Report - Analytical Methods in TCS Constructive Discrepancy Minimization for Convex Sets

Matheus V. X. Ferreira Corey Sinnamon

May 2019

1 Introduction

This report relates the main results and methods of *Constructive Discrepancy Minimization for Convex Sets* [6] by Thomas Rothvoß, a beautiful and largely self-contained paper in algorithmic discrepancy theory. This paper ties together or extends many previous results, including those of Gluskin [4], Giannopoulos [3], Bansal [1], and Lovett and Meka [5].

The primary motivation and application for this work is in solving the discrepancy minimization problem optimally, up to a constant factor. The problem is as follows: Given sets $S_1, \ldots, S_n \subseteq [n]$, find a *discrepancy function* $\chi \colon [n] \to \{-1, 1\}$ that minimizes the *discrepancy*,

$$\max_{j \in [n]} |\chi(S_j)|$$

where $\chi(S)$ denotes the sum $\sum_{i \in S} \chi(i)$.

It was shown by Spencer [8] in 1985 that any instance of the problem has a discrepancy function achieving $\max_{j \in [n]} |\chi(S_j)| \leq 6\sqrt{n}$, and it is well known that this is optimal in the worst case, up to a small constant factor. However, Spencer's proof does not yield an efficient algorithm to find such a discrepancy function - indeed, his proof involves the pigeonhole principle on exponentially many items.

Rothvoß's paper gives an elegant method to construct a discrepancy function with $O(\sqrt{n})$ discrepancy in polynomial time. The basic algorithm that achieves this is almost embarrassingly simple, but the analysis requires some real insight, and there are some details that need to be looked after. We shall also see that the analysis, in its current state, does not yield a constant anywhere close to Spencer's existential guarantee.

2 Preliminaries

A half-space. Given $a \in R^n$, $\lambda \in R$, a n-dimensional half-space is a set of the form $H := \{x | \langle a, x \rangle \leq \lambda\}$.

A strip. Given $\lambda \in R^n$, $b \in R$, an n-dimensional strip is a set of the form $S := \{x | |\langle a, x\rangle| \leq b\}$.

Given a set $S \subseteq \mathbb{R}^n$, we define the distance from a point $y \in \mathbb{R}^n$ to S as $d(x, S) = \min_{y \in S} ||x - y||_2$. We say $y \in S$ is the projection of x to S if $d(x, S) = ||x - y||_2$. In other words, y is the point in S closest to x. Let $S_\delta = \{x \in \mathbb{R}^n | d(x, S) \leq \delta\}$ be the set of points at distance at most δ from S.

Gaussian Measure. Let $N^n(0, 1)$ be the n-dimensional standard Gaussian distribution. For all measurable set S,

$$\gamma_n(S) = Pr_{x \leftarrow N(0,1)^n}[x \in S]$$

The definition of Gaussian Measure allow us to compute the probability that a random n-dimensional Gaussian point is at distance at most δ from S:

$$Pr_{x \leftarrow N^n(0,1)}[d(x, S) \leq \delta] = \gamma_n(S_\delta)$$

Let us now cover the ingredients that will be used in proving the main theorem of Rothvoß's work. The proof is a consequence of three essential lemmas, here listed as Lemma 2.1, Lemma 2.2, and Lemma 2.3.

Lemma 2.1 (Gaussian Isoperimetric Inequality)**.** *Assume $K \subseteq R^n$ is measurable set and H is a half-space such that $\gamma_n(K) = \gamma_n(H)$, then*

$$\forall \delta \geq 0, \gamma_n(K_\delta) \geq \gamma_n(H_\delta)$$

Lemma 2.2. *For all $\epsilon > 0$, if K is a measurable set and $\gamma_n(K) \geq e^{-\epsilon n}$, then $\gamma_n(K_{3\sqrt{\epsilon n}}) \geq 1 - e^{-\epsilon n}$.*

Proof. For a half-plane $H = \{x | x_1 \leq \lambda\}$, pick λ such that $\gamma_n(H) = \gamma_n(K)$. If $e^{-\epsilon n} \geq 1/2$, then the conclusion is trivially true. If $e^{-\epsilon n} < 1/2$, we have $\lambda < 0$. Suppose for contradiction, $\lambda < -3\sqrt{\epsilon n}/2$, then

$$\begin{aligned}
\gamma_n(H) &= \int_{-\infty}^{\lambda} \frac{e^{-x^2/2}}{\sqrt{2\pi}} dx < \int_{-\infty}^{3\sqrt{\epsilon n}/2} \frac{e^{-x^2/2}}{\sqrt{2\pi}} dx \\
&\leq e^{-9\epsilon n/4} \\
&< e^{-\epsilon n} \Rightarrow\Leftarrow
\end{aligned}$$

