

TIGHTER BOUNDS FOR THE DISCREPANCY OF BOXES AND POLYTOPES

ALEKSANDAR NIKOLOV

Dedicated to Klaus Roth

Abstract. Combinatorial discrepancy is a complexity measure of a collection of sets which quantifies how well the sets in the collection can be simultaneously balanced. More precisely, we are given an n -point set P , and a collection $\mathcal{F} = \{F_1, \dots, F_m\}$ of subsets of P , and our goal is color P with two colors, red and blue, so that the maximum over the F_i of the absolute difference between the number of red elements and the number of blue elements (the discrepancy) is minimized. Combinatorial discrepancy has many applications in mathematics and computer science, including constructions of uniformly distributed point sets, and lower bounds for data structures and private data analysis algorithms. We investigate the combinatorial discrepancy of geometrically defined systems, in which P is an n -point set in d -dimensional space, and \mathcal{F} is the collection of subsets of P induced by dilations and translations of a fixed convex polytope B . Such set systems include systems of sets induced by axis-aligned boxes, whose discrepancy is the subject of the well-known Tusnády problem. We prove new discrepancy upper and lower bounds for such set systems by extending the approach based on factorization norms previously used by the author, Matoušek, and Talwar. We also outline applications of our results to geometric discrepancy, data structure lower bounds, and differential privacy.

§1. *Introduction.* As usual, we define the combinatorial discrepancy of a system \mathcal{F} of subsets of some finite set P as:

$$\text{disc } \mathcal{F} := \min_{\chi: P \rightarrow \{-1, 1\}} \max_{F \in \mathcal{F}} \sum_{p \in F} \chi(p).$$

Matoušek [26] provides more background on combinatorial and geometric discrepancy. For a large number of applications of combinatorial discrepancy to geometric discrepancy, numerical methods, and computer science, see Chazelle [12]. A recent application to private data analysis is described in the author's PhD Thesis [29]. We give more details on applications to geometric discrepancy, numerical integration, data structures, and differential privacy in §2.

Let \mathcal{A}_d be the family of *anchored* axis-aligned boxes in \mathbb{R}^d : $\mathcal{A}_d := \{A(x) : x \in \mathbb{R}^d\}$, where $A(x) := \{y \in \mathbb{R}^d : y_i \leq x_i \text{ for all } i \in [d]\}$ and $[d] := \{1, \dots, d\}$. This is a slight abuse of terminology: $A(x)$ is not a box, but rather a polyhedral

Received 3 June 2017.

MSC (2010): 11K38 (primary), 52B11, 42B05 (secondary).

cone. Usually $A(x)$ is defined to be anchored at 0, i.e. it is defined as the set $\{y \in \mathbb{R}^d : 0 \leq y_i \leq x_i \text{ for all } i \in [d]\}$. However, we prefer to “anchor” $A(x)$ at infinity, since this convention does not affect any of the results in the paper, and allows us to avoid some minor technicalities.

For an n -point set $P \subset \mathbb{R}^d$, we denote by $\mathcal{A}_d(P) := \{A(x) \cap P : x \in \mathbb{R}^d\}$ the set system induced by anchored boxes on P . (Note that this set system is finite, and, in fact, can have at most n^d sets.) Tusnády’s problem asks for tight bounds on the largest possible combinatorial discrepancy of $\mathcal{A}_d(P)$ over all n -point sets P , i.e. for the order of growth of the function $\text{disc}(n, \mathcal{A}_d) := \sup\{\text{disc } \mathcal{A}_d(P) : P \subset \mathbb{R}^d, |P| = n\}$ with n . The best known upper bound for the Tusnády problem is $\text{disc}(n, \mathcal{A}_d) = O_d(\log^d n)$, and was recently proved by Bansal and Garg [6]. Their result improved on the prior work of Larsen [22] (see also the proof in [28]), who showed that $\text{disc}(n, \mathcal{A}_d) = O_d(\log^{d+1/2} n)$. Here, and in the rest of this paper, we use the notation $O_p(\cdot)$, $\Omega_p(\cdot)$, $\Theta_p(\cdot)$ to denote the fact that the implicit constant in the asymptotic notation depends on the parameter p .

In our first result, we improve the upper bounds above.

