrict attention to the upper bound. In particular, given $u_1,\dots,u_N \in \mathbb{R}^n$ and a convex body K in $\mathbb{R}^n$, it suffices to show that $\mathrm{disc}((u_i)_{i=1}^N) \le C(1+\log n)\,\mathrm{volLB}^{\mathrm{h}}((u_i)_{i=1}^N, K)$.

To begin, we compute a basic solution to the linear program $\sum_{i=1}^N x_i u_i = 0$, $x \in [-1,1]^N$. After relabeling, we may assume the variables not hitting their $\{-1,1\}$ bounds are $x_1,\dots,x_l$, noting that if there are no such variables we already have a 0 discrepancy coloring. Since x is basic, we know that the vectors $u_1,\dots,u_l$ must be linearly independent. Therefore, we may apply a invertible linear transformation $T : \mathbb{R}^n \to \mathbb{R}^n$ sending $u_1,\dots,u_l$ to $e_1,\dots,e_l$. In particular, letting $K' := TK$, we have that

$$\begin{aligned}\mathrm{disc}((u_i)_{i=1}^N, K) &\le \min_{z\in\{-1,1\}^l} \Big\|\sum_{i=1}^l z_i e_i + \sum_{i=l+1}^N x_i T u_i\Big\|_{K'} \\ &= \min_{z\in\{-1,1\}^l} \Big\|\sum_{i=1}^l (z_i - x_i) e_i\Big\|_{K'},\end{aligned}$$

Furthermore, a direct computation shows that

$$\mathrm{volLB}^{\mathrm{h}}((u_i)_{i=1}^l, K) = \mathrm{volLB}^{\mathrm{h}}((e_i)_{i=1}^l, K') = \max_{S\subseteq[l]} \mathrm{vol}_{|S|}(K'_S)^{-1/|S|}\ .$$

Let us assume that we have computed $M > 0$ satisfying

$$M/2 \le \mathrm{volLB}^{\mathrm{h}}((e_i)_{i=1}^l, K') \le M.$$

From here, it suffices to compute $z \in \{-1,1\}^l$ such that $\sum_{i=1}^l (z_i - x_i)e_i \in O(M \log n)K'$. Note that by assumption on M, $\mathrm{vol}(MK'_S) \geq 1, \forall S \subseteq [l]$. Therefore, repeatedly applying Lemma 8 on MK', letting $x^0 := x$ we compute a sequence $x^1, \dots, x^T \in [-1,1]^l$, $T = \lceil \log l/\varepsilon_0 \rceil = O(\log n)$, such that $\forall t \in [T]$, we have that

1. $\sum_{i=1}^l (x_i^t - x_i^{t-1})e_i \in O(M)K'$.
2. $|\{i \in [l] : x_i^t \in (-1,1)\}| \leq (1-\varepsilon_0)|\{i \in [l] : x_i^{t-1} \in (-1,1)\}|$.

By our choice of T, it is direct to check that $x^T \in \{-1,1\}^l$ and by the triangle inequality that $\sum_{i=1}^l (x_i^T - x_i^0)e_i \in TMK' = O(M \log n)K'$. Thus, setting $z = x^T$ satisfies the requirements.

We now discuss the computation of M. We first note that

$$M_1 := \max_{i \in [l]} \mathrm{vol}_1(K'_i)^{-1} = \max_{i \in [l]} \|e_i\|_{K'} = \max_{i \in [l]} \|u_i\|_K \ ,$$

and thus the restricted maximum can be efficiently computed. Note that by construction

$$\mathrm{conv}\{\pm e_1, \dots, \pm e_l\}/M_1 \subseteq K'_{[l]}.$$

Thus, for any $S \subseteq [l]$,$|S| = k$, we see that

$$\mathrm{vol}_k(lM_1K'_S) \geq \mathrm{vol}_k(l \cdot \mathrm{conv}(\pm e_i : i \in S)) = \frac{(2l)^k}{k!} \geq 1 \ .$$

In particular, we get that $M_1 \leq \mathrm{volLB^h}((e_i)_{i=1}^l, K') \leq lM_1$. Hence, as input to the stage above we may successively try the values $M_1 2^k$, $k \in \{0, \dots, \log_2 l\}$, stopping the first time we find a valid coloring. $\square$

Lemma 8 should be viewed as a volumetric analogue of a theorem of Rothvoss [Rot14], who both extended and made algorithmic vector balancing results of Giannopoulos [Gia97a]. We state a slight variant of [Rot14, Lemma 9] below.

Theorem 9 (Partial Colorings via Gaussian Measure)**.** *Let $0 < \varepsilon \leq 1/60000$ and $\delta = \frac{3}{2}\varepsilon \log_2 \frac{1}{\varepsilon}$. Let $K \subseteq \mathbb{R}^n$ be a symmetric convex body given by a membership oracle, and assume that for some subspace $H \subseteq \mathbb{R}^n$ of dimension at least $(1-\delta)n$, we have that $\gamma_H(K) \geq e^{-\delta n}$. Then for any $y \in (-1,1)^n$, there exists a polynomial time algorithm which with high probability finds $x \in [-1,1]^n$ satisfying $|\mathrm{bounds}(y)| \geq \varepsilon n/2$ and $x - y \in K$.*

We recall that Rothvoss' algorithm, for the special case $y = 0$ (the general case is similar), works by computing the Euclidean projection of a Gaussian random vector onto $K \cap [-1,1]^n$. The above statement deviates from the corresponding Lemma in [Rot14] in that it does not assume that $K \subseteq H$ or that H is known to the algorithm. It is not hard to verify however that this condition is not needed in analysis, so we defer discussion of the proof of this statement to the full version. The flexibility gained by not needing to know the subspace in advance will be very useful in the sequel. We note that one can also adapt the analysis of the algorithm of Singh and Eldan [ES14], which maximizes over $K \cap [-1,1]^n$ using the Gaussian as the objective vector instead of projecting it, to work in the above setting.

Our proof of Lemma 8 will in fact be a direct reduction to Rothvoss' theorem. The core of our reduction is the following geometric theorem, which shows that if all the coordinate sections of K of proportional dimension have large volume, then there exists a subspace H of proportional dimension on which K has large Gaussian measure.

Theorem 10 (Gaussian Measure via Volume)**.** *There exists a decreasing function $\eta : (0,1) \to \mathbb{R}_+$, $\eta(\delta) = e^{O(1/\delta)}$, such that the following holds. For any $n \in \mathbb{N}$, $K \subseteq \mathbb{R}^n$ symmetric convex body, $2 \le k \le n-1$, $\alpha = k/n$, such that $\forall S \subseteq [n]$, $|S| = \alpha n$, $\mathrm{vol}_{\delta n}(K_S) \ge 1$, there exists a linear subspace H of dimension $(1-\delta)n$ for which $\gamma_H(\eta(\delta)K) \ge e^{-\delta n}$.*

We note that the above theorem is primarily interesting in the case where δ is a fixed constant, as $\eta(\delta) = e^{O(1/\delta)}$ blows up quite quickly as $\delta \to 0$. Lemma 8 now follows directly combining the above with Theorem 9, as shown below.

Proof of Lemma 8. Let $\varepsilon = 1/60000$ and $\delta = (3/2)\varepsilon \log_2(1/\varepsilon)$. By Theorem 10, $\eta(\delta)K$ satisfies the conditions for applying Theorem 9 on any $x \in (-1,1)^n$ with 'parameters ε and δ as in the last sentence. This yields Lemma 8 with $\varepsilon_0 = \varepsilon/2$, $\delta = \delta$ and $C = O(\eta(\delta)) = O(1)$, as needed. $\square$

While there are examples where the volume lower bound is a $O(\log n)$ factor off from hereditary discrepancy (e.g. 3 permutations), we conjecture that the volume lower bound actually characterizes a hereditary version of partial coloring discrepancy. Namely, if the volume lower bound is D, we conjecture that there exists a subset of vectors for which every (fractional) partial coloring has discrepancy $\Omega(D)$. Recall that Lemma 8 gives the other direction, i.e. that there always exist partial colorings of discrepancy $O(D)$. If this conjecture were true, then Rothvoss' algorithm, which we use as a blackbox, would in a weak sense be optimal for finding partial colorings. We discuss this conjecture in more detail in the next subsection.

Comparing to prior works, Theorem 10 provides a useful and different route for proving that a body (or at least a large section of it) has exponentially small Gaussian measure. In the context of discrepancy, to the authors' knowledge, only two main techniques were used to prove such bounds, neither of which is directly applicable in the above setting. The first technique consists of combining chaining techniques and moment bounds, which can generally only measure when the body has Gaussian measure close to 1/2. This approach loses the leverage we have in only needing exponentially small bounds and thus often incurs additional logarithmic factors. The second strategy is based on the positive correlation properties of Gaussian measure, first proved for the intersection of symmetric slabs (i.e. Sidak's lemma), and with the recent resolution of the Gaussian correlation conjecture [Roy14], for the intersection of arbitrary symmetric convex bodies. More precisely, one tries to show that K contains (or equals) the intersection of "simpler" symmetric convex bodies $K_1, \dots, K_T$ (most often slabs) to deduce that $\gamma_n(K) \ge \prod_{i=1}^T \gamma_n(K_i)$, from which one can usefully get exponentially small bounds.

In contrast, our proof of the $e^{-\delta n}$ lower bounds on Gaussian measure in the above theorem proceeds via a direct covering argument. Namely, if one can show that $e^{\delta n - 1}$ translates of K cover the Euclidean ball of radius $\sqrt{n}$, which has Gaussian measure $\ge 1/2$, then one can directly deduce that

$$1 \le 2\gamma_n(\sqrt{n}B_2^n) \le 2N(\sqrt{n}B_2^n, K) \max_{t \in \mathbb{R}^n} \gamma_n(K+t) \le e^{\delta n}\gamma_n(K),$$

where we have used that $\max_{t \in \mathbb{R}^n} \gamma_n(K+t) = \gamma_n(K)$ for a centrally symmetric convex body, which follows by the symmetry and logconcavity of Gaussian measure. We will adopt this strategy on a section of K, which is chosen to align with the longest axes of a so-called regular M-ellipsoid for K. The volumetric condition in Lemma 8 will in fact be used to guarantee that these axes have length $\Omega(\sqrt{n})$, which makes the above strategy plausible. We recall that an M-ellipsoid E of K is an ellipsoid which approximates K well from the perspective of covering, i.e. $2^{O(n)}$ shifts of K suffice to cover E and vice versa. The existence of such ellipsoids was first proven by Milman [Mil86]. We give precise definitions below.

For any two sets $A, B \subseteq \mathbb{R}^n$, let

$$N(A,B) = \min\left\{|\Lambda| : \Lambda \subset \mathbb{R}^n, A \subseteq \Lambda + B\right\},$$

denote the minimum number of shifts of B needed to cover A. The following theorem of Pisier [Pis89], gives the existence of M-ellipsoids whose covering estimates have polynomial decay. The decay estimate will be used to make the Gaussian measure of the large section of K we find as close to 1 as we like after a sufficient scaling.

Theorem 11 (Regular M-ellipsoid)**.** *There exists an absolute constant $c_0 > 0$, such that any $0 < \alpha < 2$, letting $\sigma(\alpha) = c_0(2-\alpha)^{-1/2}$, $n \in \mathbb{N}$, and symmetric convex body $K \subseteq \mathbb{R}^n$, there exists an ellipsoid $E \subseteq \mathbb{R}^n$, $\mathrm{vol}_n(E) = \mathrm{vol}_n(K)$, such that for all $t \geq 1$*

$$\max\left\{N(K,tE), N(E,tK), N(K^\circ,tE^\circ), N(E^\circ,tK^\circ)\right\} \leq e^{\sigma(\alpha)n/t^\alpha} .$$

Such an ellipsoid will be referred to as an α-regular M-ellipsoid.

To show that the axes of the M-ellipsoid of K are long, we will need to relate the axis lengths to the volumes of coordinate projections of the polar ellipsoid. For this, purpose we will require the following formula for coordinate projection volumes.

Lemma 12. *Let $E := E(Q) \subseteq \mathbb{R}^n$ be an origin center ellipsoid. Then, for any $S \subseteq [n]$, $|S| = k$, we have that*

$$\mathrm{vol}_k(\pi_S(E^\circ)) = \kappa_k \det(Q_{S,S})^{1/2} .$$

Proof. Recall that $E^\circ = E(Q^{-1}) = Q^{1/2}B_2^n$, where $Q^{1/2}$ is the positive definite square root of Q. To begin, we recall that the support function of E° can be computed by

$$h_{E^\circ}(w) \triangleq \max_{y\in E^\circ}\langle w, y\rangle = \max_{z\in B_2^n}\langle w, Q^{1/2}z\rangle = \|Q^{1/2}w\| = \sqrt{w^\mathsf{T} Q w}.$$

Let $W_S = \mathrm{span}\{e_i : i \in S\}$. Note that by construction, $\pi_S(E^\circ) \subseteq W_S$ and $h_{\pi_S(E^\circ)}(w) = h_E(w)$, $\forall w \in W_S$. Furthermore, by duality, among convex bodies these conditions uniquely define $\pi_S(E^\circ)$.

Let $s_1 < s_2 < \cdots < s_k$ be the elements of S and let $P_S = (e_{s_1}, \ldots, e_{s_k})$, noting that $P_S P_S^\mathsf{T} = \pi_S$. Let $T = P_S(Q_{S,S})^{1/2} \in \mathbb{R}^{n\times k}$. We now show that $TB_2^k = \pi_S(E^\circ)$ using the aforementioned conditions. Clearly $TB_2^k \subseteq W_S$ since $\mathrm{span}(P_S) = W_S$. For $w \in W_S$, letting $w_S = (w_{s_1}, \ldots, w_{s_k})^\mathsf{T}$ denote the restriction to the coordinates in S, we have that

$$\begin{aligned} h_{TB_2^k}(w) &= \max_{z\in B_2^k}\langle T^\mathsf{T} w, z\rangle = \max_{z\in B_2^k}\langle (Q_{S,S})^{1/2}w_S, z\rangle \\ &= \sqrt{w_S^\mathsf{T} Q_{S,S} w_S} = \sqrt{w^\mathsf{T} Q w}, \end{aligned}$$

where the last equality follows since $w_i = 0$ for $i \notin S$. Thus $\pi_S(E^\circ) = TB_2^k$ as claimed. The volume can now be computed as follows:

$$\begin{aligned} \mathrm{vol}_k(TB_2^k) &= \kappa_k \det(T^\mathsf{T} T)^{1/2} = \kappa_k \det((Q_{S,S})^{1/2}(P_S^\mathsf{T} P_S)(Q_{S,S})^{1/2})^{1/2} \\ &= \kappa_k \det(Q_{S,S})^{1/2} , \end{aligned}$$

as needed. □

We now have all the ingredients needed to prove our main geometric estimate.

Proof of Theorem 10. Let $E := E(Q)$ denote a 1-regular M-ellipsoid for K and let $\sigma := \sigma(1)$. Let $l_1 \geq \cdots \geq l_n > 0$ denote the length of the principal axes of E, where we recall that $l_n^{-2} \geq \cdots \geq l_1^{-2} > 0$ are then the eigenvalues of Q.