where we use the fact for all $t \geq 0$, $\int_t^\infty \frac{e^{-x^2/2}}{\sqrt{2\pi}} dx \leq e^{-t^2/2}$. We have $\lambda + 3\sqrt{\epsilon n} \geq 3\sqrt{\epsilon n}/2$, and by symmetry $H_{3\sqrt{\epsilon n}} \geq 1 - e^{-\epsilon n}$ which by the Gaussian Isoperimetric Inequality implies $\gamma_n(K_{3\sqrt{\epsilon n}}) \geq 1 - e^{-\epsilon n}$. □

Lemma 2.3 (Gaussian Correlation Inequality)**.** *Let $K \subseteq \mathbb{R}^n$, $S \subseteq \mathbb{R}^n$ be symmetric convex bodies, then*

$$\gamma_n(K \cap S) \geq \gamma_n(K)\gamma_n(S)$$

The Gaussian Correlation Inequality was proved by Royen [7] in 2014, after being open for over four decades. We shall only need (and shall only prove) a simple version of the inequality, wherein S is a strip.[1]

Proof assuming S is a strip. Write $S = \{x \in \mathbb{R}^n \mid |\langle x, z\rangle| \leq \lambda\}$, where z is a unit vector in $\mathbb{R}^n$ and $\lambda \geq 0$. Without loss of generality, we may assume $z = e_n$, so that

$$S = \{x \in \mathbb{R}^n \mid |x_n| \leq \lambda\}$$

First, notice that the lemma is trivial when $n = 1$ as both K and S must be strips, and so $K \cap S$ is just the smaller of the two sets.

Now let

$$K_\alpha = \{(x_1, x_2, \ldots, x_{n-1}) \in \mathbb{R}^{n-1} \mid (x_1, x_2, \ldots, x_{n-1}, \alpha) \in K\}$$

and define

$$I_s(K) = \{\alpha \in \mathbb{R} \mid \gamma_{n-1}(K_\alpha) \geq s\}.$$

Then the measure of $K \cap S$ can be formulated as

$$\gamma_n(K \cap S) = \int_0^\infty \gamma_1(I_s(K) \cap [-\lambda, \lambda])ds.$$

Since the lemma holds in one dimension, we have

$$\int_0^\infty \gamma_1(I_s(K) \cap [-\lambda, \lambda])ds \geq \left(\int_0^\infty \gamma_1(I_s(K))ds\right)\gamma_1([-\lambda, \lambda]) = \gamma_n(K)\gamma_n(S)$$

as required. $\square$

3 The Main Theorem

Now we are prepared to describe and analyze Rothvoß's algorithm, which chooses a point in a sufficiently large symmetric convex body $K \subseteq \mathbb{R}^n$ such that the point also lies the hypercube $[-1, 1]^n$ and, with high probability, has many entries in $\{-1, 1\}$.

The algorithm consists of selecting a symmetric convex set K (i.e $x \in K$, then $-x \in K$) with large enough Gaussian measure. We first sample a random point x^* from the n-dimensional Gaussian distribution and project x^* to $K \cap [-1, 1]^n$. The main theorem states that this algorithm outputs a point where a constant fraction of the entries are either -1 or 1 with high probability. More formally:

- Sample $x^* \leftarrow N^n(0, 1)$.
- Output $y^* = \arg\min_{y \in K \cap [-1,1]^n} ||x^* - y||_2$.

Theorem 1. *For all $0 < \epsilon < \frac{1}{9000}$, $\delta = 3/2\epsilon \log 1/\epsilon$, if $K \subseteq \mathbb{R}^n$ is a symmetric convex set with Gaussian measure at least $e^{-\delta n}$, then with probability at least $1 - e^{-\Omega(n)}$, at least ϵn of the coordinates of y^* are $\{-1, 1\}$.*

[1] This proof is based on that given in [3].

Proof. Let $I^* = \{i \in [n] | y_i^* \in \{-1,1\}\}$. In the first step of the proof, we show that $Pr[d(x^*, K \cap [-1,1]^n) \geq \sqrt{n}/5] \geq 1 - e^{-\Omega(n)}$. We then argue that $Pr[d(x^*, K \cap [-1,1]^n) < \sqrt{n}/5 || I^*| < \epsilon n] \geq 1 - e^{-\Omega(n)}$. By Bayes' rule,

$$\begin{aligned} Pr[|I^*| < \epsilon n] &\leq \frac{Pr[d(x^*, K \cap [-1,1]^n) < \sqrt{n}/5]}{Pr[d(x^*, K \cap [-1,1]^n) < \sqrt{n}/5 || I^*| < \epsilon n]} \\ &\leq e^{-\Omega(n)} \end{aligned}$$

which allows us to conclude that $|I^*| \geq \epsilon n$ with high probability.