THEOREM 1. *For any $d \geq 2$,*

$$\text{disc}(n, \mathcal{A}_d) = O_d(\log^{d-1/2} n).$$

It was shown in [27, 28] that $\text{disc}(n, \mathcal{A}_d) = \Omega_d(\log^{d-1} n)$, so the upper bound above is within a $O_d(\sqrt{\log n})$ factor from the lower bound. This brings us tantalizingly close to a full resolution of Tusnády’s problem.

More generally, let \mathcal{F} be a collection of subsets of \mathbb{R}^d , and, for an n -point set $P \subset \mathbb{R}^d$, let $\mathcal{F}(P)$ be the set system induced by \mathcal{F} on P , i.e. $\mathcal{F}(P) := \{F \cap P : F \in \mathcal{F}\}$. We are interested in how the worst-case combinatorial discrepancy of such set systems grows with n . This is captured by the function $\text{disc}(n, \mathcal{F}) := \sup\{\text{disc } \mathcal{F}(P) : P \subset \mathbb{R}^d, |P| = n\}$.

Let $K \subseteq \mathbb{R}^d$, and let’s define \mathcal{T}_K to be the family of images of K under translations and homothetic transformations: $\mathcal{T}_K := \{tK + x : t \in \mathbb{R}_+, x \in \mathbb{R}^d\}$, where \mathbb{R}_+ is the set of positive reals. If we take $Q := [0, 1]^d$, then it is well known that $\text{disc}(n, \mathcal{T}_Q) \leq 2^d \text{disc}(n, \mathcal{A}_d)$ and, therefore, Theorem 1 implies $\text{disc}(n, \mathcal{T}_Q) = O_d(\log^{d-1/2} n)$. Our next result shows that this bound holds for any convex polytope B in \mathbb{R}^d .

THEOREM 2. *For any $d \geq 2$, and any closed convex polytope $B \subset \mathbb{R}^d$,*

$$\text{disc}(n, \mathcal{T}_B) = O_{d,B}(\log^{d-1/2} n).$$

With essentially the same proof we can establish the stronger fact that $\text{disc}(n, \text{POL}(\mathcal{H})) = O_{d,\mathcal{H}}(\log^{d-1/2} n)$, where \mathcal{H} is a family of hyperplanes in \mathbb{R}^d , and $\text{POL}(\mathcal{H})$ is the set of all polytopes which can be written as $\bigcap_{i=1}^m H_i$, with each H_i a halfspace whose boundary is parallel to some $h \in \mathcal{H}$. The best previously known upper bound in this setting is also due to Bansal and Garg [6], and is equal to $O_{d,\mathcal{H}}(\log^d n)$.

Our final result is a nearly matching lower bound for any generic convex polytope B . (See Definition 2 for the meaning of “generic”.)

THEOREM 3. *For any generic convex polytope with non-empty interior $B \subseteq \mathbb{R}^d$,*

$$\text{disc}(n, \mathcal{T}_B) = \Omega_{d,B}(\log^{d-1} n).$$

The best previously known lower bound on $\text{disc}(n, \mathcal{T}_B)$ was $\Omega_{d,B}(\log^{(d-1)/2} n)$, which follows from a result of Drmota [13] in geometric discrepancy theory. This lower bound holds for any convex polytope B , but is quadratically weaker than our lower bound.

We conjecture that the genericity condition in Theorem 3 is not necessary, and, furthermore, that the asymptotic constant in the lower bound need not depend on B . These conjectures would be implied by certain estimates on the Fourier spectrum of convex polytopes: see §6 for more details. By contrast, the asymptotic constant in the upper bound in Theorem 2 has to depend on B , because we can approximate the unit Euclidean ball D^d in \mathbb{R}^d arbitrarily well with a convex polytope, and $\text{disc}(n, \mathcal{T}_{D^d}) = \Omega_d(n^{1/2-1/(2d)})$ [2].

Our proofs combine the approach to combinatorial discrepancy based on factorization norms, used previously by the author, Matoušek and Talwar [28, 30], with several new ideas. In the proof of Theorem 1 we use a result from Banaszczyk [4] on balancing series of vectors, and in the proof of Theorem 2 we further use a decomposition of polytopes given by Matoušek [25]. The proof of Theorem 3 combines the factorization norms approach with the Fourier analytic method pioneered by Roth [33]. In §3.4 we discuss the techniques in more detail, after we have introduced the relevant background.