By the Blaschke-Santaló inequality, for all $S \subseteq [n]$, $|S| = \delta n$, we have that

$$\mathrm{vol}_{\delta n}(K_S)\,\mathrm{vol}_{\delta n}((K_S)^\circ) = \mathrm{vol}_{\delta n}(K_S)\,\mathrm{vol}_{\delta n}(\pi_S(K^\circ)) \leq \kappa_{\delta n}^2 \ .$$

Since we assume $\mathrm{vol}_{\delta n}(K_S) \geq 1$, the above implies that $\mathrm{vol}_{\delta n}(\pi_S(K^\circ)) \leq \kappa_{\delta n}^2$. The coordinate projections of E° thus have volume at most

$$\mathrm{vol}_{\delta n}(\pi_S(E^\circ)) \leq N(E^\circ, K^\circ)\,\mathrm{vol}_{\delta n}(\pi_S(K^\circ)) \leq e^{\sigma n}\kappa_{\delta n}^2 \ .$$

Combining Lemma 7 and 12, we have that

$$\begin{aligned}
\prod_{i=(1-\delta)n+1}^{n} l_i^{-1} &\leq \binom{n}{\delta n}^{1/2} \max_{S \subseteq [n], |S| = \delta n} \det(Q_{S,S})^{1/2} \\
&= \binom{n}{\delta n}^{1/2} \max_{S \subseteq [n], |S| = \delta n} \mathrm{vol}_{\delta n}(\pi_S(E^\circ))\kappa_{\delta n}^{-1} \\
&\leq \binom{n}{\delta n}^{1/2} e^{\sigma n}\kappa_{\delta n} \ .
\end{aligned}$$

From here, we conclude that

$$l_{(1-\delta)n} \geq \prod_{i=(1-\delta)n+1}^{n} l_i^{\frac{1}{\delta n}} \geq \binom{n}{\delta n}^{-\frac{1}{2\delta n}} e^{-\frac{\sigma}{\delta}} \kappa_{\delta n}^{-\frac{1}{\delta n}} \geq \frac{1}{e^{2\frac{\sigma}{\delta}}} \cdot \sqrt{\frac{\delta n}{2\pi e}} := c(\delta)^{-1}\sqrt{n} \ . \tag{11}$$

Letting H be the span of the first $(1-\delta)n$ principal axes of E, we thus conclude that

$$\sqrt{n}(B_2^n \cap H) \subseteq c(\delta)(E \cap H).$$

Using the 1-regularity of E, letting $t = 2\sigma/\delta$, we derive the following covering estimate

$$\begin{aligned}
N(\sqrt{n}(B_2^n \cap H), 2tc(\delta)(K \cap H)) &\leq N(c(\delta)(E \cap H), 2tc(\delta)(K \cap H)) \\
&\leq N(E, tK) \leq e^{\sigma n/t} = e^{\delta n/2} \ .
\end{aligned}$$

Since $\gamma_H(\sqrt{n}B_2^n \cap H) \geq 1/2$, setting $\eta := \eta(\delta) = \frac{2c(\delta)\sigma}{\delta} = e^{O(1/\delta)}$, we get that $\gamma_H(\eta K \cap H) \geq \frac{1}{2}e^{-\delta n/2} \geq e^{-\delta n}$, as needed. Lastly, η as defined above is easily checked to be decreasing in δ. $\square$

3.1 The Discrepancy of Partial Colorings

In this section we discuss a geometric conjecture which would imply that a tight relationship between the discrepancy of partial colorings and the volume lower bound, and thus a weak form of optimality for Rothvoss' partial coloring algorithm. For this purpose, we formally define the partial coloring discrepancy as well as its hereditary version. Given $(u_i)_{i=1}^N \in \mathbb{R}^n$, symmetric convex body $K \subseteq \mathbb{R}^n$ and $\alpha \in (0,1]$, we define

$$\mathrm{disc}_\alpha((u_i)_{i=1}^N, K) \triangleq \min_{\substack{x \in [-1,1]^N \\ |\,\mathrm{bounds}(x)| \geq \alpha N}} \Big\| \sum_{i=1}^N x_i u_i \Big\|_K. \tag{12}$$

and

$$\mathrm{hd}_\alpha((u_i)_{i=1}^N, K) \triangleq \max_{S\subseteq[N]} \mathrm{disc}_\alpha((u_i)_{i\in S}, K). \tag{13}$$

We recall that (repeated applications) of Lemma 8 implies the upper bound

$$\mathrm{hd}_{1/2}((u_i)_{i=1}^N, K) \leq O(1)\, \mathrm{volLB}^{\mathrm{h}}((u_i)_{i=1}^N, K).$$

Here we conjecture that the reverse inequality should also hold.

Conjecture 13. *There exists a universal constant $c \geq 1$, such that for any $n \in \mathbb{N}$, $u_1, \dots, u_n \in \mathbb{R}^n$ linearly independent and symmetric convex body $K \subseteq \mathbb{R}^n$:*

$$\mathrm{volLB}^{\mathrm{h}}((u_i)_{i=1}^n, K) \leq c\, \mathrm{hd}_{1/2}((u_i)_{i=1}^n, K) \tag{14}$$

Note that we restrict above to linear independent subsets of vectors, but as is well-known (e.g. see proof of Theorem 2), this is without loss of generality. As a pathway to prove the conjecture, we suggest the following natural geometric analog of the so-called spectral lower bound for discrepancy into ℓ_2.

Lemma 14. *Let $U = (u_1, \dots, u_n) \in \mathbb{R}^{n\times n}$ be linearly independent, $K \subseteq \mathbb{R}^n$ be a symmetric convex body, and $\alpha \in [0,1]$. For any subset $S \subseteq [n]$, $|S| = k$, letting $W_S := \mathrm{span}\{u_i : i \in S\}$, define*

$$\mathrm{specLB}((u_i)_{i\in S}, K) := \max\left\{ r \geq 0 : rK \cap W_S \subseteq \sqrt{|S|}\, U_S B_2^k \right\}.$$

Then, we have that

$$\mathrm{disc}_\alpha((u_i)_{i\in S}, K) \geq \sqrt{\alpha}\, \mathrm{specLB}((u_i)_{i\in S}, K). \tag{15}$$

In particular, defining

$$\mathrm{specLB}^{\mathrm{h}}((u_i)_{i=1}^n, K) := \max_{S\subseteq[n]} \mathrm{specLB}((u_i)_{i\in S}, K)$$

we have that

$$\mathrm{hd}_\alpha((u_i)_{i=1}^n, K) \geq \sqrt{\alpha}\, \mathrm{specLB}^{\mathrm{h}}((u_i)_{i=1}^n, K). \tag{16}$$

Proof. We prove only (15) since then (16) follows trivially. For (15), by replacing K by $K \cap W_S$, we may wlog assume that $|S| = n$ and $W_S = \mathbb{R}^n$.

Let $x \in [-1,1]^n$, $|\,\mathrm{bounds}(x)| \geq \alpha n$, and $u_x = \sum_{i=1}^n x_i u_i$. Our goal is to show that $\beta := \|u_x\|_K \geq \sqrt{\alpha} r$, where $r := \mathrm{specLB}((u_i)_{i=1}^n, K)$.

Let $(u_i^*)_{i=1}^n$ denote the corresponding dual basis of $(u_i)_{i=1}^n$, i.e. satisfying $\langle u_i^*, u_j\rangle = 1$ if $i = j$ and 0 otherwise, which exists by linear independence. Now letting $v_x = \sum_{i=1}^n \frac{x_i}{\|x\|_2} u_i^*$, it is easy to check that

$$\langle v_x, u_x\rangle = \|x\|_2 \geq \sqrt{|\,\mathrm{bounds}(x)|} \geq \sqrt{\alpha n}.$$

Since $u_x \in \beta K$ and $rK \subseteq \sqrt{n} U B_2^n$, we have that

$$\sqrt{\alpha n} \leq \beta \max_{z\in K} \langle v_x, z\rangle \leq \frac{\beta}{r}\sqrt{n} \max_{z\in UB_2^n} \langle v_x, z\rangle = \frac{\beta}{r}\sqrt{n} \max_{z\in B_2^n} \langle \frac{x}{\|x\|_2}, z\rangle = \frac{\beta}{r}\sqrt{n}.$$

The desired inequality now follows by rearranging. $\square$

We note that as with $\mathrm{volLB^h}$, one may extend the $\mathrm{specLB^h}$ to an arbitrary sequence of vectors $(u_i)_{i=1}^N$, however one must take care to optimize only over subsets of linearly independent vectors, since otherwise the conclusion of Lemma 14 is false.

Given the above, it suffices to prove Conjecture 13 with $\mathrm{hd}_{1/2}$ replaced by $\mathrm{specLB^h}$. The resulting stronger conjecture has a very natural geometric interpretation which we expand on below. For $(u_i)_{i=1}^n \in \mathbb{R}^n$ linearly independent and $K \subseteq \mathbb{R}^n$ a symmetric convex body, letting T denote the linear map sending $(u_i)_{i=1}^n$ to $(e_i)_{i=1}^n$, it is direct to check that $\tau((u_i)_{i=1}^n, K) = \tau((e_i)_{i=1}^n, TK)$ for $\tau \in \left\{\mathrm{specLB^h}, \mathrm{volLB^h}\right\}$. Thus, for the purpose of the conjecture, it suffices to consider the setting where the vectors are the standard basis. In this setting, we see that

$$\mathrm{specLB^h}((e_i)_{i=1}^n, K) = \max_{S \subseteq [n]} \max \left\{ r \geq 0 : rK_S \subseteq \sqrt{|S|} B_2^S \right\}$$

and that

$$\mathrm{volLB^h}((e_i)_{i=1}^n, K) = \max_{S \subseteq [n]} \mathrm{vol}_{|S|}(K_S)^{-1/|S|} \ .$$

The goal is now to show that for every $S_0 \subseteq [n]$, there exists $S_1 \subseteq [n]$, such that

$$\mathrm{vol}_{|S_0|}(K_{S_0})^{-1/|S_0|} \leq c \max \left\{ r \geq 0 : rK_{S_1} \subseteq \sqrt{|S_1|} B_2^{|S_1|} \right\} \ . \tag{17}$$

Since this must hold for every symmetric convex body K, we may assume that $S_0 = [n]$ (and thus $S_1 \subseteq S_0$). Furthermore, by homogeneity, we may also assume that $\mathrm{vol}_n(K) = 1$. In this case (17), and hence Conjecture 13, directly reduces to the following geometric conjecture.

Conjecture 3 (Restricted Invertibility for Convex Bodies)**.** *There exists an absolute constant $c \geq 1$, such that for any $n \in \mathbb{N}$ and symmetric convex body $K \subseteq \mathbb{R}^n$ of volume at most 1, there exists $S \subseteq [n]$, $S \neq \emptyset$, such that $K_S \subseteq c\sqrt{|S|} B_2^S$.*

Two natural weakenings of Conjecture 3 are to ask whether (a) it holds for ellipsoids and (b) whether it holds for general bodies but with coordinate sections replaced by arbitrary sections. As our main evidence for the conjecture, we show that both statements are true. We note that (a) indeed implies Conjecture 13 when K is an ellipsoid. We state our results formally below.

Theorem 15. *There exists a universal constant $c \geq 1$, such that any $n \in \mathbb{N}$, the following holds:*

1. *For any origin centered ellipsoid $E \subseteq \mathbb{R}^n$ of volume at most 1, there exists $S \subseteq [n]$, $S \neq \emptyset$, such that $E_S \subseteq c\sqrt{|S|} B_2^S$.*
2. *For any symmetric convex body $K \subseteq \mathbb{R}^n$ of volume at most 1, there exists a linear subspace $W \subseteq \mathbb{R}^n$, such that $K \cap W \subseteq c\sqrt{W} B_2^W$.*

To prove the above theorem, we will need the following two lemmas.

Lemma 16. *Let $K \subseteq \mathbb{R}^n$ be a symmetric convex body of volume 1 satisfying for all $S \subseteq [n]$, $\mathrm{vol}_{|S|}(K_S) \geq 1$. Let $E \subseteq \mathbb{R}^n$ be a 1-regular M-ellipsoid of K. Then there exists $S \subseteq [n]$, $|S| \geq a_1 n$, such that $E_S \subseteq a_2\sqrt{|S|} B_2^S$, where $a_1, a_2 > 0$ are universal constants.*

Proof. Let $E = E(Q)$ be a 1-regular M-ellipsoid of K. Recall that $N(K, E) \leq N(E, K) \leq e^{\sigma n}$ where $\sigma := \sigma(1)$ and that $\mathrm{vol}_n(K) = \mathrm{vol}_n(E) = 1$. Let $\lambda_1 \geq \cdots \geq \lambda_n > 0$ denote the axes of the M-ellipsoid, where $1/\lambda_n^2 \geq \cdots \geq 1/\lambda_1^2 > 0$ are the eigenvalue of Q.

Recall by the proof of Theorem 10 (equation (11)), we have for that

$$\lambda_{\lceil (3/4)n \rceil} \geq e^{-O(4/3)} \sqrt{n} \geq c_1 \sqrt{n}, \tag{18}$$

where $c_1 > 0$ is an absolute constant.