We first show that $d(x^*, K \cap [-1,1]^n) \geq \sqrt{n}/5$ with high probability. For all $i \in [n]$, $Pr_{x \leftarrow N^n(0,1)}[|x_i^*| \geq 2] = 2\int_2^\infty \frac{e^{-x^2/2}}{\sqrt{2\pi}} dx > \frac{1}{25}$ by the 1-dimensional density function for the standard Gaussian distribution. We have,

$$\begin{aligned} d(x^*, K \cap [-1,1]^n) &= \min_{y \in K \cap [-1,1]^n} ||x^* - y||_2 \\ &\geq \sqrt{\sum_{i=1}^n (x_i^* - y_i)^2}, \text{ for some } y \in K \cap [-1,1]^n \\ &\geq \sqrt{\sum_{i=1}^n \mathbb{I}[|x_i^*| \geq 2](2-1)}, \text{ by the fact } y \in [-1,1]^n \text{ and } |x_i^*| \geq 2 \end{aligned}$$

Next, we apply Chernoff bound, Lemma 6.1 on $\sum_{i=1}^n \mathbb{I}[|x_i^*| \geq 2]$. Because $\mu = E[\sum_{i=1}^n \mathbb{I}[|x_i^*| \geq 2]] > n/25$, there is a constant $\delta > 0$ such that $(1-\delta)\mu = n/25$.

We conclude that with probability at least $1 - e^{-\Omega(n)}$, $\sum_{i=1}^n \mathbb{I}[|x_i| \geq 2] \geq n/25$. It follows,

$$\begin{aligned} d(x^*, K \cap [-1,1]^n) &\geq \sqrt{n/25} \\ &= \sqrt{n}/5 \end{aligned}$$

Let's now bound the probability that $d(x^*, K \cap [-1,1]^n) < \sqrt{n}/5$ when $|I^*| < \epsilon n$. Let $K(I^*) := K \cap \{x \in R^n | \forall i \in I^*, |x_i| \leq 1\}$. We will argue that for all x, $d(x, K \cap [-1,1]^n) = d(x, K(I^*))$. This follows by the fact that for all $x \in \mathbb{R}^n$, $d(x, \cdot)$ is a minimization problem of a strictly convex function over a convex domain.

Strictly convex. A function $f : \mathbb{R}^n \to \mathbb{R}$ is strictly convex if for all $x \in R^n$, $y \in R^n$, $x \neq y$, and for all $\alpha \in (0,1)$, $f(\alpha x + (1-\alpha)y) < \alpha f(x) + (1-\alpha) f(y)$.

Lemma 3.1. *Let P and Q be convex sets. Let $x^* = \arg\min_{x \in P \cap Q} f(x)$. Assume x^* is in the interior of P. If $f : \mathbb{R}^n \to \mathbb{R}$ is a strictly convex function, then x^* minimizes $f(x)$ subject $x \in Q$.*

Proof. Suppose for contradiction there is $y^* \in Q$ such that $f(y^*) < f(x^*)$. For all $\alpha \in (0, 1)$,

$$\begin{aligned} f(\alpha x^* + (1-\alpha)y^*) &< \alpha f(x^*) + (1-\alpha) f(y^*) \\ &\leq \alpha f(x^*) + (1-\alpha) f(x^*) \\ &= f(x^*) \end{aligned}$$

Because x^* is in the interior of P, there exists $\alpha \in (0, 1)$ such that $\alpha x^* + (1 - \alpha)y^* \in P \cap Q$, contradicting x^* minimizes $f(x)$ subject to $x \in P \cap Q$. $\square$

To apply this Lemma, observe the ℓ_2 norm is strictly convex, Lemma 6.2. Let $P = \{x \in R^n || x_i| \leq 1, \forall i \notin I^*\}$, and let $Q = K(I^*)$. Observe that $P \cap Q = K \cap [-1, 1]^n$ and if x^* is the projection of a point x to $P \cap Q$, x^* must be in the interior of P since all points in the boundary of P are not points in Q and vice-versa. It follows by Lemma 3.1, that $d(x, K \cap [-1, 1]^n) = d(x, K(I^*))$.