Together, these results give nearly tight estimates on the discrepancy function $\text{disc}(n, \mathcal{T}_B)$ for almost any convex polytope B . Perhaps surprisingly, they imply that the order of growth of $\text{disc}(n, \mathcal{T}_B)$ with n is essentially independent of the particular structure of B , at least when B is generic, and, moreover, that order of growth is achieved by the simplest possible example, a cube. In the next section we describe applications of our results to geometric discrepancy, quasi-Monte Carlo methods, data structure lower bounds, and differential privacy.

§2. Applications. In this section we outline several applications of our discrepancy upper and lower bounds.

2.1. Geometric discrepancy and numerical integration. Geometric discrepancy measures the irregularity of a distribution of n points in $[0, 1]^d$ with respect to a family of distinguishing sets. In particular, for an n -point set $P \subseteq [0, 1]^d$ and a family of measurable subsets \mathcal{F} of \mathbb{R}^d , we define the discrepancy

$$D(P, \mathcal{F}) := \sup_{F \in \mathcal{F}} | |P \cap F| - \lambda_d(F \cap [0, 1]^d) |,$$

where λ_d is the Lebesgue measure on \mathbb{R}^d . The smallest achievable discrepancy over n -point sets with respect to \mathcal{F} is denoted $D(n, \mathcal{F}) := \inf \{ D(P, \mathcal{F}) :$

$P \subset \mathbb{R}^d, |P| = n$. A famous result from Schmidt [34] shows that $D(n, \mathcal{A}_2) = \Theta(\log n)$. The picture is much less clear in higher dimensions. In a seminal paper [32], Roth showed that $D(n, \mathcal{A}_d) = \Omega_d(\log^{(d-1)/2} n)$ for any $d \geq 2$; the best known lower bound in $d \geq 3$ is given by Bilyk, Lacey, and Vagharshakyan [9] and is $D(n, \mathcal{A}_d) = \Omega_d(\log^{(d-1)/2 + \eta_d} n)$, where η_d is a positive constant depending on d and going to 0 as d goes to infinity. On the other hand, the best known upper bound is $D(n, \mathcal{A}_d) = O_d(\log^{d-1} n)$ and can be achieved in many different ways, one of the simplest being the Halton–Hammersley construction [19, 20]. Beck and Chen [8] call the problem of closing this significant gap “the Great Open Problem” (in geometric discrepancy theory). See Matoušek [26] for further background on geometric discrepancy.

There is a known connection between combinatorial and geometric discrepancy. Roughly speaking, combinatorial discrepancy is an upper bound on geometric discrepancy. More precisely, we have the following transference lemma, which goes back to the work of Beck on Tusnády’s problem [7] (see [26] for a proof).

LEMMA 4. *Let \mathcal{F} be a family of measurable sets in \mathbb{R}^d such that there is some $F \in \mathcal{F}$ which contains $[0, 1]^d$. Assume that $D(n, \mathcal{F})/n$ goes to 0 as n goes to infinity, and that $\text{disc}(n, \mathcal{F}) \leq f(n)$ for a function f that satisfies $f(2n) \leq (2 - \delta)f(n)$ for all n and some fixed $\delta > 0$. Then there exists a constant C_δ that only depends on δ , for which $D(n, \mathcal{F}) \leq C_\delta f(n)$.*

Lemma 4 and Theorem 2 imply that $D(n, \mathcal{T}_B) = O_{d,B}(\log^{d-1/2} n)$ for any convex polytope B in \mathbb{R}^d . This bound gets within an $O_{d,B}(\sqrt{\log n})$ factor from the best bound known for axis-aligned boxes in d dimensions. The tightest bound known prior to our work was $O_{d,B}(\log^d n)$ and was also implied by the best previously known upper bound on combinatorial discrepancy.

Geometric discrepancy can be defined with any Borel probability measure ν on $[0, 1]^d$ in place of the Lebesgue measure: let’s call the resulting quantities $D(P, \mathcal{F}, \nu)$ and $D(n, \mathcal{F}, \nu)$. It turns out that Lemma 4 holds with $D(n, \mathcal{F})$ replaced by $D(n, \mathcal{F}, \nu)$, and, together with Theorem 1 we get that $D(n, \mathcal{A}_d, \nu) = O_d(\log^{d-1/2} n)$ for any Borel probability measure ν on $[0, 1]^d$. This bound has an application to the quasi-Monte Carlo method in numerical integration. A version of the Koksma–Hlawka inequality for general measures given by Götz [18] shows that for any real-valued function f on $[0, 1]^d$ of bounded total variation in the sense of Hardy and Krause, we have

$$\left| \int_{[0,1]^d} f(x) d\nu(x) - \frac{1}{n} \sum_{x \in P} f(x) \right| \leq \frac{1}{n} V(f) D(P, \mathcal{A}_d, \nu).$$