We now show that assuming $\mathrm{vol}_n(E) = 1$, that the central axes also have length $O(\sqrt{n})$. Let $k = \lfloor \frac{1}{4} n \rfloor$. By Lemma 7 applied to Q^{-1}, there exists $S \subseteq [n]$, such that

$$\prod_{i=1}^{k} \lambda_i^2 \le \binom{n}{k} \det((Q^{-1})_{S,S}). \tag{19}$$

Let $R = [n] \setminus S$. Using the fact that $(Q^{-1}_{S,S})^{-1}$ is the Schur complement of Q with respect to $Q_{R,R}$ block, we have the identity

$$\det(Q) = \det(Q_{R,R}) \det((Q^{-1}_{S,S})^{-1}) \Leftrightarrow \frac{\det(Q)^{-1}}{\det((Q_{R,R}))^{-1}} = \det((Q^{-1})_{S,S}). \tag{20}$$

From here, since $\mathrm{vol}_n(E) = 1$, we have that

$$\begin{aligned} \frac{\det(Q)^{-1}}{\det((Q_{R,R}))^{-1}} &= \frac{\kappa_{n-k}^2 \,\mathrm{vol}_n(E)^2}{\kappa_n^2 \,\mathrm{vol}_{n-k}(E_R)^2} = \frac{\kappa_{n-k}^2}{\kappa_n^2} \frac{1}{\mathrm{vol}_{n-k}(E_R)^2} \\ &\le \frac{\kappa_{n-k}^2}{\kappa_n^2} \frac{e^{2\sigma n}}{\mathrm{vol}_{n-k}(2K_R)^2} \le \frac{\kappa_{n-k}^2}{\kappa_n^2} e^{4\sigma n}. \end{aligned} \tag{21}$$

Combining (19) and (21), using that $k = \lfloor n/4 \rfloor$, we get that

$$\lambda_k \le (\prod_{i=1}^{k} \lambda_i)^{1/k} \le \binom{n}{k}^{1/(2k)} (\frac{\kappa_{n-k}}{\kappa_n})^{1/k} e^{2\sigma n/k} \le c_2 \sqrt{n}, \tag{22}$$

where $c_2 \ge 1$ is an absolute constant. Given the above, we have that

$$c_1 \sqrt{n} \le \lambda_{\lceil (3/4)n \rceil} \le \lambda_{\lfloor (1/4)n \rfloor} \le c_2 \sqrt{n}\ . \tag{23}$$

Let W denote the span of axes of E associated with $\lambda_{\lfloor (1/4)n \rfloor}, \dots, \lambda_{\lceil (3/4)n \rceil}$ and let π_W denote the corresponding orthogonal projection. Let $Q' = Q^{1/2} \pi_W Q^{1/2} \preceq Q$, noting that Q' preserves all the eigenvalues associated with W while setting the others to zero. Applying Theorem 6 to Q' with $\varepsilon = 1/2$, we get a subset $S \subseteq [n]$, which by (23) and our choice of W has size at least

$$|S| = \frac{1}{4} \frac{\mathrm{tr}(Q')}{\lambda_{\max}(Q')} \ge c_3 n$$

where $c_3 > 0$ is an absolute constant. Furthermore,

$$\lambda_{\min}(Q_{S,S}) \ge \lambda_{\min}(Q'_{S,S}) \ge \frac{\mathrm{tr}(Q')}{n} \ge \frac{1}{2} \lambda^{-2}_{\lfloor (1/4)n \rfloor} \ge c_4/n\ ,$$

where $c_4 > 0$ is an absolute constant. Noting that the above implies that $E_S \subseteq \sqrt{n/c_4} B_2^S$ completes the proof, setting $a_1 = c_3$ and $a_2 = 1/\sqrt{c_4}$. □

The next lemma is essentially a consequence of Milman's quotient of subspace theorem, whose proof we defer to the full version.

Lemma 17. *Let $K \subseteq \mathbb{R}^n$ be a symmetric convex body such that $N(K, \sqrt{n} B_2^n) \le 2^{O(n)}$. Then there exists a linear subspace $W \subseteq \mathbb{R}^n$, $\dim(W) = \Theta(n)$, such that $K \cap W \subseteq O(\sqrt{|W|}) B_2^W$.*

Given the above two lemmas, we can now prove our main theorem.

Proof of Theorem 15. Firstly, for both statements, we note that we may in fact that additionally assume that $\mathrm{vol}_{|S|}(K_S) \ge 1$ for all $S \subset [n]$. This follows by noting that if there exists $S \subset [n]$ with $\mathrm{vol}_{|S|}(K_S) < 1$, we may simply apply induction on K_S, after scaling it up to have volume 1 (which only makes the task more difficult).

Proof of 1. When K is an ellipsoid, the statement follows immediately by applying Lemma 16 with $E = K$.

Proof of 2. First apply Lemma 16 to K, noting that the produced section K_S now satisfies the conditions of Lemma 17, from which we derive the result. □

4 The Factorization Approach

In this section we develop the approach to vector balancing via linear operators and prove that, roughly speaking, applying Banaszczyk's vector balancing theorem optimally gives nearly tight bounds on vector balancing constants. This is the basis of our polynomial time approximation algorithm for hereditary discrepancy in any norm.

Let us recall the definition of the vector balancing constant of an operator $U : X \to Y$ between two n-dimensional normed space X and Y, given in (10):

$$\mathrm{vb}(U) = \sup \left\{ \min_{\varepsilon_1,\ldots,\varepsilon_N \in \{-1,1\}} \left\| \sum_{i=1}^{N} \varepsilon_i U(x_i) \right\|_Y : N \in \mathbb{N}, x_1, \ldots, x_N \in B_X \right\}$$

Once again our main tool for giving lower bounds on $\mathrm{vb}(U)$ is the volume lower bound, which we reformulate in the operator setting.

Lemma 18. *Let $(X, \|\cdot\|_X)$ and $(Y, \|\cdot\|_Y)$ be two n-dimensional normed spaces, and let $U : X \to Y$ be an invertible linear operator. Define*

$$\mathrm{volLB}^{\mathrm{h}}_k(U) \triangleq \sup\Big\{\mathrm{vol}_k\Big(\Big\{a \in \mathbb{R}^k : \|UVa\|_Y \le 1\Big\}\Big)^{-1/k}\Big\},$$

where the supremum is over all operators $V : \ell_1^k \to X$ of operator norm 1 and rank k. Then, letting

$$\mathrm{volLB}^{\mathrm{h}}(U) \triangleq \max_{k \in [n]} \mathrm{volLB}^{\mathrm{h}}_k(U),$$

we have that

$$\mathrm{vb}(U) \ge \mathrm{volLB}^{\mathrm{h}}(U).$$

Lemma 18 is just a reformulation of Lemma 1 in the language of linear operators. Because

$$\max_{i=1}^{k} \|Ve_i\|_X = \|V\| = 1,$$

the points $u_1, \ldots, u_k$ defined by $u_i = Ve_i$ lie in B_X, where e_i is the i-th standard basis vector of X. Then the lemma follows directly from Lemma 1 with $C = B_X$ and $K = \{x \in X : Ux \in B_Y\}$.

Before we present the details of our approach, we give some relevant preliminaries on normed spaces and operators.

4.1 Basic Concepts in Normed Spaces

To every finite dimensional normed space $(X, \|\cdot\|)$ over $\mathbb{R}$ we associate its dual space $(X^*, \|\cdot\|_*)$ defined over linear maps from X to $\mathbb{R}$ (i.e. linear functionals) with the dual norm $\|x^*\|_* = \sup\{\langle x, x^*\rangle : \|x\|_X \le 1\}$. Here the notation $\langle x, x^*\rangle$ means "the functional x^* applied to x", i.e. $x^*(x)$. For any finite dimensional space X we have $X^{**} = X$. Any vector $y \in \mathbb{R}^n$ gives a linear functional over ℓ_2^n via the standard inner product, i.e. $\langle x, y\rangle = \sum_{i=1}^n x_i y_i$, and by the

Cauchy-Schwarz inequality this linear functional has norm equal to $\|y\|_2$. For this reason, we can identify $(\ell_2^n)^*$ with ℓ_2^n, and for $x, y \in \mathbb{R}^n$ identify $\langle x, y\rangle$ with the standard inner product.

As usual, given an operator $A : X \to Y$, we define its dual operator $A^* : Y^* \to X^*$ on every $y^* \in Y^*$ by $\langle x, A^*y^*\rangle = \langle Ax, y^*\rangle\ \forall x \in X$. We use the shorthand A^{-*} for $(A^*)^{-1} = (A^{-1})^*$.

We will use the tensor product notation $x^* \otimes y$ for the rank-1 linear operator from a normed space X to a normed space Y, defined by $(x^* \otimes y)(x) = \langle x, x^*\rangle y$, where $x^* \in X^*$ and $y \in Y$. Note that if X and Y are normed spaces over $\mathbb{R}^n$, the matrix of $x^* \otimes y$ with respect to the standard basis is yx^{T}. Any linear operator $A : X \to X$ on an n-dimensional normed space X can be written as as sum of rank-1 operators $A = \sum_{i=1}^n x_i^* \otimes y$ for $x_1^*, \dots, x_n^* \in X^*$ and $y_1, \dots, y_n \in X$. The trace of A is then defined by

$$\operatorname{tr}(A) \triangleq \sum_{j=1}^n \langle y_j, x_j^*\rangle.$$

This abstract definition agrees with the usual one, i.e. if the matrix of A with respect to a basis of X and the corresponding dual basis of X^* is M, then $\operatorname{tr}(A) = \operatorname{tr}(M)$. This also shows that $\operatorname{tr}(A)$ is uniquely defined, independent of how we write A as a sum of rank-1 operators. We will identify linear functionals on operators $A : X \to Y$, where X and Y are n-dimensional, with operators $B : Y \to X$ via $f_B(A) = \operatorname{tr}(BA)$.

A linear operator $A : X \to X^*$ on normed space X defines a bilinear form B on $X \times X$, given by $B(x, y) = \langle x, Ay\rangle$. We will say that A is positive definite if the corresponding bilinear form is symmetric and positive definite, i.e. if $\langle x, Ay\rangle = \langle Ax, y\rangle$ for all $x, y \in X$, and $\langle x, Ax\rangle > 0$ for all nonzero $x \in X$; A is positive semidefinite if instead we have $\langle x, Ax\rangle \geq 0$. In the case of $X = \ell_2^n$ this is equivalent to stating that the matrix M of A with respect to the standard basis is positive definite, i.e. is symmetric and all its eigenvalues are positive. We write $A \succ 0$ to denote that A is positive definite, and $A \succeq 0$ to denote that it is positive semidefinite.

For a positive definite operator $A : \ell_2^n \to \ell_2^n$, and a positive integer k, there exists a unique positive definite operator $B : \ell_2^n \to \ell_2^n$ such that $B^k = A$. We use the notation $A^{1/k}$ for B. We also use the shorthand notation $A^{\ell/k} \triangleq (A^\ell)^{1/k}$ for (positive or negative) integers ℓ. Equivalently, we can derive $A^{\ell/k}$ by raising every eigenvalue of A to the power ℓ/k in the spectral decomposition of A.

We also recall that the operator norm $\|A\|$ of a linear operator $A : X \to Y$ is defined by

$$\|A\| = \sup\{\|Ax\|_Y : \|x\|_X \leq 1\},$$

where $X = (\mathbb{R}^n, \|\cdot\|_X)$ and $Y = (\mathbb{R}^n, \|\cdot\|_Y)$ are normed spaces.

A related norm on operators is the nuclear norm. Here we only use the nuclear norm $\nu(A)$ of an operator $A : \ell_2^n \to \ell_2^n$, which equals the sum of its singular values. It is easy to see that

$$\nu(A) = \operatorname{tr}((AA^*)^{1/2}).$$

The nuclear norm is dual to the operator norm, and in particular we have the identity

$$\nu(A) = \sup\{\operatorname{tr}(AO) : O \text{ orthogonal transformation}\},$$

where the supremum is over orthogonal transformations on ℓ_2^n.

4.2 The Factorization Constant λ

In what follows we fix two n-dimensional normed spaces $(X, \|\cdot\|_X)$ and $(Y, \|\cdot\|_Y)$, and an invertible linear operator $U : X \to Y$. Since we work with general finite-dimensional normed space, restricting to spaces of equal dimension and to invertible operators is without loss of generality: given an operator $U : X \to Y$ which has a nontrivial kernel W, we can replace X by the quotient space X/W, and Y by its subspace given by the range of U, and U by the induced operator $\tilde{U} : X/W \to U(X)$ which sends $x + W$ to Ux. It is straightforward to check that this does not change any of the quantities we study.

In this section we reduce all upper bounds on vector balancing constants to the following deep theorem of Banaszczyk.

Theorem 19 ([Ban98])**.** *Let K be a convex body in $\mathbb{R}^n$ such that $\gamma_n(K) \geq \frac{1}{2}$. Then, $\mathrm{vb}(B_2^n, K) \leq 5$.*

Recall the definition of the vector balancing constant $\lambda(U)$ of $U : X \to Y$:

$$\lambda(U) \triangleq \inf\{\ell(S)\|T\| : T : X \to \ell_2^n,\ S : \ell_2^n \to Y, U = ST\}.$$

Note that a standard compactness argument shows that the infimum is in fact achieved.

We have the following theorem, which shows that $\lambda(U)$ is, up to constants, an upper bound on $\mathrm{vb}(U)$ for any operator U.

Theorem 20. *There exists a constant C such that for any linear operator $U : X \to Y$ between two n-dimensional normed spaces X, Y, we have*

$$\mathrm{vb}(U) \leq C\lambda(U).$$

Proof. As mentioned above, we can assume U to be invertible. Let $x_1, \dots, x_N \in B_X$ be arbitrary, and let $T : X \to \ell_2^n$, $S : \ell_2^n \to Y$ be such that $ST = U$, $\ell(S) \leq \lambda(U)$ and $\|T\| \leq 1$. For $i \in \{1, \dots, N\}$, define $u_i \triangleq Tx_i$; since $\|T\| \leq 1$ by assumption, we have $u_i \in B_2^n$ for all i. Let $K = \sqrt{2}\ell(S)S^{-1}(B_Y)$, and let, as usual, $\|\cdot\|_K$ be the norm with unit ball K. Observe that for any $x \in \mathbb{R}^n$, $\|x\|_K = \frac{1}{\sqrt{2}\ell(S)}\|Sx\|_Y$. By Chebyshev's inequality, for a standard Gaussian Z,

$$1 - \gamma_n(K) \leq \mathbb{E}\|Z\|_K^2 = \frac{1}{2\ell(S)^2}\mathbb{E}\|Z\|_Y^2 = \frac{1}{2}.$$

We can, therefore, apply Theorem 19, and have that there exist signs $\varepsilon_1, \dots, \varepsilon_N$ such that

$$\begin{aligned}\sum_{i=1}^N \varepsilon_i u_i \in 5K \iff \left\|\sum_{i=1}^N \varepsilon_i u_i\right\|_K \leq 5 &\iff \left\|\sum_{i=1}^N \varepsilon_i S u_i\right\|_Y \leq 5\sqrt{2}\lambda(U)\\ &\iff \left\|\sum_{i=1}^N \varepsilon_i U x_i\right\|_Y \leq 5\sqrt{2}\lambda(U)b,\end{aligned}$$

where we used that $Su_i = STx_i = Ux_i$. This completes the proof. □

Theorem 20 refines and generalizes the connection between the γ_2 norm and hereditary discrepancy from [MNT18]. The γ_2 norm of an operator $U : X \to Y$ is defined by

$$\gamma_2(U) = \inf\{\|S\|\|T\| : T : X \to \ell_2^n,\ S : \ell_2^n \to Y, U = ST\}.$$

In [MNT18], the authors studied the special case in which $X = \ell_1^n$ and $Y = \ell_\infty^m$. In that case, $\|S\|$ is the largest Euclidean norm of a column in the matrix of S, and $\|T\|$ is the largest Euclidean norm of a row of the matrix of T. It then follows from standard concentration of measure arguments that, for an absolute constant C,

$$\|S\| \leq \ell(S) \leq C\sqrt{1 + \log m} \cdot \|S\|, \tag{24}$$

for any operator $S : \ell_2^n \to \ell_\infty^m$. Therefore, for any $U : \ell_1^n \to \ell_\infty^m$, we have

$$\gamma_2(U) \le \lambda(U) \le C\sqrt{1 + \log m} \cdot \gamma_2(U).$$

This inequality and Theorem 20 recover the upper bound on hereditary discrepancy in terms of the γ_2 norm from [MNT18]. It also shows that λ provides at least as good an approximation to hereditary discrepancy as γ_2. However, (24) is often not tight, and Theorem 20 provides a tighter upper bound.