By the Gaussian correlation inequality, for an index set $|I| < \epsilon n$,

$$\begin{aligned} \gamma_n(K(I)) &\geq \gamma_n(K) \prod_{i \in I} \gamma_n(\{x \in \mathbb{R}^n || x_i| \leq 1\}) \\ &\geq e^{-\delta n} \gamma_n(K) \prod_{i \in I} \gamma_n(\{x \in \mathbb{R}^n || x_i| \leq 1\}), \text{ by assumption } \gamma(K) \geq e^{-\delta n} \\ &\geq e^{-\delta n} e^{-1/2|I|}, \text{ for 1-dimensional strip of width 2, } \gamma_n(S) \geq e^{-1/2} \\ &\geq e^{-2\delta n}, \text{ by the fact } \epsilon \leq \delta \end{aligned}$$

By Lemma 2.2, $\gamma_n(K(I)_{3\sqrt{2\delta n}}) \geq 1 - e^{-2\delta n}$. Let $B := \cap_{|I| \leq \epsilon n}(K(I)_{3\sqrt{2\delta n}}))$.

By Lemma 6.3, the number of sets of size at most ϵn is at most $2^{nh(\epsilon)} \leq e^{3/4\epsilon \log 1/\epsilon n} = e^{\delta n}$. It follows by the Gaussian Correlation Inequality that

$$\begin{aligned} \gamma_n(B) &= 1 - \gamma_n(\cup_{|I| \leq \epsilon n}(\mathbb{R}^n \backslash K(I_{3\sqrt{2\delta n}})) \\ &\geq 1 - \sum_{|I| \leq \epsilon n} \gamma_n(\mathbb{R}^n \backslash K(I_{3\sqrt{2\delta n}})), \text{ by Union bound} \\ &\geq 1 - e^{\delta n} e^{-2\delta n}, \text{ by the fact } |\{I || I| \leq \epsilon n\}| \leq e^{\delta n} \\ &= 1 - e^{-\delta n} \end{aligned}$$

This implies, for all $|I| \leq \epsilon n$, $\gamma_n(K(I)_{3\sqrt{2\delta n}}) \geq 1 - e^{-\delta n}$ and by definition

$$\begin{aligned} Pr[d(x, K(I)) \leq \sqrt{n}/5] &\geq Pr[d(x, K(I)) \leq 3\sqrt{2\delta n}] \\ &\geq 1 - e^{-\delta n}, \text{ by our choice of } \delta, \ 3\sqrt{2\delta n} \geq \sqrt{n}/5. \end{aligned}$$

In particular, when $I^* \leq \epsilon n$, we have

$$\begin{aligned} Pr[d(x, K \cap [-1, 1]^n) \leq \sqrt{n}/5 || I^*| \leq \epsilon n] &= Pr[d(x, K(I^*)) \leq \sqrt{n}/5 || I^*| \leq \epsilon n] \\ &\geq 1 - e^{-\delta n} \end{aligned}$$

Together with the fact $Pr[d(x, K \cap [-1, 1]^n) \geq \sqrt{n}/5] \geq 1 - e^{-\Omega(n)}$, implies $Pr[|I^*| \leq \epsilon n] \leq e^{-\Omega(n)}$ which concludes the proof. $\square$

4 Minimizing Discrepancy By Bootstrapping

In this section, we see how to apply Theorem 1 to the discrepancy problem defined in the introduction. If we view the discrepancy function as a colouring of the elements in $[n]$ by either -1 or 1, our goal is to construct a full colouring with low discrepancy; along the way, we shall critically deal with *partial colourings*, where some elements are assigned values in $\{-1,1\}$ and others have values in $(-1,1)$.

The algorithm of the last section quickly finds a partial colouring having low discrepancy and at least ϵn ± 1-valued entries (with high probability). To solve Spencer's problem effectively, we need a way to bootstrap Theorem 1 to extend a partial colouring $x \in [-1,1]^n$ to a new vector that still has low discrepancy, but has $\approx \epsilon n$ additional entries in $\{-1,1\}$.

4.1 Controlling Discrepancy

For the purposes of minimizing discrepancy, the symmetric convex body we are interested in has the form

$$K^* = \{x \in \mathbb{R}^n \mid \left|\sum_{i \in S_j} x_i\right| \leq \sqrt{n} \text{ for all } j \in \{1,\dots,n\}\}$$

More specifically, we are interested in *scalings* of K^*. The set cK^* contains all solutions $x \in \{-1,1\}^n$ such that the discrepancy function $\chi(i) = x_i$ has $\max_j \chi(S_j) = \max_j |\langle x, S_j\rangle| \leq c\sqrt{n}$, but it relaxes the discrepancy function to allow "fractional colourings", where x may have real coordinates.