Here, $V(f)$ is the Hardy–Krause variation of f . So, our upper bound implies that for any function f of constant total variation, and any Borel measure ν , we can numerically estimate the integral of f with respect to ν by averaging the values of f at n points, and the estimate has error $O_d(n^{-1} \log^{d-1/2}(n))$. Contrast this with the $O_d(n^{-1/2})$ error rate achieved by the Monte Carlo method. The

error rate we achieve is within $O_d(\sqrt{\log n})$ of the best error rate known for the Lebesgue measure. Integration with respect to measures other than the Lebesgue measure arises often, and a constructive proof of our upper bound could have significant impact in practice. We refer to the note [1] for an exposition of these connections.

2.2. Range searching lower bounds. In the dynamic range searching problem in computer science, we are given a *range space* $\mathcal{F}(P)$, where \mathcal{F} is a collection of subsets of \mathbb{R}^d , and P is an n -point set in \mathbb{R}^d ; our goal is to design a data structure which keeps a set of weights $w \in G^P$ under updates, where the weights come from a commutative group G . An update specifies a point p of P and an element $g \in G$, and asks to change the weight of p to $w_p + g$. The data structure should be able to answer range searching queries, where a query is specified by a range $F \in \mathcal{F}(P)$, and must return the answer $\sum_{p \in F} w_p$. The main question for this data structure problem is to identify a tight trade-off between the update time and the query time.

Fredman [17] first observed that many data structures for the dynamic range searching problem can be identified with a matrix factorization $A = UV$ of the incidence matrix A of $\mathcal{F}(P)$ into two matrices U and V with integer entries. Following Larsen [22], we define an oblivious data structure with multiplicity Δ for the dynamic range searching problem for the range space $\mathcal{F}(P)$ as a factorization $A = UV$ of the incidence matrix of $\mathcal{F}(P)$ into matrices U and V with integer entries bounded in absolute value by Δ . The update time t_u for such a data structure equals the maximum number of non-zero entries in any column of V ; the query time t_q equals the maximum number of non-zero entries in any row of U .

Our arguments imply the following result.

THEOREM 5. *For any generic convex polytope B in \mathbb{R}^d there exists a family of point sets P in \mathbb{R}^d so that for any family of oblivious data structures with multiplicity Δ for the dynamic range counting problem with range space $\mathcal{T}_B(P)$, we have $\sqrt{t_u t_q} = \Omega_{d,B}(\Delta^{-1} \log^d n)$. Conversely, for any convex polytope B and any n -point set P in \mathbb{R}^d there exists an oblivious data structure with multiplicity $\Delta = 1$ and $\sqrt{t_u t_q} = O_{d,B}(\log^d n)$.*

Theorem 5 implies that, as for the discrepancy question, the geometric mean of query and update time grows with n at the same rate for any (generic) convex polytope B , and that order of growth is already achieved for orthogonal range searching. Our upper and lower bounds are tight up to constants when the multiplicity is bounded. Our main contributions are the lower bounds, while the upper bounds follow from standard techniques.

2.3. Differential privacy. Range searching and range counting problems also naturally arise in the field of private data analysis. The setting here is that, given the range space $\mathcal{F}(P)$ as above, we have as input a database D which is a multiset of points from P . The goal is to output the number (with multiplicity)

of points in D that fall in each range $F \in \mathcal{F}(P)$. However, the database D could be sensitive because it may, for example, encode the locations of different people. For this reason, we want to approximate the range counts under the constraints of differential privacy [14, 15]. Formally, a randomized algorithm \mathcal{M} (“mechanism” in the terminology of differential privacy) is (ε, δ) -differentially private if, for any two databases D, D' that differ in the location of single point, and all measurable subsets S of the range of \mathcal{M} , we have

$$\mathbb{P}(\mathcal{M}(D) \in S) \leq e^\varepsilon \mathbb{P}(\mathcal{M}(D') \in S) + \delta.$$

We define the error of a mechanism \mathcal{M} as the maximum of

$$\mathbb{E} \max_{F \in \mathcal{F}(P)} |\mathcal{M}(D)_F - |D \cap F||$$

over databases D , where the expectation is with respect to the randomness of \mathcal{M} , and $\mathcal{M}(D)_F$ is the output that \mathcal{M} gives on input D for the range F . Let $\text{opt}_{\varepsilon, \delta}(\mathcal{F}(P))$ be the smallest achievable error of an (ε, δ) -differentially private algorithm on $\mathcal{F}(P)$, and let $\text{opt}_{\varepsilon, \delta}(N, \mathcal{F}) = \sup_P \text{opt}_{\varepsilon, \delta}(\mathcal{F}(P))$, where the supremum is over all N -point sets P in \mathbb{R}^d .