Another interesting special case is $X = \ell_1^n$ and $Y = \ell_2^m$. Since for any $S : \ell_2^n \to \ell_2^n$, $\ell(S) = \|S\|_{HS}$, where $\|A\|_{HS} = \mathrm{tr}(SS^*)^{1/2}$ is the Hilbert-Schmidt norm of A, we have

$$\lambda(U) = \inf\{\|S\|_{HS}\|T\| : T : \ell_1^n \to \ell_2^n,\ S : \ell_2^n \to \ell_2^n, U = ST\}.$$

This function was studied by two of the authors in [NT15], where they showed that it approximates hereditary discrepancy with respect to ℓ_2^n up to a factor of $O(\log n)$.

Our goal in the remainder of the section is to prove that the inequality in Theorem 20 holds in the reverse direction as well, as captured in the following theorem.

Theorem 21. *There exists a constant C such that the following holds. Let X and Y be two n-dimensional normed spaces and let $U : X \to Y$ be a linear operator between them. Then*

$$\lambda(U) \le CK(Y)(1 + \log n)^{3/2}\,\mathrm{volLB}^{\mathrm{h}}(U),$$

where $K(Y) = O(\log n)$ is the K-convexity constant of Y. [2]

Moreover, there exists an integer $k \le n$, and a rank k operator $V : \ell_1^k \to X$ of operator norm 1 so that for any standard basis vector e_i, Ve_i is an extreme point of B_X, and

$$\lambda(U) \le CK(Y)(1 + \log n)^{3/2}\,\mathrm{vol}_k\Big(\Big\{a \in \mathbb{R}^k : \|UVa\|_Y \le 1\Big\}\Big)^{-1/k}.$$

Theorems 20 and 21 together with Lemma 18 give a characterization of $\mathrm{vb}(U)$ in terms of $\lambda(U)$.

Corollary 22. *There exists a constant C such that for any two n-dimensional normed spaces X and Y, and any linear operator $U : X \to Y$ between them, we have*

$$\frac{1}{C} \le \frac{\lambda(U)}{\mathrm{vb}(U)} \le CK(Y)(1 + \log n)^{3/2},$$

where $K(Y) = O(\log n)$ is the K-convexity constant of Y.

The statement after "moreover" in Theorem 21 allows us to also show that $\lambda(U)$ approximately characterizes hereditary discrepancy as well.

Corollary 23. *Given a sequence of vectors $(u_i)_{i=1}^N$ in $\mathbb{R}^n$ and a convex body K in $\mathbb{R}^n$, define X to be the normed space with unit ball $B_X = \mathrm{conv}\{\pm u_1, \dots \pm u_N\}$, Y to be the normed space with unit ball K, and $I : X \to Y$ to be the identity map. Then, for an absolute constant C,*

$$\frac{\lambda(I)}{CK(Y)(1 + \log n)^{3/2}} \le \mathrm{hd}((u_i)_{i=1}^N, K) \le \mathrm{vb}(I) \le C\lambda(I)$$

[2] See Section 4.4 for a definition of $K(Y)$.

Proof. The inequality $\mathrm{hd}((u_i)_{i=1}^N, K) \leq \mathrm{vb}(I)$ is trivial from the definitions, and then the final inequality follows from Theorem 20. For the first inequality, observe that the points $Ve_1, \dots, Ve_k$ belong to $(u_i)_{i=1}^n$, so there is some subset $S \subseteq [N]$ of size k for which $(Ve_i)_{i=1}^k$ equals $(u_i)_{i \in S}$ (possibly after rearrangement), so that

$$\mathrm{vol}_k\Big(\Big\{a \in \mathbb{R}^k : \|UVa\|_Y \leq 1\Big\}\Big)^{-1/k} = \mathrm{volLB}((u_i)_{i \in S}, K).$$

Then, the first inequality follows from Lemma 1 and the statement after "moreover" in Theorem 21. □

Finally, in Section 5 we show that, under reasonable assumptions on how we are given access to the normed spaces X and Y, $\lambda(U)$ can be computed in polynomial time. Given Corollaries 22 and 23, this also implies a polynomial time approximation algorithm for vector balancing and hereditary discrepancy with arbitrary norms, and proves Theorem 4.

4.3 Convex And Dual Formulations

Here we give an equivalent formulation of $\lambda(U)$ as a convex minimization problem, i.e. as the infimum of a convex function over a convex domain. Then we use this convex formulation to derive another equivalent dual reformulation as a maximization problem.

The objective function $f(A)$ of our convex optimization formulation of $\lambda(U)$ is defined as follows: for any positive definite operator $A : X \to X^*$, we set

$$f(A) \triangleq \ell(UT^{-1}), \tag{25}$$

where $T : X \to \ell_2^n$ is an invertible linear operator such that $T^*T = A$.

A couple of clarifications are in order. First, we claim that such an operator T exists, by the positive definiteness of A. Indeed, we can choose a basis $e_1, \dots, e_n$ of X, and a corresponding dual basis $e_1^*, \dots, e_n^*$ of X^*, and define M to be the matrix of A with respect to these bases. Then M is a positive definite matrix, so it admits a Cholesky decomposition $M = L^\mathsf{T} L$. We can then define T to be the operator whose matrix with respect to $e_1, \dots, e_n$ and the standard basis of ℓ_2^n is L; the matrix of the dual operator T^* is L^T, so we have $T^*T = A$ as required.

A second concern is whether $f(A)$ is well-defined. To see that this is the case, observe that $A = T^*T = S^*S$ implies that there is an orthogonal transformation $O : \ell_2^n \to \ell_2^n$ for which $S = OT$. Then, $S^{-1} = T^{-1}O^{-1}$, and, for a standard Gaussian random variable Z,

$$\ell(US^{-1}) = (\mathbb{E}\|UT^{-1}O^{-1}Z\|_Y^2)^{1/2} = (\mathbb{E}\|UT^{-1}Z\|_Y^2)^{1/2} = \ell(UT^{-1}),$$

where we used the fact that $O^{-1}Z$ and Z are identically distributed because O^{-1} is an orthogonal transformation.

Having defined the objective function, we are now ready to specify our convex optimization problem for $\lambda(U)$.

Lemma 24. *For any two n-dimensional normed spaces X and Y, and any invertible linear operator $U : X \to Y$, $\lambda(U)$ equals*

$$\inf f(A) \tag{26}$$

$$s.t.$$

$$A : X \to X^*, \|A\| \leq 1 \tag{27}$$

$$A \succ 0. \tag{28}$$

The function f is the one defined in (25).

Moreover, the objective (26) and the constraints (27)–(28) are convex in A.

The dual formulation of $\lambda(U)$ is given in the following lemma.

Lemma 25. *Let X and Y be two n-dimensional normed spaces, such that the unit ball of X is $B_X = \operatorname{conv}\{\pm x_1, \dots, \pm x_m\}$. Then, for any linear operator $U : X \to Y$, $\lambda(U)$ equals*

$$\sup \operatorname{tr}((RU(\sum_{i=1}^{m} p_i x_i \otimes x_i)U^* R^*)^{1/3})^{3/2} \tag{29}$$

s.t.

$$R : Y \to \ell_2^n, \ell^*(R) \le 1 \tag{30}$$

$$\sum_{i=1}^{m} p_i = 1 \tag{31}$$

$$p_1, \dots, p_m \ge 0. \tag{32}$$

We prove Lemmas 24 and 25 in Section 4.5. Before doing so, we use them to prove Theorem 21 in the following section.

4.4 Proof of Theorem 21

We will use Lemma 25 to prove Theorem 21. In the proof, we need to relate the objective function (29) of the dual formulation to volumetric information about X and Y. We do so in two steps. In the first step, we consider a solution R, p of (29)–(32), and, using purely linearly algebraic techniques, we find a subset S of $\{1, \dots, m\}$ so that the Gram matrix of the vectors $(RUx_i)_{i \in S}$ has large determinant in relation to the value of (29). This allows us to define an operator V from ℓ_1^k (for $k = |S|$) to X so that the set $V^*U^*R^*(B_2^n))$ has large volume. In the second step of the proof we give a lower bound on the volume of the set $V^*U^*(B_Y)$ in terms of the volume of $V^*U^*R^*(B_2^n))$. Here we use classical connections between the ℓ^* norm and the ℓ norm (via K-convexity) and between the ℓ norm and covering numbers (via the dual Sudakov inequality). Since $V^*U^*(B_Y)$ is polar to the set $\{a : \|UVa\|_Y \le 1\}$ appearing in the volume lower bound, we can finish the proof by appealing to the Blaschke-Santaló inequality.

In the context of linear operators it is often convenient to use the notion of entropy numbers instead of covering numbers. The entropy number $e_k(A)$ of a linear operator $A : X \to Y$ is defined by

$$e_k(u) = \inf\{\varepsilon : N(u(B_X), \varepsilon B_Y) \le 2^{k-1}\}.$$

It is well known that covering numbers give both upper and lower estimates for the supremum of a Gaussian process. Here we use the dual Sudakov inequality, which in the language of entropy numbers has the following simple form: there exists a constant C such that for any linear operator $A : \ell_2^n \to X$, we have

$$\max_{k=1}^{n} \sqrt{k} e_k(u) \le C\ell(u). \tag{33}$$

This inequality is due to [PTJ85]. See [LT11, Section 3.3] for an easy proof.

Another important tool in the proof of Theorem21 is K-convexity, introduced by Maurey and Pisier [MP76]. The K-convexity constant $K(Y)$ of an n-dimensional normed space Y is the infimum over all constants K for which the inequality

$$\ell(A^*) \le K\ell^*(A) \tag{34}$$

holds for every operator $A : Y \to \ell_2^n$. (See [Pis89] or [TJ89] for an equivalent definition.) An important estimate of Pisier [Pis80] shows that that there exists an absolute constant C such that for any n-dimensional normed space Y,

$$K(Y) \leq C(1 + \log d(Y, \ell_2^n)) \leq C(1 + \log n). \tag{35}$$

Above $d(Y, \ell_2^n)$ is he Banach-Mazur distance between Y and ℓ_2^n, equal to the minimum of $\|T\|\|T^{-1}\|$ over linear operators $T : Y \to \ell_2^n$. Equivalently, it is equal to the smallest d for which there exists a linear operator T such that $B_2^n \subseteq T(B_Y) \subseteq dB_2^n$. For any n-dimensional normed space Y, $d(Y, \ell_2^n)$ is bounded by $\sqrt{n}$ by John's theorem, which implies the second inequality.

We also make use of a weighted version of Lemma 7.

Lemma 26. *Let $u_1, \dots, u_m \in \mathbb{R}^n$, and let $p_1, \dots, p_m \geq 0$, $\sum_{i=1}^m p_i = 1$. Let $\lambda_1 \geq \dots \geq \lambda_n$ be the eigenvalues of the matrix $\sum_{i=1}^m p_i u_i^\mathsf{T} u_i$, and let G be the Gram matrix of $u_1, \dots, u_m$, i.e. $g_{ij} = \langle u_i, u_j \rangle$. For any integer k such that $1 \leq k \leq n$, there exists a set $S \subseteq [m]$ of size k such that*

$$\frac{\det(G_{S,S})}{k!} \geq \lambda_1 \dots \lambda_k.$$

Proof. Consider the matrix $H = (\sqrt{p_i p_j} \langle u_i, u_j \rangle)_{i,j=1}^m$. This matrix has the same nonzero eigenvalues as $\sum_{i=1}^m p_i u_i^\mathsf{T} u_i$, and, therefore,

$$\sum_{S \subseteq [m] : |S| = k} \left(\prod_{i \in S} p_i \right) \det(G_{S,S}) = \sum_{S \subseteq [m] : |S| = k} \det(H_{S,S}) = s_{k,n}(\lambda),$$

where $s_{k,n}$ is the degree k elementary symmetric polynomial in n variables. (See the proof of Lemma 7 for a justification of the final equality.) Therefore,

$$\max_{S \subseteq [m] : |S| = k} \det(G_{S,S}) \geq \frac{s_{k,n}(\lambda)}{s_{k,m}(p)}.$$

We have the trivial inequality

$$s_{k,n}(\lambda) \geq \lambda_1 \dots \lambda_k,$$

since $\lambda_1 \dots \lambda_k$ is one of the terms of $s_{k,n}(\lambda)$.

To bound $s_{k,m}(p)$ from above, observe that

$$s_{k,m}(p) \leq \frac{(p_1 + \dots + p_k)^k}{k!} = \frac{1}{k!},$$

since each term of $s_{k,m}(p)$ appears exactly $k!$ times in $(p_1 + \dots + p_k)^k$. Combining the inequalities finishes the proof. $\square$

Proof of Theorem 21. We can approximate the unit ball B_X of X arbitrarily well by a symmetric polytope, so we may assume that $B_X = \text{conv}\{\pm x_1, \dots, \pm x_m\}$. Then, by Lemma 25, there exists an operator $R : \ell_2^n \to Y$, $\ell^*(R) \leq 1$ and non-negative reals $p_1, \dots, p_m \geq 0$, $\sum_{i=1}^m p_i = 1$, such that

$$\lambda(U)^{2/3} = \text{tr}((RU(\sum_{i=1}^m p_i x_i \otimes x_i)U^* R^*)^{1/3}) = \text{tr}((\sum_{i=1}^m p_i RU(x_i) \otimes RU(x_i))^{1/3}).$$

Let $u_i \triangleq RU(x_i) \in \mathbb{R}^n$, and let $\lambda_1 \geq \ldots \geq \lambda_n \geq 0$ be the eigenvalues of $\sum_{i=1}^m p_i u_i \otimes u_i$, so that $\lambda(U)^{2/3} = \sum_{i=1}^n \lambda_i^{1/3}$. We have the following elementary but very useful inequality, which is an approximate reverse of the AM-GM inequality:

$$\sum_{i=1}^n \lambda_i^{1/3} \leq \sum_{i=1}^n \left(\prod_{j=1}^i \lambda_j\right)^{1/3i} = \sum_{i=1}^n \frac{1}{i} \cdot i \left(\prod_{j=1}^i \lambda_j\right)^{1/3i}$$
$$\leq \left(\sum_{i=1}^n \frac{1}{i}\right) \max_{i=1}^n i \left(\prod_{j=1}^i \lambda_j\right)^{1/3i}.$$

Let us fix a value k so that the maximum on the right hand side is achieved. Then, the above inequality implies

$$\lambda(U)^{2/3} \leq C_0(1+\log n)k(\lambda_1 \ldots \lambda_k)^{1/(3k)},$$

for an absolute constant C_0. Observe that the matrix of $\sum_{i=1}^m p_i u_i \otimes u_i$ with respect to the standard basis is $\sum_{i=1}^m p_i u_i^\mathsf{T} u_i$, so, by Lemma 26, there exists a set $S \subseteq [m]$ of size k such that

$$(\lambda_1 \ldots \lambda_k)^{1/k} \leq \frac{\det(G_{S,S})^{1/k}}{(k!)^{1/k}},$$

where G is the Gram matrix of $u_1, \ldots, u_m$. By Stirling's estimate, this implies that

$$\lambda(U)^{2/3} \leq C_1(1+\log n)k^{2/3} \det(G_{S,S})^{1/(3k)},$$

or, equivalently,

$$\lambda(U) \leq C_1^{3/2}(1+\log n)^{3/2} k \det(G_{S,S})^{1/(2k)}, \tag{36}$$

for an absolute constant C_1.