Clearly K^* is symmetric and convex, as is cK^* for any $c \in \mathbb{R}$. In order to apply the techniques of Theorem 1 to cK^*, we need to ensure that $\gamma_n(cK^*)$ is sufficiently large.

Let $\operatorname{strip}(z, M)$ denote $\{x \in \mathbb{R}^n \mid |\langle x, z\rangle| \leq M\}$, and observe that

$$cK^* = \bigcap_{j=1}^{n} \operatorname{strip}(\mathbb{1}_{S_j}, c\sqrt{n}).$$

By the Gaussian Correlation inequality, $\gamma_n(cK^*) \geq \prod_{j=1}^n \gamma_n(\operatorname{strip}(\mathbb{1}_{S_j}, c\sqrt{n}))$. Since the "width" of $\operatorname{strip}(\mathbb{1}_{S_j}, c\sqrt{n})$ is $2\frac{c\sqrt{n}}{\sqrt{|S_j|}}$, by rotational symmetry we have

$$\gamma_n(\operatorname{strip}(\mathbb{1}_{S_j}, c\sqrt{n})) = \gamma_1\left(\left[-\frac{c\sqrt{n}}{\sqrt{|S_j|}}, \frac{c\sqrt{n}}{\sqrt{|S_j|}}\right]\right) \geq \gamma_1([-c,c])$$

Let us estimate this value (assuming $c \geq 1$) using Hoeffding's inequality:

$$\gamma_1([-c,c]) \geq 1 - 2\exp(-c^2/2) \geq \exp(-2\exp(-c^2/2)) \qquad (*)$$

Thus $\gamma_n(cK^*) \geq \prod_{j=1}^n \exp(-2\exp(-c^2/2)) = \exp(-2n\exp(-c^2/2))$.

Theorem 1 only requires that $\gamma_n(cK^*) \geq \exp(-\delta n)$, which is satisfied as long as

$$c \geq \sqrt{2 \ln \frac{1}{2\delta}}$$

As $\delta < 1/500$, we can safely use any c larger than $\sqrt{2 \ln \frac{500}{2}} \geq 3.33$.

4.2 Extending a Partial Colouring

The first issue to overcome is how to use the algorithm to substantially extend an existing solution without changing any entries that are already in $\{-1, 1\}$. The following lemma accomplishes this, at the small cost that only $\frac{\epsilon}{2}n$ entries are made tight, rather than ϵn.

Lemma 4.1. *Let $z \in (-1,1)^n$. If $K \subseteq \mathbb{R}^n$ is a symmetric convex set with $\gamma_n(K) \geq \exp(-\delta n)$, then there is a polynomial time algorithm to find a point $y \in (z + K) \cap [-1,1]^n$ with at least $\frac{\epsilon}{2}n$ entries in $\{-1,1\}$.*

Proof. Define the linear map $F\colon \mathbb{R}^n \to \mathbb{R}^n$ by

$$F(e_i) = \frac{\operatorname{sgn}(z_i)}{1 - |z_i|} e_i$$

where e_i denotes the ith standard basis vector.

We stretch K according F: Since F is linear, $F(K)$ is still a symmetric convex body. Moreover, for all i, we have $\frac{1}{1-|z_i|} \geq 1$ since $0 \leq |z_i| < 1$. Thus $K \subseteq F(K)$, and in particular $\gamma_n(F(K)) \geq \gamma_n(K) \geq \exp(-\delta n)$.

Now apply the main algorithm to $F(K)$:

- Sample $x^* \leftarrow N^n(0,1)$.
- Output $y^* = \arg\min_{y \in F(K) \cap [-1,1]^n} ||x^* - y||_2$.

By Theorem 1, $y^* \in F(K) \cap [-1,1]^n$ and y^* has at least ϵn ± 1-valued entries with high probability (if this fails, then resample and repeat until y^* has ϵn tight entries).

Now let $y = z + F^{-1}(y^*)$ so that $y_i = z_i + \operatorname{sgn}(z_i)(1 - |z_i|)y_i^*$. Since $y_i^* \in [-1,1]$ and $z_i \in (-1,1)$ for all $i \in [n]$, we have $y \in [-1,1]^n$. Thus, the algorithm always finds $y \in (z + K) \cap [-1,1]^n$.