Our techniques imply the following result.

THEOREM 6. *For any generic convex polytope B in \mathbb{R}^d , for all small enough ε , and all δ small enough with respect to ε , we have*

$$\begin{aligned} \text{opt}_{\varepsilon, \delta}(N, \mathcal{T}_B) &= \Omega_{d, B}(\varepsilon^{-1} \log^{d-1} N), \\ \text{opt}_{\varepsilon, \delta}(N, \mathcal{T}_B) &= O_{d, B}(\varepsilon^{-1} \sqrt{\log 1/\delta} \log^{d+1/2} N). \end{aligned}$$

Moreover, the upper bound holds for any (not necessarily generic) convex polytope B .

Once again, our result shows that the growth of the best possible error under differential privacy with N for the range counting problem with ranges induced by a convex polytope B does not depend strongly on the structure of the particular polytope B .

§3. Preliminaries and techniques. In what follows C_p and c_p are constants that depend only on the parameter p and may change from one line to the next. All logarithms are assumed to be in base e (although usually this does not matter). We use $\langle \cdot, \cdot \rangle$ for the standard inner product on \mathbb{R}^n , and $\|\cdot\|_2$ for the corresponding Euclidean norm. We use capital letters to denote matrices, and lower case letters with indexes to denote matrix entries, e.g. a_{ij} to denote the entry in row i column j of the matrix A . We also use μ_n for the standard Gaussian measure on \mathbb{R}^n . (We avoid the more standard γ_n to avoid confusion with the γ_2 factorization norm.) We use σ_{d-1} for the uniform (rotation invariant, Hausdorff) probability measure on the $(d-1)$ -dimensional unit Euclidean sphere S^{d-1} in \mathbb{R}^d . We use θ_d for the Haar probability measure on the orthogonal group $\mathbf{O}(d)$.

3.1. Hereditary discrepancy. Hereditary discrepancy is a robust version of combinatorial discrepancy. For a set system \mathcal{F} of subsets of a set P , the hereditary discrepancy of \mathcal{F} is defined by $\text{herdisc } \mathcal{F} = \max_{Q \subseteq P} \text{disc } \mathcal{F}(Q)$. Hereditary discrepancy is often more tractable than discrepancy itself, both analytically, and computationally. For example, while a non-trivial approximation to discrepancy is in general NP-hard [11], hereditary discrepancy can be approximated up to polylogarithmic factors [28, 30]. Importantly for us, this robustness comes at no additional cost: for all collections of sets \mathcal{F} that we study, it is easy to see that, for any n -point set P in \mathbb{R}^d , $\text{disc}(n, \mathcal{F}) \geq \text{herdisc } \mathcal{F}(P)$.

3.2. The γ_2 factorization norm. The γ_2 norm was introduced in functional analysis to study operators that factor through Hilbert space. We say that an operator $u : X \rightarrow Y$ between Banach spaces X and Y factors through a Hilbert space if there exists a Hilbert space H and bounded operators $u_1 : X \rightarrow H$ and $u_2 : H \rightarrow Y$ such that $u = u_2 u_1$. Then the γ_2 norm of u is

$$\gamma_2(u) := \inf \|u_1\| \|u_2\|,$$

where the infimum is taken over all Hilbert spaces H , and all operators u_1 and u_2 as above. Here $\|u_1\|$ and $\|u_2\|$ are the operator norms of u_1 and u_2 , respectively. Tomczak-Jaegermann's book [35] is an excellent reference on factorization norms and their applications in Banach space theory.