To finish the proof, we need to relate the right hand side of (36) to the volume lower bound. Let $V : \ell_1^S \to X$ be the operator defined by $V(a) \triangleq \sum_{i \in S} a_i x_i$, where ℓ_1^S is the coordinate subspace of ℓ_2^n spanned by the standard basis vectors $\{e_i : i \in S\}$. We have that

$$\det(G_{S,S})^{1/2} = \frac{\mathrm{vol}_k(RUV(B_2^n \cap \mathbb{R}^S))}{\mathrm{vol}_k(B_2^n \cap \mathbb{R}^S)} = \frac{\mathrm{vol}_k(V^*U^*R^*(B_2^n \cap W))}{\mathrm{vol}_k(B_2^k)},$$

where W is the range of RUV in $\mathbb{R}^n$. By the definition of entropy numbers,

$$\mathrm{vol}_k(V^*U^*R^*(B_2^n \cap W)) \leq \mathrm{vol}_k(V^*U^*R^*(B_2^n)) \leq 2^{k-1} e_k(R^*)^k \,\mathrm{vol}_k(V^*U^*(B_{Y^*})),$$

since $R^*(B_2^n)$ can be covered by 2^{k-1} translates of $e_k(R^*)B_{Y^*}$. By (33) and (34), we have

$$e_k(R^*) \leq C_2 \frac{\ell(R^*)}{\sqrt{k}} \leq C_2 K(Y) \frac{\ell^*(R)}{\sqrt{k}} \leq C_2 K(Y) \frac{1}{\sqrt{k}},$$

for an absolute constant C_2. Combining the inequalities so far, we have

$$\det(G_{S,S})^{1/2k} \leq 2C_2K(Y) \frac{\mathrm{vol}_k(V^*U^*(B_{Y^*}))^{1/k}}{\sqrt{k}\,\mathrm{vol}_k(B_2^k)^{1/k}}$$
$$\leq 2C_2K(Y) \frac{\mathrm{vol}_k(B_2^k)^{1/k}}{\sqrt{k}\,\mathrm{vol}_k(\{a : \|UVa\|_Y \leq 1\})^{1/k}},$$

where in the final step we used the Blaschke-Santaló inequality and the fact that

$$(V^*U^*(B_{Y^*}))^\circ = \{a : \|UVa\|_Y \le 1\}.$$

By a standard estimate, $\frac{\mathrm{vol}_k(B_2^k)^{1/k}}{\sqrt{k}} \le \frac{C_3}{k}$ for a constant C_3, and, combining with (36), we get

$$\lambda(U) \le \frac{CK(Y)(1+\log n)^{3/2}}{\mathrm{vol}_k(\{a : \|UVa\|_Y \le 1\})^{1/k}} \le CK(Y)(1+\log n)^{3/2}\,\mathrm{volLB}^{\mathrm{h}}(U),$$

for an absolute constant C, as desired. The statement after "moreover" follows by observing that we can assume $x_1, \dots, x_m$ to be extreme points of B_X, and since the points Ve_i are a subset of them, they are extreme as well. $\square$

4.5 Proofs of Convexity and Duality

In this section we supply the missing proofs of Lemmas 24 and 25, i.e. the fact that the convex optimization problem (26)–(28) is indeed convex and equal to $\lambda(U)$, and also the derivation of the dual maximization problem.

We first prove some key technical properties of the function f.

Lemma 27. *The following statements hold for the function f defined on positive definite operators $A : X \to X^*$ as in* (25)*:*

- *f is a differentiable convex function on positive definite operators $A : X \to X^*$;*
- *f is given by the formula*

$$f(A) = \sup\{\mathrm{tr}((RUA^{-1}U^*R^*)^{1/2}) : R : Y \to \ell_2^n, \ell^*(R) \le 1\}; \tag{37}$$

- *the derivative of f at A is*

$$\nabla f(A) = -\frac{1}{2}(RU)^{-1}((RU)^{-*}A(RU)^{-1})^{-3/2}(RU)^{-*}, \tag{38}$$

  where $R : Y \to \ell_2^n$, $\ell^*(R) \le 1$ is such that $f(A) = \mathrm{tr}((RUA^{-1}U^*R^*)^{1/2})$.

Note that the derivative $\nabla f(A)$ is a linear functional on operators from X to X^*, and, via the trace, we identify such functionals with operators from X^* to X. In what follows we will, therefore, treat $\nabla f(A)$ as a linear operator from X^* to X.

Before we prove Lemma 27, we need an auxiliary lemma.

Lemma 28. *Let Y be an n-dimensional normed space, and let $S : \ell_2^n \to Y$ be an invertible linear operator. Then any operator $R : Y \to \ell_2^n$ such that $\mathrm{tr}(RS) = \ell(S)$ is invertible.*

Proof. Assume for contradiction that R is not invertible, i.e. it has a non-trivial kernel. Let $k < n$ be the dimension of the kernel of R, and let W be the k-dimensional subspace of $\mathbb{R}^n$ such that $S(W)$ is the kernel of R. Then, if π is the orthogonal projection onto the orthogonal complement $W^\perp$ of W, we have

$$\ell(S) = \mathrm{tr}(RS) = \mathrm{tr}(RS\pi) \le \ell(S\pi).$$

Define the convex body $K = S^{-1}(B_Y)$. Using integration by parts, we have

$$\begin{aligned}\ell(S\pi)^2 &= \int_{t=0}^{\infty}(1-\gamma_{n-k}(\sqrt{t}K\cap W^{\perp}))dt\\ \ell(S)^2 &= \int_{t=0}^{\infty}(1-\gamma_n(\sqrt{t}K))dt\\ &= \int_{t=0}^{\infty}\int_W (1-\gamma_{n-k}(\sqrt{t}K\cap(y+W^{\perp}))d\gamma_k(y))dt\end{aligned}$$

By the log-concavity of Gaussian measure and the symmetry of K, $\gamma_{n-k}(\sqrt{t}K\cap(y+W^{\perp})) \le \gamma_{n-k}(\sqrt{t}K\cap W^{\perp})$ for all $y\in W$, and, for all y outside a compact set we have $\gamma_{n-k}(\sqrt{t}K\cap(y+W^{\perp})) = 0 < \gamma_{n-k}(\sqrt{t}K\cap W^{\perp})$. Therefore, $\ell(S\pi) < \ell(S)$, a contradiction. $\square$

Proof of Lemma 27. We begin with the proof of differentiability. Because A is positive definite, it defines an inner product $\langle y,x\rangle_A = \langle y,Ax\rangle$ on X, and a corresponding Gaussian measure γ_A. Let us fix some positive definite operator $A_0 : X\to X^*$. Then, for any positive definite $A : X\to X^*$ we can write

$$f(A) = \int_X \|Ux\|_Y d\gamma_A(x) = \int_X \|Ux\|_Y \frac{d\gamma_A}{d\gamma_{A_0}}(x)\ d\gamma_{A_0}(x).$$

For any $x\in X$,

$$\frac{d\gamma_A}{d\gamma_{A_0}}(x) = \sqrt{\frac{\det(A)}{\det(A_0)}}e^{-\langle x,(A-A_0)x\rangle/2}.$$

Since $\frac{d\gamma_A}{d\gamma_{A_0}}(x)$ is easily seen to be continuously differentiable in A, by the dominated convergence theorem the derivative of f also exists and is given by differentiating under the integral sign.

Next we prove the identity (37). Observe that for any orthogonal transformation $O:\ell_2^n\to\ell_2^n$ and any operator $V:\ell_2^n\to Y$, $\ell(VO)=\ell(V)$ by the rotational invariance of the Gaussian measure. Therefore,

$$\begin{aligned}f(A) = \ell(UT^{-1}) &= \sup\{\ell(UT^{-1}O) : O \text{ orthogonal}\}\\ &= \sup\{\mathrm{tr}(RUT^{-1}O) : R : Y\to\ell_2^n, \ell^*(R)\le 1, O \text{ orthogonal}\}\\ &= \sup\{\nu(RUT^{-1}) : R : Y\to\ell_2^n, \ell^*(R)\le 1\}\\ &= \sup\{\mathrm{tr}((RUA^{-1}U^*R^*)^{1/2}) : R : Y\to\ell_2^n, \ell^*(R)\le 1\}.\end{aligned}$$

Moreover, since we assumed that U is invertible, by Lemma 28, we can assume that R is invertible. Given this formula, in order to prove convexity, it is enough to prove that for any invertible operator $V : X\to\ell_2^n$ (which, in our case, equals RU), the function $\mathrm{tr}((VA^{-1}V^*)^{1/2})$ is convex in A for $A : X\to X^*$ positive definite. Then, we would have that $f(A)$ is a supremum of convex functions, and, therefore, convex.

The convexity of $\mathrm{tr}((VA^{-1}V^*)^{1/2})$ follows from a standard argument based on majorization and Schur convexity, which we give next. Since V is invertible, we have $VA^{-1}V^* = (V^{-*}AV^{-1})^{-1}$, and $\mathrm{tr}((VA^{-1}V^*)^{1/2}) = \mathrm{tr}(((V^{-*}AV^{-1})^{-1/2})$. Let $\alpha\in(0,1)$ be arbitrary, and let A_1, A_2 be two positive definite operators from X to X^*. Let μ be the function that maps a self-adjoint operator on ℓ_2^n to the vector of its eigenvalues. By the Ky-Fan inequalities, $\mu(V^{-*}(\alpha A_1+(1-\alpha)A_2)V^{-1})$ is majorized by $\alpha\mu(V^{-*}A_1V^{-1})+(1-\alpha)\mu(V^{-*}A_2V^{-1})$. Let g be the function defined on vectors $x\in\mathbb{R}^n$ with positve coordinates by $g(x)=\sum_{i=1}^n x_i^{-1/2}$. It is easy to verify that g is convex and

Schur-convex, so

$$\begin{aligned} g(\mu(V^{-*}(\alpha A_1 + (1-\alpha)A_2)V^{-1})) &\le g(\alpha\mu(V^{-*}A_1V^{-1}) + (1-\alpha)\mu(V^{-*}A_2V^{-1})) \\ &\le \alpha g(\mu(V^{-*}A_1V^{-1})) + (1-\alpha)g(\mu(V^{-*}A_2V^{-1})). \end{aligned}$$

Since the left hand side above equals

$$\operatorname{tr}((V^{-*}(\alpha A_1 + (1-\alpha)A_2)V^{-1})^{-1/2}) = \operatorname{tr}((V(\alpha A_1 + (1-\alpha)A_2)^{-1}V^*)^{1/2}),$$

and the right hand side equals

$$\begin{aligned} &\alpha \operatorname{tr}((V^{-*}A_1V^{-1})^{-1/2}) + (1-\alpha)\operatorname{tr}((V^{-*}A_2V^{-1})^{-1/2}) \\ &\qquad = \alpha \operatorname{tr}((VA_1^{-1}V^*)^{1/2}) + (1-\alpha)\operatorname{tr}((VA_2^{-1}V^*)^{1/2}), \end{aligned}$$

we have established convexity.

From (37), we can see that the subgradient of f at A is

$$\begin{aligned} \partial f(A) &= \operatorname{conv}\{\nabla \operatorname{tr}((RUA^{-1}U^*R^*)^{1/2}) : \\ &\qquad R : Y \to \ell_2^n, \ell^*(R) \le 1, f(A) = \operatorname{tr}((RUA^{-1}U^*R^*)^{1/2})\} \\ &= \operatorname{conv}\{\nabla \operatorname{tr}(((RU)^{-1}A(RU)^{-*})^{-1/2}) : \\ &\qquad R : Y \to \ell_2^n, \ell^*(R) \le 1, f(A) = \operatorname{tr}((RUA^{-1}U^*R^*)^{1/2})\} \\ &= \operatorname{conv}\{-\frac{1}{2}(RU)^{-1}((RU)^{-*}A(RU)^{-1})^{-3/2}(RU)^{-*} : \\ &\qquad R : Y \to \ell_2^n, \ell^*(R) \le 1, f(A) = \operatorname{tr}((RUA^{-1}U^*R^*)^{1/2})\}. \end{aligned}$$

Above we used Lemma 28, and the fact that

$$\nabla \operatorname{tr}(X^{-1/2}) = -\frac{1}{2}X^{-3/2}$$

for any positive definite $X : \ell_2^n \to \ell_2^n$ (see [Lew95]). Since f is differentiable, we have that $\partial f(A)$ is a singleton set, i.e.

$$\nabla f(A) = -\frac{1}{2}(RU)^{-1}((RU)^{-*}A(RU)^{-1})^{-3/2}(RU)^{-*}$$

for the an invertible operator $R : Y \to \ell_2^n$, such that $\ell^*(R) \le 1$ and $f(A) = \operatorname{tr}((RUA^{-1}U^*R^*)^{1/2})$. This finishes the proof of the lemma. □

We are now ready to prove Lemma 24, restated for convenience below.

Lemma 24. *For any two n-dimensional normed spaces X and Y, and any invertible linear operator $U : X \to Y$, $\lambda(U)$ equals*

$$\inf f(A) \tag{26}$$

s.t.

$$A : X \to X^*, \|A\| \le 1 \tag{27}$$

$$A \succ 0. \tag{28}$$

The function f is the one defined in (25).

Moreover, the objective (26) *and the constraints* (27)–(28) *are convex in A.*

Proof. The convexity of the constraints (27)–(28) is apparent from the definition, and the convexity of the objective was proved in Lemma 27. We proceed to show that the value of (26)–(28) equals $\lambda(U)$.

Let $T : X \to \ell_2^n$ and $S : \ell_2^n \to Y$, $ST = U$ be a factorization achieving $\lambda(U)$ such that $\|T\| = 1$ and $\ell(S) = \lambda(U)$. Then we claim that $A \triangleq T^*T$ satisfies (27)–(28) and $f(A) = \ell(S)$, so the value of (26)–(28) is at most $\ell(S) = \lambda(U)$. Indeed, A is clearly positive semidefinite, and must be positive definite, as T is invertible, because U is invertible. Furthermore, since $\|T\| \le 1$, we have that for any $x \in B_X$,

$$\langle x, Ax\rangle = \langle Tx, Tx\rangle = \|Tx\|_2^2 \le 1.$$

On the other hand, since A is self-adjoint, $\|A\| = \sup\{\langle x, Ax\rangle : x \in B_X\}$, and we have shown that $\|A\| \le 1$. Moreover, since $A = T^*T$, by the definition of $f(A)$

$$f(A) = \ell(UT^{-1}) = \ell(S) = \lambda(U).$$

This finishes the proof of the claim that the value of (26)–(28) is at most $\lambda(U)$.