Moreover, if $y_i^* = 1$, then

$$y_i = z_i + \operatorname{sgn}(z_i)(1 - |z_i|)y_i^* = \operatorname{sgn}(z_i) \in \{-1, 1\}.$$

That is, whenever $y_i^* = 1$, the corresponding entry of y is tight (either 1 or -1). Since at least ϵn entries of y^* are either 1 or -1, we expect $\frac{\epsilon}{2}n$ entries of y to be in $\{-1,1\}$; however, it could be that most of the tight entries in y^* have the value -1, in which case the corresponding entries of y may not be tight.

To complete the algorithm, we add one more step: If fewer than $\frac{\epsilon}{2}n$ entries of y^* have the value 1, then replace y^* by $-y^*$. As K and $F(K)$ are symmetric, the difference between y^* and $-y^*$ is insignificant. Since we required y^* to have ϵn tight entries, either y^* or $-y^*$ must have $\frac{\epsilon}{2}n$ entries with value 1. After this operation, y will have at least $\frac{\epsilon}{2}n$ tight entries, as required. $\square$

We can also apply this lemma on a subspace of $\mathbb{R}^n$ in order to extend a vector with some tight entries and some nontight entries. For any subspace $U \subset \mathbb{R}^n$, let γ_U denote the Gaussian measure on U.

Corollary 4.1. *Given $y^{(t)} \in [-1,1]^n$, let $L = \{i \in [n] \mid y_i^{(t)} \notin \{-1,1\}\}$ and $U = \operatorname{span}(\{e_i \mid i \in L\})$; then U is the subspace of $\mathbb{R}^n$ on the unfixed entries of $y^{(t)}$.*

If $K \subseteq U$ is a symmetric convex set with $\gamma_U(K) \geq \exp(-\delta \dim(U))$, then there is a polynomial time algorithm to find a point $y^{(t+1)} \in (y^{(t)} + K) \cap [-1,1]^L$ such that $y_i \in \{-1,1\}$ for at least $\frac{\epsilon}{2}\dim(U)$ indices i in L.

4.3 Putting It Together

We are now prepared to use the partial colouring method to find a complete colouring.

Theorem 2. *Given $S_1, S_2, \ldots, S_n \subseteq [n]$, there is a polynomial time algorithm to find a point $y \in \{-1,1\}^n$ such that $|\sum_{i \in S_j} y_i| \leq C\sqrt{n}$ for all $j \in [n]$ for some absolute constant C.*

Proof. Beginning with $y^{(1)} = 0^n$, we shall find a sequence of partial colourings $y^{(2)}, y^{(3)}, \ldots$ that have more and more fixed coordinates, while each stays within cK^* for some constant c, and hence the discrepancy is always at most $c\sqrt{n}$. When the number of unfixed coordinates is sufficiently small, say $\log n$, then those coordinates can be set to ± 1 arbitrarily.

For all $t \geq 1$,

- let $L^{(t)}$ denote the set of unfixed coordinates of $y^{(t)}$:
$$L^{(t)} = \{i \mid y_i^{(t-1)} \in (-1,1)\}$$
- let $U^{(t)}$ denote the subspace of unfixed coordinates of $y^{(t-1)}$:
$$U^{(t)} = \operatorname{span}\{e_i \mid i \in L^{(t)})\}$$
- let $m^{(t)} = |L^{(t)}| = \dim(U^{(t)})$, and
- let $K^{(t)} = C^{(t)} K^* \cap U^{(t)}$, where
$$C^{(t)} = \sqrt{2\left(\ln\frac{1}{\delta} + \ln\frac{2n}{m^{(t)}}\right)\frac{m^{(t)}}{n}}$$

To obtain $y^{(t+1)}$ from $y^{(t)}$, we apply Corollary 4.1 using the symmetric convex body $K^{(t)} = C^{(t)} K^* \cap U^{(t)}$.

Let us first check that $K^{(t)}$ has sufficiently large measure in $U^{(t)}$ to apply the corollary: Observe that

$$K^{(t)} = \bigcap_{j \in [n]} \text{strip}\left(\frac{\mathbb{1}_{S_j}}{C^{(t)}\sqrt{n}}\right) \cap U^{(t)} = \bigcap_{j \in [n]} \text{strip}\left(\frac{\mathbb{1}_{S_j \cap L^{(t)}}}{C^{(t)}\sqrt{n}}\right)$$