In this work we will use the γ_2 norm of an $m \times n$ matrix A , which is defined as the γ_2 norm of the linear operator $u : \ell_1^n \rightarrow \ell_\infty^m$ with matrix A (in the standard bases of \mathbb{R}^m and \mathbb{R}^n). In the language of matrices, this means that

$$\gamma_2(A) := \inf \{\|U\|_{2 \rightarrow \infty} \|V\|_{1 \rightarrow 2} : A = UV\},$$

where the infimum is taken over matrices U and V , $\|V\|_{1 \rightarrow 2}$ equals the largest ℓ_2 norm of a column of V , and $\|U\|_{2 \rightarrow \infty}$ equals the largest ℓ_2 norm of a row of U . By a standard compactness argument, the infimum is achieved; moreover, we can take $U \in \mathbb{R}^{m \times r}$ and $V \in \mathbb{R}^{r \times n}$, where r is the rank of A . Yet another equivalent formulation, which will be convenient for us, is that $\gamma_2(A)$ is the smallest non-negative real t for which there exist vectors $u_1, \dots, u_m, v_1, \dots, v_n \in \mathbb{R}^r$ such that for any $i \in [m]$, $j \in [n]$, $a_{ij} = \langle u_i, v_j \rangle$ and $\|u_i\|_2 \leq t$, $\|v_j\|_2 \leq 1$. For a proof of the not completely trivial fact that γ_2 is a norm, see Tomczak-Jaegermann [35]. Given a matrix A , $\gamma_2(A)$ can be computed efficiently by solving a semidefinite program [24].

Let us further overload the meaning of γ_2 by defining $\gamma_2(\mathcal{F}) = \gamma_2(A)$ for a set system \mathcal{F} with incidence matrix A . We recall that the incidence matrix of a system $\mathcal{F} = \{F_1, \dots, F_m\}$ of subsets of a set $P = \{p_1, \dots, p_n\}$ is defined as

$$a_{ij} := \begin{cases} 1 & p_j \in F_i \\ 0 & p_j \notin F_i. \end{cases}$$

In other words, $\gamma_2(\mathcal{F})$ is the smallest non-negative real t such that there exist vectors $u_1, \dots, u_m, v_1, \dots, v_n$ satisfying

$$\langle u_i, v_j \rangle = \begin{cases} 1 & p_j \in F_i \\ 0 & p_j \notin F_i, \end{cases}$$

and $\|u_i\|_2 \leq t$, $\|v_j\|_2 \leq 1$ for all $i \in [m]$, $j \in [n]$.

In [28, 30] it was shown that $\gamma_2(\mathcal{F})$ is, up to logarithmic factors, equivalent to hereditary discrepancy:

$$\text{herdisc } \mathcal{F} \leq C \sqrt{\log m} \gamma_2(\mathcal{F}), \quad (1)$$

$$\text{herdisc } \mathcal{F} \geq c \frac{\gamma_2(\mathcal{F})}{\log \text{rank } A}, \quad (2)$$

where $C, c > 0$ are absolute constants.

In [27, 28] it was also shown that γ_2 satisfies a number of nice properties which help in estimating the norm of specific matrices or set systems. Here we only need the following inequality, which holds for a set system $\mathcal{F} = \mathcal{F}_1 \cup \dots \cup \mathcal{F}_k$, where $\mathcal{F}_1, \dots, \mathcal{F}_k$ are set systems over the same set P :

$$\gamma_2(\mathcal{F}) \leq \sqrt{\sum_{i=1}^k \gamma_2(\mathcal{F}_i)^2}. \quad (3)$$

We also need the following simple lemma, which follows, e.g. from the results in [27, 28], but also appears in a similar form in [22], and follows from standard dyadic decomposition techniques.

LEMMA 7. *For any $d \geq 1$ there exists a constant C_d such that for any n -point set $P \subset \mathbb{R}^d$, $\gamma_2(\mathcal{A}_d(P)) \leq C_d(1 + \log n)^d$. Moreover, there exists a factorization of the incidence matrix of $\mathcal{A}_d(P)$ achieving this bound into matrices U and V with entries in the set $\{0, 1\}$.*

To prove lower bounds on discrepancy, we use (2) and a dual formula for γ_2 :

$$\gamma_2(A) = \max\{\|PAQ\|_{\text{tr}} : P, Q \text{ diagonal, } \text{tr } P^2 = \text{tr } Q^2 = 1\}, \quad (4)$$

where $\|M\|_{\text{tr}}$ is the trace norm of the matrix M , equal to the sum of singular values. For a proof of this formula, see [23, 30]. In this paper, we use the easily proved special case derived from (4) by setting $P = Q = (1/n)I$:

$$\gamma_2(A) \geq \frac{1}{n} \|A\|_{\text{tr}}. \quad (5)$$

See [24] for a short direct proof of this inequality.