Next we prove the reverse inequality. Given a feasible solution A to (26)–(28), we take an operator $T : X \to \ell_2^n$ such that $A = T^*T$. Then we construct a factorization $U = ST$ by setting $S = UT^{-1}$. By (27)–(28), for any $x \in B_X$ we have

$$\|Tx\|_2^2 = \langle Tx, Tx\rangle = \langle x, Ax\rangle \le 1,$$

so $\|T\| \le 1$. Moreover, $\ell(S) = f(A)$ by definition. This proves that the value of (26)–(28) is at least $\lambda(U)$, and, since we already showed that it is also at most $\lambda(U)$, the two are equal. $\square$

The derivation of the dual formulation from the convex program (26)–(28) is mostly routine using Lagrange duality (see, e.g. [BV04]). Nevertheless, because of the complicated nature of our objective, the derivation is quite technical. Once again, we restated Lemma 25, which gives our dual formulation, for convenience before the proof.

Lemma 25. *Let X and Y be two n-dimensional normed spaces, such that the unit ball of X is $B_X = \operatorname{conv}\{\pm x_1, \dots, \pm x_m\}$. Then, for any linear operator $U : X \to Y$, $\lambda(U)$ equals*

$$\sup \operatorname{tr}((RU(\sum_{i=1}^m p_i x_i \otimes x_i)U^*R^*)^{1/3})^{3/2} \tag{29}$$

s.t.

$$R : Y \to \ell_2^n, \ell^*(R) \le 1 \tag{30}$$

$$\sum_{i=1}^m p_i = 1 \tag{31}$$

$$p_1, \dots, p_m \ge 0. \tag{32}$$

Proof. Since $B_X = \operatorname{conv}\{\pm x_1, \dots, \pm x_m\}$, we can can rewrite (26)–(28) as

$$\inf f(A) \quad \text{s.t.} \tag{39}$$

$$\langle x_i, Ax_i\rangle \le 1 \;\; \forall i \in [m] \tag{40}$$

$$A \succ 0, \tag{41}$$

By Lemma 24 this is a convex minimization problem and its value equals $\lambda(U)$.

The conjugate function f^* of f is defined on self-adjoint linear operators $Q : X^* \to X$ by:

$$f^*(Q) \triangleq \sup\{\operatorname{tr}(QA) - f(A) : A \succ 0\}.$$

Since $f^*(Q)$ is the supremum of affine functions, it is convex and lower semicontinuous. The set on which f^* takes a finite value is called its domain. The significance of f^* is the fact

$$\lambda(U) = \sup\left\{-\sum_{i=1}^m q_i - f^*\left(-\sum_{i=1}^m q_i x_i \otimes x_i\right) : q_1, \ldots, q_m \geq 0\right\}. \tag{42}$$

This follows from the duality theory of convex optimization, since (39)–(41) satisfies Slater's condition, which in this case reduces to just checking the existence of a feasible A (see [BV04, Chapter 5]). We proceed to compute $f^*(Q)$.

It is easy to see that unless $-Q \succeq 0$, $f^*(Q) = \infty$, and, conversely, if $-Q \succeq 0$, $f^*(Q) \leq 0 < \infty$. Therefore, the domain of f^* is $\{Q : -Q \succeq 0\}$. We will first handle the case $-Q \succ 0$, and then we will extend our formula for $f^*(Q)$ to $-Q \succeq 0$ by continuity. Assume then that $-Q \succ 0$. As f is a differentiable convex function, the range of the derivative ∇f includes the relative interior of the domain of f^*, i.e. the set of linear operators $\{X : -X \succ 0\}$ (see Corollary 26.4.1. in [Roc70]); from (38) it is also apparent that $\nabla f(A)$ is negative definite for any positive definite A, so the range of $\nabla f(A)$ is exactly $\{X : -X \succ 0\}$. This means that the equation

$$0 = \nabla(\operatorname{tr}(QA) - f(A)) = Q - \nabla f(A)$$

has a solution over $A \succ 0$, and $f^*(Q)$ is achieved at this solution. Fix $A : X \to X^*$ to be a solution to this equation, and let $R : Y \to \ell_2^n$ be an invertible map such that $\ell^*(R) \leq 1$ and

$$f(A) = \operatorname{tr}((RUA^{-1}U^*R^*)^{1/2}) = \operatorname{tr}(((RU)^{-*}A(RU)^{-1})^{-1/2}).$$

The equations $\nabla f(A) = Q$ and (38) imply

$$((RU)^{-*}A(RU)^{-1})^{-3/2} = -2RUQU^*R^*,$$

and, therefore,

$$\begin{aligned}
f^*(Q) &= \operatorname{tr}(QA) - \operatorname{tr}((RUA^{-1}U^*R^*)^{1/2})\\
&= \operatorname{tr}((RUQU^*R^*)((RU)^{-*}A(RU)^{-1})) - \operatorname{tr}(((RU)^{-*}A(RU)^{-1})^{-1/2})\\
&= -\frac{3}{2^{2/3}}\operatorname{tr}((-RUQU^*R^*)^{1/3}).
\end{aligned}$$

We have proved that

$$f^*(Q) \geq \inf\left\{-\frac{3}{2^{2/3}}\operatorname{tr}((-RUQU^*R^*)^{1/3}) : R : Y \to \ell_2^n, \ell^*(R) \leq 1\right\}, \tag{43}$$

for any $Q : X^* \to X$ such that $-Q \succ 0$. Let $\mathcal{D}$ be the set of invertible maps $R : Y \to \ell_2^n$ such that $\ell^*(R) \leq 1$. By (37) and the definition of the conjugate function f^* we also have

$$\begin{aligned}
f^*(Q) &= \sup_{A:A\succ 0}\inf_{R\in\mathcal{D}} \operatorname{tr}(QA) - \operatorname{tr}((RUA^{-1}U^*R^*)^{1/2})\\
&\leq \inf_{R\in\mathcal{D}}\sup_{A:A\succ 0} \operatorname{tr}(QA) - \operatorname{tr}((RUA^{-1}U^*R^*)^{1/2}).
\end{aligned}$$

For any $R \in \mathcal{D}$, the supremum on the right hand side equals the value at Q of the conjugate h_R^* of the function $h_R(A) = \operatorname{tr}((RUA^{-1}U^*R^*)^{1/2})$. A calculation analogous to the one above for f^* shows that $h_R^*(Q) = -\frac{3}{2^{2/3}} \operatorname{tr}((-RUQU^*R^*)^{1/3})$.

Therefore,

$$f^*(Q) \le \inf\left\{-\frac{3}{2^{2/3}} \operatorname{tr}((-RUQU^*R^*)^{1/3}) : R : Y \to \ell_2^n, \ell^*(R) \le 1\right\},$$

and, together with (43), we have established

$$f^*(Q) = \inf\left\{-\frac{3}{2^{2/3}} \operatorname{tr}((-RUQU^*R^*)^{1/3}) : R : Y \to \ell_2^n, \ell^*(R) \le 1\right\}, \tag{44}$$

for any $Q : X^* \to X$ such that $-Q \succ 0$. Since f^* is a a proper lower-semicontinuous function, it is continuous on any line segment contained in its domain by Corollary 7.5.1. in [Roc70]. Therefore, (44) holds for any $Q \succeq 0$ as well.

By (42) and (44),

$$\begin{aligned}\lambda(U) = \sup\bigg\{ -\sum_{i=1}^m q_i + \frac{3}{2^{2/3}} \operatorname{tr}\bigg(\Big(-RU\Big(\sum_{i=1}^m q_i x_i \otimes x_i\Big)U^*R^*\Big)^{1/3}\bigg) :& \\ R : Y \to \ell_2^n, \ell^*(R) \le 1, q_1, \dots, q_m \ge 0\bigg\}.& \end{aligned} \tag{45}$$

Let us write $q = tp$ where $t \ge 0$ is a real number, and $p_1, \dots, p_m \ge 0$ satisfy $\sum_i p_i = 1$. Then we can rewrite the equation above as

$$\lambda(U) = \sup\left\{-t + \frac{3t^{1/3}}{2^{2/3}} \operatorname{tr}\bigg(\Big(-RU\Big(\sum_{i=1}^m p_i x_i \otimes x_i\Big)U^*R^*\Big)^{1/3}\bigg) : t, p_1, \dots, p_m \ge 0, \sum_{i=1}^m p_i = 1\right\}.$$

Maximizing over t finishes the proof. □

5 Algorithm for the Factorization Constant

In this section we use our convex formulation (26)–(28) of the $\lambda(U)$ factorization constant in order to compute an approximately optimal factorization. In order to use known results in convex optimization, we need to make sure that we are optimizing a Lipschitz function over a sufficiently bounded feasible region, and, moreover, that we have a strictly feasible point. These conditions are not automatically satisfied for (26)–(28), but we can modify the optimization problem so that they are, at the cost of a small constant factor approximation to the optimum.

In this section we assume that both normed spaces X and Y are defined over $\mathbb{R}^n$. We assume that the unit ball of X is $B_X = \operatorname{conv}\{\pm x_1, \dots, \pm x_m\}$, and that X is specified by giving the points $x_1, \dots, x_m$ as input to the algorithm. We assume that Y is specified by an evaluation oracle, which takes a point $y \in \mathbb{R}^n$ and returns $\|y\|_Y$. Moreover, we will assume that U is the identity map; otherwise we can define a new norm Z by $\|z\|_Z = \|Uz\|_Y$ and we have $\lambda(U) = \lambda(I)$ for the identity map $I : X \to Z$. An evaluation oracle for Y easily gives an evaluation oracle for Z.

As a first step, we transform the problem so that B_X is well-rounded. Roughly speaking, this makes the problem well-conditioned. For example, it will allow us to determine $\lambda(I)$ up to a factor of $O(\sqrt{n})$. To round B_X, we use a classical algorithm of Khachiyan.

Theorem 29 ([Kha96]). *There exists an algorithm running in time $O(mn^2(\log n + \log\log m)$ that, given a set of points $x_1, \dots, x_m \in \mathbb{R}^n$, computes a linear map T such that*

$$\frac{1}{2\sqrt{n}} B_2^n \subseteq T(\mathrm{conv}\{\pm x_1, \dots, \pm x_m\}) \subseteq B_2^n.$$

We compute the linear map T for the extreme points $\pm x_1, \dots, \pm x_m$, using the algorithm guaranteed by Theorem 29, and apply T to both B_X and B_Y. I.e. we replace X with the space whose unit ball is TB_X, and Y with the space whose unit ball is TB_Y. This does not change $\lambda(I)$, $\mathrm{vb}(I)$, or the volume lower bound. With this transformation, we can assume that

$$\frac{1}{2\sqrt{n}} B_2^n \subseteq B_X \subseteq B_2^n. \tag{46}$$

In the rest of this section, we use the notation $A \succeq B$ for two symmetric matrices A and B to denote the fact that $A - B$ is positive semidefinite. The notation $A \preceq B$ is equivalent to $B \succeq A$. We also repeatedly use the fact that if $AA^{\mathsf{T}} \preceq BB^{\mathsf{T}}$ for two matrices A and B, then, for a standard Gaussian Z in $\mathbb{R}^n$ and any norm Y defined on $\mathbb{R}^n$, $\mathbb{E}\|AZ\|_Y^2 \le \mathbb{E}\|BZ\|_Y^2$. This is well-known, and is due to the fact that BZ is distributed identically to $AZ + CZ'$ for a standard Gaussian Z' independent from Z. Then, by Jensen's inequality,

$$\mathbb{E}\|AZ\|_Y^2 = \mathbb{E}\|AZ + \mathbb{E}[CZ']\|_Y^2 \le \mathbb{E}\|AZ + CZ'\|_Y^2 = \mathbb{E}\|BZ\|_Y^2.$$

Our first lemma shows that we can strengthen the constraints (27)–(28) without affecting the value of the optimization problem significantly. The stronger constraints will be helpful in showing that the objective is Lipschitz and bounded over the feasible region.

Lemma 30. *Assume that X and Y are normed spaces over $\mathbb{R}^n$ such that $B_X = \mathrm{conv}\{\pm x_1, \dots, \pm x_m\}$, and equation (46) holds, and let $I : X \to Y$ be the identity map. Then the value of the following convex optimization problem over positive definite matrices A is at least $\lambda(I)$ and at most $\sqrt{2}\lambda(I)$:*

$$\inf (\mathbb{E}\|A^{-1/2}Z\|_Y^2)^{1/2} \tag{47}$$

$$s.t.$$

$$x_i^T A x_i \le 1 \quad \forall i \in [m], \tag{48}$$

$$A \succeq \frac{1}{2} I. \tag{49}$$

Above $A^{-1/2}$ is the unique positive definite matrix such that $(A^{-1/2})^2 = A^{-1}$, and Z is a standard Gaussian random variable in $\mathbb{R}^n$.

Moreover, any positive definite matrix A satisfying (48) also satisfies $A \preceq 4nI$.

Proof. The objective function (47) equals (26), the constraints (48) are equivalent to (27), and (49) implies (28), so, trivially, the value of (47)–(49) is at least the value of (26)–(28), which, by Lemma 24, equals $\lambda(I)$. Moreover, again by Lemma 24, (47)–(49) is a convex optimization problem.

To show that the value of (47)–(49) is at most $\sqrt{2}\lambda(I)$, let us take an operator $A : X \to X^*$ achieving the optimal value $\lambda(I)$ in (26)–(28), and let us identify A with its matrix in the standard basis. Let $\tilde{A} \triangleq \frac{1}{2}(A + I)$. Since $B_X \subseteq B_2^n$ by assumption, and A satisfies (27), we have

$$\frac{1}{2} x_i^T (A + I) x_i = \frac{x_i^T A x_i + \|x_i\|_2^2}{2} \le 1,$$

and, therefore, $\tilde{A}$ satisfies (48). Because A is positive definite, we have $\tilde{A} \succeq \frac{1}{2}I$, and (49) is also satisfied, and $\tilde{A}$ is feasible. Because $\tilde{A} \succeq \frac{1}{2}A$, we have $\tilde{A}^{-1} \preceq 2A^{-1}$, so

$$\mathbb{E}\|\tilde{A}^{-1/2}Z\|_Y^2 \le 2\mathbb{E}\|A^{-1/2}Z\|_Y^2 = 2\lambda(I)^2.$$

This shows that the value of (47)–(49) is at most $\sqrt{2}\lambda(I)$.

The statement after "moreover" follows because if there exists some $x \in \mathbb{R}^n$ for which $x^\mathsf{T}Ax > 4n\|x\|_2^2$ then for $y = \frac{x}{2\sqrt{n}\|x\|_2}$ we have $y \in \frac{1}{2\sqrt{n}}B_2^n \subseteq B_X = \mathrm{conv}\{\pm x_1, \ldots, \pm x_m\}$ but $x^\mathsf{T}Ax > 1$. This would imply that for at least one extreme point x_i of B_X we have $x_i^\mathsf{T}Ax_i > 1$, in contradiction with (48). $\square$

Our second lemma shows that the objective function (47) is Lipschitz over the feasible region (48)–(49).

Lemma 31. *Under the assumptions of Lemma 30, for any A, B satisfying* (48)–(49) *we have*

$$(\mathbb{E}\|A^{-1/2}Z\|_Y^2)^{1/2} - (\mathbb{E}\|B^{-1/2}Z\|_Y^2)^{1/2} \le 16\sqrt{n}\lambda(I)\|A - B\|_{op},$$

where $\|A - B\|_{op}$ is the largest singular value of $A - B$.