Since $\|\mathbb{1}_{S_j \cap L^{(t)}}\|_2 \leq \|\mathbb{1}_{L^{(t)}}\|_2 = \sqrt{m^{(t)}}$, we have

$$\gamma_{U^{(t)}}\left(\text{strip}\left(\frac{\mathbb{1}_{S_j \cap L^{(t)}}}{C^{(t)}\sqrt{n}}\right)\right) \geq \gamma_1\left(\left[-C^{(t)}\sqrt{\frac{n}{m^{(t)}}}, C^{(t)}\sqrt{\frac{n}{m^{(t)}}}\right]\right) \geq \exp\left(-2\exp\left(-\frac{(C^{(t)})^2}{2}\frac{n}{m^{(t)}}\right)\right)$$

by $(*)$. Thus, by Lemma 2.3,

$$\begin{aligned}
\gamma_{U^{(t)}}(K^{(t)}) &\geq \exp\left(-2n\exp\left(-\frac{(C^{(t)})^2}{2}\frac{n}{m^{(t)}}\right)\right) \\
&= \exp\left(-2n\exp\left(-\frac{2\left(\ln\frac{1}{\delta} + \ln\frac{2n}{m^{(t)}}\right)\frac{m^{(t)}}{n}}{2}\frac{n}{m^{(t)}}\right)\right) \\
&= \exp\left(-2n\exp\left(-\ln\frac{1}{\delta} - \ln\frac{2n}{m^{(t)}}\right)\right) \\
&= \exp\left(-2n\delta\frac{m^{(t)}}{2n}\right) \\
&= \exp(-\delta m^{(t)}).
\end{aligned}$$

Thus, $\gamma_{U^{(t)}}(K^{(t)})$ is sufficiently large to apply Corollary 4.1.

In each step, the number of unfixed entries decreases by a factor of $\frac{\epsilon}{2}$; hence, after $T = O(\log n)$ steps, the number of unfixed entries of $y^{(T)}$ will be at most $\log n$. The $\log n$ remaining entries can then be assigned values in ± 1 arbitrarily; the resulting increase in discrepancy will be at most $2\log n$.

It only remains to bound the discrepancy of $y^{(T)}$ Since $y^{(t)} \in y^{(t-1)} + C^{(t)}K^*$ for all $t \geq 1$, it follows that

$$\begin{aligned}
y^{(T)} &\in y^{(T-1)} + C^{(T)}K^* \\
&\subseteq y^{(T-2)} + C^{(T-1)}K^* + C^{(T)}K^* \\
&\subseteq y^{(T-3)} + C^{(T-2)}K^* + C^{(T-1)}K^* + C^{(T)}K^* \\
&\vdots \\
&\subseteq y^{(0)} + \left(\sum_{h=1}^{T} C^{(h)}\right)K^* = \left(\sum_{h=1}^{T} C^{(h)}\right)K^*
\end{aligned}$$

Let us bound $\sum_{h=1}^{T} C^{(h)}$. Since $m^{(t+1)} \leq (1-\epsilon/2)\,m^{(t)}$ for all $t \geq 1$, we have

$m^{(t)} \leq (1-\epsilon/2)^{t-1} n$. Hence,

$$C^{(t)} \leq \sqrt{2\left(\ln\frac{1}{\delta} + 1 + (t-1)\ln\frac{1}{1-\epsilon/2}\right)(1-\epsilon/2)^{t-1}}$$

Notice that $\sum_{h=1}^{T} C^{(h)}$ converges to a constant because it decreases geometrically, approximately mirroring $\sqrt{1-\epsilon/2}^{t-1} \approx 0.9999^{t-1}$. Therefore, $y^{(T)}$ is in CK^* for some constant C (which depends on ϵ), and so the discrepancy of the final colouring is $O(\sqrt{n})$. □

In this report, we have attempted to manage the constants involved in the analysis in order to gauge the discrepancy of the final output. Unfortunately, if ϵ is as small as $\frac{1}{9000}$ (as required by Theorem 1), then $\sum_{h=1}^{T} C^{(h)}$ may be very large indeed - larger than 18000, and certainly not comparable with the constant in Spencer's bound, which is about 5.32. We could find an exact upper bound by carefully analyzing the sum in this last step, but it hardly seems worthwhile; it will be tens of thousands of times larger than 5.32.

5 A Few Words About Eldan-Singh

In 2014, a similar algorithm to find a point in a symmetric convex body having many ± 1 entries was proposed by Eldan and Singh [2][2]. Rothvoß's algorithm chooses a normally distributed point $x* \in \mathbb{R}^n$ and then returnes $y^* = \arg\min_{y \in K} \|x^* - y^*\|$; their approach was to choose $y* = \arg\max_{y\in K}\langle x^*, y\rangle$ instead. The theorem in this work corresponding to Theorem 1 is as follows.

Theorem 3 (Eldan and Singh, 2014)**.** *For any constant* $0 < \epsilon < \left(\frac{1-\sqrt{2/\pi}}{32}\right)^4$, *there exists a constant* $0 < \delta < 1$ *such that every symmetric convex body* $K \subseteq \mathbb{R}^n$ *with* $\gamma_n(K) \geq e^{-\epsilon n}$, *the point* $y^* = \arg max_{y\in K}\langle x^*, y\rangle$ *where* x^* *is a standard Gaussian in* $\mathbb{R}^n$, *satisfies* $\#\{i \in [n] \mid |x_i| = 1\} \geq \delta n$ *with probability at least* $\frac{1}{2}$.

We looked at this alternative approach in the hopes that it would improve the constants involved, or make them easier to analyze. Unfortunately, neither seems to be the case: Their range for ϵ is smaller than that in Theorem 1, as $\left(\frac{1-\sqrt{2/\pi}}{32}\right)^4 \approx 1.6 \cdot 10^{-9} \ll \frac{1}{9000}$. The value of δ in their analysis is eventually chosen to be the same as ϵ.The probability of success is also much worse ($\frac{1}{2}$ instead of $1 - \exp(-\Omega(n))$). Though this approach has advantages - one being that the algorithm involves optimizing a linear program rather than a SOS program - the constants appear even more difficult than those of Rothvoß.

[2]The paper also includes an approach to finding a partial colouring in a non-symmetric convex body, which is interesting in its own right. It is an open question if this approach can be bootstrapped to produce a full colouring.

References

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

[2] Ronen Eldan and Mohit Singh. Efficient algorithms for discrepancy minimization in convex sets. *arXiv preprint arXiv:1409.2913*, 2014.

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

[4] 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.

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

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

[7] Thomas Royen. A simple proof of the gaussian correlation conjecture extended to multivariate gamma distributions. *arXiv preprint arXiv:1408.1028*, 2014.

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

6 Appendix

Lemma 6.1 (Chernoff Bound)**.** *For independent random variables $x_1, ..., x_n$ taking values $\{0,1\}$, let $X = \sum_{i=1}^{n} x_i$, $\mu = E[X]$, then*

$$Pr[X \leq (1-\delta)\mu] \leq e^{-\delta^2 \mu/2}$$

Lemma 6.2. *The Euclidean norm is strictly convex.*

Proof. For all $x \in R^n$, $y \in R^n$, if x and y are not linearly dependent, then for all $\alpha \in (0,1)$,

$$||\alpha x + (1-\alpha)y||_2^2 = \alpha^2 ||x||_2^2 + (1-\alpha)^2 ||y||_2^2 + 2\alpha(1-\alpha)\sum_{i=1}^{n} x_i y_i$$

$$(\alpha||x||_2 + (1-\alpha)||y||_2)^2 = \alpha^2 ||x||_2^2 + (1-\alpha)^2 ||y||_2^2 + 2\alpha(1-\alpha)||x||_2||y||_2$$

Subtracting the equations,

$$\begin{aligned} ||\alpha x + (1-\alpha)y||_2^2 - (\alpha||x||_2 + (1-\alpha)||y||_2)^2 &= 2\alpha(1-\alpha)(\langle x, y\rangle - ||x||_2||y||_2) \\ &< 0, \text{ by Cauchy–Schwarz} \end{aligned}$$

where in the last inequality, we use the fact Cauchy-Schwarz is tight only when x and y are linearly dependent. When x and y are linearly dependent, it is easy to check that $||\alpha x + (1-\alpha)y||_2 = \alpha||x||_2 + (1-\alpha y)||y||_2$ iff $x = y$. □

Lemma 6.3. *Let $S = \cup_{i=1}^{\epsilon n} \binom{[n]}{i}$ be the collection of all subsets of size at most ϵn, then $|S| \leq 2^{h(\epsilon)n}$ where $h(\epsilon) = \epsilon \log(1/\epsilon) + (1-\epsilon)\log(1/(1-\epsilon))$ is the binary entropy.*

Proof. Let $X \subset [n]$, be a uniformly random subset of size at most ϵn. The entropy of X is given by:

$$\begin{aligned} H(X) &= \sum_{I \subseteq [n] || I| \leq \epsilon n} -Pr[X = I]\log(Pr[X = I]) \\ &= -|S|/|S|\log(1/|S|) \\ &= -\log(1/|S|) \end{aligned}$$

This implies $|S| = 2^{H(X)}$. Define the indicator random variable $X_i = \mathbb{I}[i \in X]$. By additivity of entropy, $H(X) \leq \sum_{i=1}^{n} H(X_i) = nh(\epsilon)$ which concludes the proof. □