Proof. Observe that, because any A which is feasible for (48)–(49) satisfies $\frac{1}{2}I \preceq A \preceq 4nI$, we have

$$\frac{1}{2\sqrt{n}}(\mathbb{E}\|Z\|_Y^2)^{1/2} \le (\mathbb{E}\|A^{-1/2}Z\|_Y^2)^{1/2} \le \sqrt{2}(\mathbb{E}\|Z\|_Y^2)^{1/2}.$$

Therefore, for any feasible A, $(\mathbb{E}\|A^{-1/2}Z\|_Y^2)^{1/2}$ is within a factor of $2\sqrt{2n}$ from the minimum of (47)–(49). By Lemma 30, we then have

$$\lambda(I) \le (\mathbb{E}\|A^{-1/2}Z\|_Y^2)^{1/2} \le 4\sqrt{n}\lambda(I).$$

Let $\delta \triangleq \|A - B\|_{op}$, and consider first the case $\delta > \frac{1}{4}$. Then,

$$(\mathbb{E}\|A^{-1/2}Z\|_Y^2)^{1/2} - (\mathbb{E}\|B^{-1/2}Z\|_Y^2)^{1/2} \le (\mathbb{E}\|A^{-1/2}Z\|_Y^2)^{1/2} \le 4\sqrt{n}\lambda(I) \le 16\sqrt{n}\lambda(I)\delta.$$

Consider now the case $\delta < \frac{1}{4}$. Then, $A - B \succeq -\delta I$, and we have

$$A = B + (A - B) \succeq B - \delta I \succeq (1 - 2\delta)B.$$

Therefore,

$$(\mathbb{E}\|A^{-1/2}Z\|_Y^2)^{1/2} \le (1 - 2\delta)^{-1/2}(\mathbb{E}\|B^{-1/2}Z\|_Y^2)^{1/2} \le (1 + 4\delta)(\mathbb{E}\|B^{-1/2}Z\|_Y^2)^{1/2}.$$

Finally,

$$(\mathbb{E}\|A^{-1/2}Z\|_Y^2)^{1/2} - (\mathbb{E}\|B^{-1/2}Z\|_Y^2)^{1/2} \le 4\delta(\mathbb{E}\|B^{-1/2}Z\|_Y^2)^{1/2} \le 16\sqrt{n}\lambda(I)\delta.$$

This completes the proof $\square$

Our final lemma verifies that, given a positive definite matrix A and an evaluation oracle for Y, we can approximate evaluate the objective (47). We do so in the natural way: we sample a random Gaussian Z and use the oracle to compute $\|A^{-1/2}Z\|_Y$. The next lemma (which is standard) shows that the resulting estimate is concentrated around its mean.

Lemma 32. *For any norm Y on $\mathbb{R}^n$, $n \times n$ matrix M, and a standard Gaussian Z in $\mathbb{R}^n$,*

$$\Pr(|\|MZ\|_Y - \mathbb{E}\|MZ\|_Y| > t\mathbb{E}\|MZ\|_Y) \leq 2e^{-4t^2/\pi^3}.$$

Proof. As usual, we can assume that M is invertible (otherwise we project to a subspace of Y) and that in fact $M = I$, by replacing Y with the norm Y' given by $\|x\|_{Y'} = \|Mx\|_Y$. Then, we have that for any $x \in \mathbb{R}^n$, of Euclidean norm $\|x\|_2 = 1$,

$$\|x\|_Y = \sqrt{\frac{\pi}{2}}\mathbb{E}\|Z_1 x\|_Y \leq \sqrt{\frac{\pi}{2}}\mathbb{E}\|Z\|_Y.$$

I.e. for all x in $\mathbb{R}^n$, $\|x\|_Y \leq \left(\sqrt{\frac{\pi}{2}}\mathbb{E}\|Z\|_Y\right) \cdot \|x\|_2$. Then the lemma follows by the Maurey-Pisier inequality (see Theorem 4.7 in [Pis89]). $\square$

We are now ready to prove our main algorithmic result.

Theorem 33. *There exists an algorithm that, given $x_1, \dots, x_m \in \mathbb{R}^n$, an evaluation oracle for a norm Y on $\mathbb{R}^n$, and a linear operator $U : X \to Y$ specified by its matrix, where X is the space with unit ball $B_X = \operatorname{conv}\{\pm x_1, \dots, \pm x_m\}$, computes in time polynomial in m and n a factorisation $U = ST$, $S : \ell_2^n \to Y$, $T : X \to \ell_2^n$, such that*

$$\ell(S)\|T\| \leq C\lambda(U),$$

for an absolute constant C.

Proof. As discussed above, we can reduce to the case when U is the identity and B_X satisfies (46). We are going to solve the optimization problem

$$\begin{aligned}
&\inf \mathbb{E}\|A^{-1/2}Z\|_Y \\
&\text{s.t.} \\
&x_i^T A x_i \leq 1 \quad \forall i \in [m], \\
&A \succeq \frac{1}{2}I.
\end{aligned}$$

This is the same problem as (47)–(49), but with a slightly modified objective. However, this new objective is the same as (47) up to a constant factor (see Corollary 4.9 in [Pis89]). By Lemma 30, any feasible A satisfies $\|A\|_{op} \leq 4n$. Moreover, the solution $\frac{3}{4}I$ is strictly feasible in the sense that any A satisfying $\|A - \frac{3}{4}I\|_{op} \leq \frac{1}{4}$ satisfies the constraints. Indeed, $A \succeq \frac{1}{2}I$ is immediate, and we also have $A \preceq I$, which implies $x_i^{\mathsf{T}} A x_i \leq \|x_i\|_2^2 \leq 1$, by (46). Then our problem reduces to optimizing a convex function over a convex set with a stochastic zero-order oracle with subgaussian error. A polynomial time algorithm for this problem is given, for example, in Section 6 of [BLNR15]. $\square$

Theorem 33 and Corollaries 22 and 23 imply that we can efficiently approximate both the vector balancing constant and hereditary discrepancy in any norm.

Corollary 34. *Let $x_1, \dots, x_m$, X, Y, and U be as in Theorem 33. Then the vector balancing constant $\operatorname{vb}(U)$ can be approximated in time polynomial in n and the number of vertices of C up to a factor of $O(K(Y)(1 + \log n)^{3/2})$, where $K(Y) = O(\log n)$ is the K-convexity constant of Y.*

Moreover, for any points $u_1, \dots, u_N$, the hereditary discrepancy $\operatorname{hd}((u_i)_{i=1}^N, B_Y)$ can be approximated in time polynomial in n and N up to the same factor of $O(K(Y)(1 + \log n)^{3/2})$.

6 Properties of the Bounds

6.1 The Volume Lower Bound and Convex Hulls

Here we show that the volume lower bound is maximized at the extreme points of a convex set. This shows that the fact that Theorem 21 relates λ to volume lower bounds given by extreme points is not an accident. In general, it shows that the volume lower bound does not distinguish between vector balancing and hereditary discrepancy.

Theorem 35. *Let $v_1, \ldots, v_m$ be points in $\mathbb{R}^n$, and let $C \triangleq \mathrm{conv}\{\pm v_1, \ldots, \pm v_m\}$. Then, for any $k \in \mathbb{N}$, $1 \le k \le n$, and any symmetric convex body K in $\mathbb{R}^n$,*

$$\sup_{u_1,\ldots,u_k \in C} \mathrm{volLB}((u_i)_{i=1}^k, K) \le \mathrm{volLB}_k^{\mathrm{h}}((v_i)_{i=1}^m, K).$$

We will use a theorem of K. Ball [Bal88], which allows us to define a norm associated with an arbitrary logarithmically concave function f.

Theorem 36. *Let $f : \mathbb{R}^k \to [0, \infty)$ be an even logarithmically concave function such that $0 < \int_{\mathbb{R}^k} f < \infty$. Then, for any $p \ge 1$,*

$$\|x\|_{f,p} \triangleq \begin{cases} \left(\int_0^\infty f(rx) r^{p-1} dr\right)^{-1/p}, & x \ne 0, \\ 0, & x = 0. \end{cases}$$

defines a norm on $\mathbb{R}^n$.

Proof of Theorem 35. Let us fix a sequence of linearly independent vectors $u_1, \ldots, u_{k-1} \in C$ for the remainder of the proof, and, for any $x \in \mathbb{R}^n$, define the matrix $U_x = (u_1, \ldots u_{k-1}, x)$. To prove the theorem, it is enough to show that the function $g : \mathbb{R}^n \to [0, \infty)$ defined by

$$g(x) = \mathrm{volLB}((u_1, \ldots, u_{k-1}, x), K)^k = \frac{1}{\mathrm{vol}_k(\{a \in \mathbb{R}^k : U_x a \in K\})}$$

achieves its maximum on C at one of the extreme points $\pm v_1, \ldots, \pm v_m$. This follows immediately if g is convex. Below, we use Theorem 36 to prove the convexity of g.

Let $W = \mathrm{span}\{u_1, \ldots, u_{k-1}\}$ and let π be the orthogonal projection onto the orthogonal complement $W^\perp$ of W. Notice that

$$\mathrm{vol}_k(\{a \in \mathbb{R}^k : U_x a \in K\}) = \mathrm{vol}_k(\{a \in \mathbb{R}^k : U_{\pi x} a \in K\}).$$

To show that that g is convex, it is, therefore, enough to show that it is convex on $W^\perp$. For any $x \in W^\perp$, define

$$L_x = \{b \in \mathbb{R}^{k-1} : Ub + x \in K\},$$

where U is the matrix $(u_1, \ldots, u_{k-1})$. It is easy to check that for any two $x, x' \in W^\perp$, and any $\alpha \in [0,1]$, $\alpha L_x + (1-\alpha) L_{x'} \subseteq L_{\alpha x + (1-\alpha)x'}$. Therefore, by the Brunn-Minkowski inequality, the function $h(x) = \mathrm{vol}(L_x)$ is logarithmically concave. Moreover, by the symmetry of K, $L_{-x} = -L_x$, so $h(x) = h(-x)$. By Theorem 36 it follows that

$$g(x) = \frac{1}{\int_{-\infty}^\infty h(tx) dt} = \frac{1}{2\int_0^\infty h(tx) dt} = \frac{1}{2} \|x\|_{h,1}.$$

is a norm on $W^\perp$, and, therefore, convex. The theorem follows. $\square$

7 Approximate Caratheodory Estimates

In this section we introduce connections between vector balancing, and approximate versions of Caratheodory's theorem. Recall that, by Caratheodory's theorem, for any N vectors $u_1, \dots, u_N \in \mathbb{R}^n$ and a point z in their convex hull, there exist $k \le n+1$ vectors $v_1, \dots, v_k \in \{u_1, \dots, u_N\}$ such that $z \in \mathrm{conv}\{v_1, \dots, v_k\}$. While this is easily seen to be optimal in general, improvements are possible if we only need to approximately represent z. For example, it is well-known that if $u_1, \dots, u_N \in B_2^n$, then for any $z \in \mathrm{conv}\{u_1, \dots, u_N\}$ and any $\varepsilon \in (0,1)$, there exist $k = O(\varepsilon^{-2})$ vectors $v_1, \dots, v_k \in \{u_1, \dots, u_N\}$ such that $\left\| z - \frac{1}{k}\sum_{i=1}^k v_i \right\|_2 \le \varepsilon$. Note that the bound on k in this approximate version of Caratheodory's theorem is independent of the dimension n, and that z is close to the avarage of $v_1, \dots, v_k$, rather than merely in their convex hull. This dimension-independent bound has been used to speed up algorithms in computational geometry...

We call bounds of the type given above for vectors in B_2^n approximate Caratheodory estimates. Such estimates are also known for B_p^n, where the distance from z is measured in the ℓ_p^n norm, and follow from type and cotype theory. (Other proofs are also known.) Here we systematically explore approximate Caratheodory estimates and connect them to vector balancing. Let us define the k-vector approximate Caratheodory constant $\mathrm{ac}_k(C, K)$ from a convex body $C \subset \mathbb{R}^n$ to a symmetric convex body $K \subset \mathbb{R}^n$ as the smallest constant a such that for any $u_1, \dots, u_N \in C$, and any $z \in \mathrm{conv}\{u_1, \dots, u_N\}$, there exist (not necessarily distinct) vectors $v_1, \dots, v_k \in \{u_1, \dots, u_N\}$ such that

$$\left\| z - \frac{1}{k}\sum_{i=1}^{k} v_i \right\|_K \le \frac{a}{k}. \tag{50}$$

We define $\mathrm{ac}(C, K) = \sup\{\mathrm{ac}_k(C, K) : k \in \mathbb{N}\}$.

Our first result shows that the approximate Caratheodory constant is bounded in terms of hereditary discrepancy, and, therefore, in terms of the vector balancing constant. Since our proof is algorithmic, together with Theorem 2 it also gives us a polynomial time algorithm to compute the vectors $v_1, \dots, v_k$ whose average is close to z.

Theorem 37. *For any convex body $C \subset \mathbb{R}^n$, and any symmetric convex body $K \subset \mathbb{R}^n$, we have*

$$\mathrm{ac}(C, K) \le \mathrm{vb}(C - C, K).$$

Moreover, there exists a randomized polynomial time algorithm that, given $u_1, \dots, u_N \in C$, and z in their convex hull, computes vectors $v_1, \dots, v_k \in \{u_1, \dots, u_N\}$ such that (50) holds with $a = O(\log n \,\mathrm{vb}(C - C, K)$.

Our proof of Theorem 37 is based on the following lemma, which is a slight extension of a result by Lovász, Spencer, and Vesztergombi [LSV86].

Lemma 38. *For any symmetric convex body $K \subset \mathbb{R}^n$, any $u_1, \dots, u_N \in \mathbb{R}^n$, and any $w \in [0,1]^N$, there exists an $x \in \{0,1\}^N$ such that $\sum_{i=1}^N x_i \le \sum_{i=1}^N w_i$, and*

$$\left\| \sum_{i=1}^{n} (w_i - x_i) u_i \right\|_K \le \mathrm{hd}((u_i)_{i=1}^N, K).$$

Proof. We first show that the lemma holds for every $w \in \{0, \frac{1}{2}, 1\}^N$ with a better constant. Specifically, we show that for any $w \in \{0, \frac{1}{2}, 1\}^N$, there exists an $x \in \{0,1\}^N$ such that

$$\left\| \sum_{i=1}^{n} (w_i - x_i) u_i \right\|_K \le \frac{1}{2}\,\mathrm{hd}((u_i)_{i=1}^N, K),$$

and $\sum_{i=1}^N x_i \le \sum_{i=1}^N w_i$. Towards this goal, let us define $S = \{i \in [n] : w_i = \frac{1}{2}\}$. By the definition of hereditary discrepancy, there exists a vector of signs $(\varepsilon_i)_{i\in S} \in \{-1,+1\}^S$ such that

$$\left\|\sum_{i\in S} \varepsilon_i u_i\right\|_K \le \mathrm{hd}((u_i)_{i=1}^N, K).$$

Let us define $x' \in \{-1,0,1\}^N$ by

$$x_i' = \begin{cases} \varepsilon_i & i \in S \\ 0 & i \notin S \end{cases}.$$

Observe, first, that $w + \frac{1}{2}x'$ and $w - \frac{1}{2}x'$ both belong to $\{0,1\}^N$, and, second, that

$$\min\left\{\sum_{i=1}^N w_i + \frac{1}{2}x_i', \sum_{i=1}^N w_i - \frac{1}{2}x_i'\right\} \le \sum_{i=1}^N w_i.$$

We can then take the vector $x \in \{w + \frac{1}{2}x', w - \frac{1}{2}x'\}$ that achieves the minimum on the right hand side above, and we have

$$\left\|\sum_{i=1}^n (w_i - x_i)u_i\right\|_K = \frac{1}{2}\left\|\sum_{i\in S} \varepsilon_i u_i\right\|_K \le \frac{1}{2}\,\mathrm{hd}((u_i)_{i=1}^N, K),$$

as desired.

To finish the proof of the lemma, it suffices to prove it for w such that w_i has a finite binary expansion for each i, i.e. $w_i = b_1 2^{-1} + \ldots + b_k 2^{-k}$ for some finite k and $b_1,\ldots,b_k \in \{0,1\}$. We do so by induction. Let us assume that the lemma is proved for all $w \in 2^{-k}\mathbb{Z}^N \cap [0,1)^N$ for some $k \ge 1$. (Note that the case $k = 1$ was already proved above.) We will show that it then also holds for every $w \in 2^{-k-1}\mathbb{Z}^N \cap [0,1)^N$. Fix such a w, and write it as $w = w' + \frac{1}{2}w''$, where $w' \in \{0,\frac{1}{2}\}^N$ and $w'' \in 2^{-k}\mathbb{Z}^N \cap [0,1)^N$. By the induction hypothesis, there exists an $x'' \in \{0,1\}^N$ such that

$$\left\|\sum_{i=1}^n (w_i'' - x_i'')u_i\right\|_K \le \mathrm{hd}((u_i)_{i=1}^N, K),$$

and $\sum_{i=1}^N x_i'' \le \sum_{i=1}^N w_i''$. We have $w' + \frac{1}{2}x'' \in \{0,\frac{1}{2},1\}$, and, as we have already shown above, there exists an $x \in \{0,1\}^N$ such that

$$\left\|\sum_{i=1}^n (w_i' + \frac{1}{2}x_i'' - x_i)u_i\right\|_K \le \frac{1}{2}\,\mathrm{hd}((u_i)_{i=1}^N, K), \tag{51}$$

and

$$\sum_{i=1}^N x_i \le \sum_{i=1}^N w_i' + \frac{1}{2}\sum_{i=1}^N x_i'' \le \sum_{i=1}^N w_i' + \frac{1}{2}\sum_{i=1}^N w_i'' = \sum_{i=1}^N w_i.$$

The inequality (51) and the triangle inequality then imply

$$\begin{aligned}\left\|\sum_{i=1}^n (w_i - x_i)u_i\right\|_K &= \left\|\sum_{i=1}^n (w_i' + \frac{1}{2}w_i'' - x_i)u_i\right\|_K \\ &\le \left\|\sum_{i=1}^n (w_i' + \frac{1}{2}x_i'' - x_i)u_i\right\|_K + \frac{1}{2}\left\|\sum_{i=1}^n (w_i'' - x_i'')u_i\right\|_K \le \mathrm{hd}((u_i)_{i=1}^N, K).\end{aligned}$$

This completes the proof of the lemma. $\square$

Proof of Theorem 37. Let us fix some $k \in \mathbb{N}$, vectors $u_0, u_1, \ldots, u_N$ and $z \in \mathrm{conv}\{u_0, \ldots, u_N\}$. (We index the vectors from 0 for notational convenience.) Define the new vectors $u'_1, \ldots u'_N$ by $u'_i = u_i - u_0$. We have that $z - u_0 \in \mathrm{conv}\{u'_1, \ldots, u'_N\}$, and we can write $z - u_0 = \sum_{i=1}^N \lambda_i u'_i$ for some $\lambda_1, \ldots, \lambda_N \geq 0$ such that $\sum_{i=1}^N \lambda_i = 1$. Let w be the fractional part of the vector $k\lambda$, i.e. $w_i = k\lambda_i - \lfloor k\lambda_i \rfloor$. By Lemma 38, there exists a vector $x \in \{0,1\}^N$ such that $\sum_{i=1}^N x_i \leq \sum_{i=1}^N w_i$, and

$$\left\| \sum_{i=1}^n (w_i - x_i) u'_i \right\|_K \leq \mathrm{hd}((u'_i)_{i=1}^N, K).$$

Let us define vectors $v'_1, \ldots, v'_\ell \in \{u'_1, \ldots, u'_N\}$ by taking $\lfloor k\lambda_i \rfloor + x_i$ copies of u'_i. Clearly,

$$\ell = \sum_{i=1}^N \lfloor k\lambda_i \rfloor + x_i \leq \sum_{i=1}^N \lfloor k\lambda_i \rfloor + w_i = k \sum_{i=1}^N \lambda_i = k.$$

We then define $v_1, \ldots, v_k \in \{u_0, \ldots, u_N\}$ by taking $k - \ell$ copies of u_0, and also taking the vectors $v'_i + u_0$ for every $i \in [\ell]$. Observe that

$$\begin{aligned} \left\| z - \frac{1}{k} \sum_{i=1}^k v_i \right\|_K &= \left\| z - u_0 - \frac{1}{k} \sum_{i=1}^\ell v'_i \right\|_K \\ &= \frac{1}{k} \left\| \sum_{i=1}^n (w_i - x_i) u'_i \right\|_K \leq \mathrm{hd}((u'_i)_{i=1}^N, K) \leq \mathrm{vb}(C - C, K). \end{aligned}$$

This completes the proof. □

References

[AAGM15] Shiri Artstein-Avidan, Apostolos Giannopoulos, and Vitali D. Milman. *Asymptotic geometric analysis. Part I*, volume 202 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 2015.

[Bal88] Keith Ball. Logarithmically concave functions and sections of convex sets in $\mathbf{R}^n$. *Studia Math.*, 88(1):69–84, 1988.

[Bal89] Keith Ball. Volumes of sections of cubes and related problems. In *Geometric aspects of functional analysis (1987–88)*, volume 1376 of *Lecture Notes in Math.*, pages 251–260. Springer, Berlin, 1989.

[Ban93] Wojciech Banaszczyk. Balancing vectors and convex bodies. *Studia Math.*, 106(1):93–100, 1993.

[Ban98] W. Banaszczyk. Balancing vectors and gaussian measures of n-dimensional convex bodies. *Random Structures & Algorithms*, 12(4):351–360, 1998.

[Ban10] Nikhil Bansal. Constructive algorithms for discrepancy minimization. In *51st Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2010*, pages 3–10. IEEE, 2010.

[BDG16] Nikhil Bansal, Daniel Dadush, and Shashwat Garg. An algorithm for Komlós conjecture matching Banaszczyk's bound. In *Proceedings of the IEEE 57th Annual Symposium on Foundations of Computer Science — FOCS 2016*, pages 788–799. IEEE, 2016.

[BDGL18] Nikhil Bansal, Daniel Dadush, Shashwat Garg, and Shachar Lovett. The Gram-Schmidt Walk: a cure for the Banaszczyk blues. In *To appear in STOC 2018*, 2018.

[BF81] J. Beck and T. Fiala. Integer-making theorems. *Discrete Applied Mathematics*, 3(1):1–8, 1981.

[BG81] I. Bárány and VS Grinberg. On some combinatorial questions in finite-dimensional spaces. *Linear Algebra and its Applications*, 41:1–9, 1981.

[BLNR15] Alexandre Belloni, Tengyuan Liang, Hariharan Narayanan, and Alexander Rakhlin. Escaping the local minima via simulated annealing: Optimization of approximately convex functions. In Peter Grünwald, Elad Hazan, and Satyen Kale, editors, *Proceedings of The 28th Conference on Learning Theory, COLT 2015, Paris, France, July 3-6, 2015*, volume 40 of *JMLR Workshop and Conference Proceedings*, pages 240–265. JMLR.org, 2015.

[BN16] Nikhil Bansal and Viswanath Nagarajan. Approximation-friendly discrepancy rounding. In *IPCO*, pages 375–386, 2016.

[BT87] J. Bourgain and L. Tzafriri. Invertibility of large submatrices with applications to the geometry of banach spaces and harmonic analysis. *Israel journal of mathematics*, 57(2):137–224, 1987.

[BV04] Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge University Press, Cambridge, 2004.

[Cha91] Bernard Chazelle. *The Discrepancy Method*. Cambridge University Press, 1991.

[CNN11] Moses Charikar, Alantha Newman, and Aleksandar Nikolov. Tight hardness results for minimizing discrepancy. In *SODA*, pages 1607–1614, 2011.

[DGLN16] Daniel Dadush, Shashwat Garg, Shachar Lovett, and Aleksandar Nikolov. Towards a constructive version of Banaszczyk's vector balancing theorem. In *Proceedings of the 20th International Workshop on Randomization and Computation — RANDOM 2016*, volume 60 of *LIPIcs. Leibniz Int. Proc. Inform.*, pages Art. No. 28, 12. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2016.

[ES14] Ronen Eldan and Mohit Singh. Efficient algorithms for discrepancy minimization in convex sets. *CoRR*, abs/1409.2913, 2014.

[GH62] Alain Ghouila-Houri. Caractérisation des matrices totalement unimodulaires. *C. R. Acad. Sci. Paris*, 254:1192–1194, 1962.

[Gia97a] Apostolos Giannopoulos. On some vector balancing problems. *Studia Mathematica*, 122(3):225–234, 1997.

[Gia97b] Apostolos A Giannopoulos. On some vector balancing problems. *Studia Mathematica*, 122(3):225–234, 1997.

[GLS88] M. Grötschel, L. Lovász, and A. Schrijver. *Geometric Algorithms and Combinatorial Optimization*. Springer-Verlag, 1988.

[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.

[HR17] Rebecca Hoberg and Thomas Rothvoss. A logarithmic additive integrality gap for bin packing. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2616–2625. SIAM, Philadelphia, PA, 2017.

[Kha96] Leonid G. Khachiyan. Rounding of polytopes in the real number model of computation. *Math. Oper. Res.*, 21(2):307–320, 1996.

[Lar17] Kasper Green Larsen. Constructive discrepancy minimization with hereditary l2 guarantees. Arxiv report 1711.02860, 2017.

[Lew95] A. S. Lewis. The convex analysis of unitarily invariant matrix functions. *J. Convex Anal.*, 2(1-2):173–183, 1995.

[LM12] S. Lovett and R. Meka. Constructive discrepancy minimization by walking on the edges. *Arxiv preprint arXiv:1203.5747*, 2012.

[LSV86] L. Lovász, J. Spencer, and K. Vesztergombi. Discrepancy of set-systems and matrices. *European Journal of Combinatorics*, 7(2):151–160, 1986.

[LT11] Michel Ledoux and Michel Talagrand. *Probability in Banach spaces*. Classics in Mathematics. Springer-Verlag, Berlin, 2011. Isoperimetry and processes, Reprint of the 1991 edition.

[Mat99] Jiří Matoušek. *Geometric Discrepancy (An Illustrated Guide)*. Springer, 1999.

[Mat11] Jiří Matoušek. The determinant bound for discrepancy is almost tight. http://arxiv.org/abs/1101.0767, 2011.

[Mil86] Vitali D. Milman. Inégalité de Brunn-Minkowski inverse et applications à la théorie locale des espaces normés. *C. R. Acad. Sci. Paris Sér. I Math.*, 302(1):25–28, 1986.

[MNT18] Jiří Matoušek, Aleksandar Nikolov, and Kunal Talwar. Factorization norms and hereditary discrepancy. *International Mathematics Research Notices*, page rny033, 2018.

[MP76] Bernard Maurey and Gilles Pisier. Séries de variables aléatoires vectorielles indépendantes et propriétés géométriques des espaces de Banach. *Studia Math.*, 58(1):45–90, 1976.

[Nik15] Aleksandar Nikolov. Randomized rounding for the largest simplex problem [extended abstract]. In *STOC'15—Proceedings of the 2015 ACM Symposium on Theory of Computing*, pages 861–870. ACM, New York, 2015.

[NNN12] Alantha Newman, Ofer Neiman, and Aleksandar Nikolov. Beck's three permutations conjecture: a counterexample and some consequences. In *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science—FOCS 2012*, pages 253–262. IEEE Computer Soc., Los Alamitos, CA, 2012.

[NT15] Aleksandar Nikolov and Kunal Talwar. Approximating hereditary discrepancy via small width ellipsoids. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 324–336. SIAM, Philadelphia, PA, 2015.

[Pis80] G. Pisier. Un théorème sur les opérateurs linéaires entre espaces de Banach qui se factorisent par un espace de Hilbert. *Ann. Sci. École Norm. Sup. (4)*, 13(1):23–43, 1980.

[Pis89] Gilles Pisier. *The volume of convex bodies and Banach space geometry*, volume 94 of *Cambridge Tracts in Mathematics*. Cambridge University Press, Cambridge, 1989.

[PTJ85] Alain Pajor and Nicole Tomczak-Jaegermann. Remarques sur les nombres d'entropie d'un opérateur et de son transposé. *C. R. Acad. Sci. Paris Sér. I Math.*, 301(15):743–746, 1985.

[Roc70] R. Tyrrell Rockafellar. *Convex analysis*. Princeton Mathematical Series, No. 28. Princeton University Press, Princeton, N.J., 1970.

[Rot12] Thomas Rothvoss. The entropy rounding method in approximation algorithms. In *Symposium on Discrete Algorithms (SODA)*, pages 356–372, 2012.

[Rot14] Thomas Rothvoß. Constructive discrepancy minimization for convex sets. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 140–145. IEEE Computer Society, 2014.

[Roy14] Thomas Royen. A simple proof of the Gaussian correlation conjecture extended to some multivariate gamma distributions. *Far East J. Theor. Stat.*, 48(2):139–145, 2014.

[Spe85] Joel Spencer. Six standard deviations suffice. *Trans. Amer. Math. Soc.*, 289:679–706, 1985.

[Spe94] Joel Spencer. *Ten lectures on the probabilistic method*, volume 64 of *CBMS-NSF Regional Conference Series in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, second edition, 1994.

[SS10] D.A. Spielman and N. Srivastava. An elementary proof of the restricted invertibility theorem. *Israel Journal of Mathematics*, pages 1–9, 2010.

[TJ89] Nicole Tomczak-Jaegermann. *Banach-Mazur distances and finite-dimensional operator ideals*, volume 38 of *Pitman Monographs and Surveys in Pure and Applied Mathematics*. Longman Scientific & Technical, Harlow; copublished in the United States with John Wiley & Sons, Inc., New York, 1989.