The γ_2 norm is also directly connected to upper and lower bounds for data structures and differentially private mechanisms for range searching and range

counting. It follows immediately from the definitions (see §2.2) that the update time t_u and the query time t_q for an oblivious range searching data structure with multiplicity Δ satisfy $\sqrt{t_u t_q} \geq \Delta^{-1} \gamma_2(\mathcal{F}(P))$; moreover, if $\gamma_2(\mathcal{F}(P))$ is achieved by a factorization into matrices with entries in $\{-\Delta, \dots, \Delta\}$, then $\sqrt{t_u t_q} \leq \gamma_2(\mathcal{F}(P))$. In differential privacy, we have the following theorem (see §2.3 for the definitions).

THEOREM 8. *There exists an absolute constant such that the following holds for all small enough ε , and δ small enough with respect to ε . Let \mathcal{F} be a collection of subsets of \mathbb{R}^d and let P be an N -point set in \mathbb{R}^d . Then,*

$$\begin{aligned} \frac{1}{C \log |\mathcal{F}(P)|} \frac{\gamma_2(\mathcal{F}(P))}{\varepsilon} &\leq \text{opt}_{\varepsilon, \delta}(\mathcal{F}(P)) \\ &\leq C \sqrt{(\log |\mathcal{F}(P)|)(\log 1/\delta)} \frac{\gamma_2(\mathcal{F}(P))}{\varepsilon}. \end{aligned}$$

Theorem 8 was essentially proved, although stated differently, in [31]. For a statement equivalent to the one above, and a proof, see [29, Theorem 7.1]. Note that [29] uses the notation $\|\cdot\|_{E_\infty}$ in place of the standard $\gamma_2(\cdot)$.

3.3. Signed series of vectors. We will use the following result of Banaszczyk.

LEMMA 9 [4]. *Let $v_1, \dots, v_n \in \mathbb{R}^m$, $\forall i : \|v_i\|_2 \leq 1$, and let $K \subset \mathbb{R}^m$ be a convex body symmetric around the origin. If $\mu_m(K) \geq 1 - 1/(2n)$, then there exists an assignment of signs $\chi : [n] \rightarrow \{-1, 1\}$ so that*

$$\forall j \in \{1, \dots, n\} : \sum_{i=1}^j \chi(i) v_i \in 5K.$$

This lemma was proved in the context of the well-known Steinitz problem: given vectors v_1, \dots, v_n , each of Euclidean norm at most 1, such that $v_1 + \dots + v_n = 0$, find a permutation π on $[n]$ such that for all integers i , $1 \leq i \leq n$, $\|v_{\pi(1)} + \dots + v_{\pi(i)}\|_2 \leq C\sqrt{m}$, where C is an absolute constant independent of m or n . Lemma 9 gives the best known partial result in this direction: it can be used to show a bound of $C(\sqrt{m} + \sqrt{\log n})$ in place of $C\sqrt{m}$.

Lemma 9 follows relatively easily from a powerful result Banaszczyk proved in [3]. Unfortunately, the proof of the latter does not suggest any efficient algorithm to find the signs $\chi(i)$, and no such algorithm is yet known, despite some partial progress [5].

3.4. Techniques. Our approach builds on the connection between the γ_2 norm and hereditary discrepancy shown in [28, 30]. The new idea which enables the tighter upper bound is the use of Banaszczyk's signed series result (Lemma 9), in order to get one dimension "for free". On a very high level, this is similar to the way one dimension comes for free in constructions of point sets

with small geometric discrepancy, such as the Halton–Hammersley construction mentioned before. The proof of Theorem 2 combines the ideas in the proof of Theorem 1 with a decomposition given by Matoušek [25].

The lower bound for Tusnády's problem in [28] relies crucially on the product structure of \mathcal{A}_d , and does not easily extend to \mathcal{T}_B when B is a polytope other than a box. Instead, for the lower bound in this paper, we combine the factorization norm approach with the Fourier method, pioneered in discrepancy theory by Roth [33] and developed further by Beck: see [8] for an exposition. In order to give a lower bound on γ_2 , we use (5), and estimate the trace norm of the incidence matrix M of a set system related to $\mathcal{T}_B(P)$, where P is a grid in $[0, 1]^d$. We define M so that it is a convolution matrix (i.e. the matrix of a convolution operator), and its eigenvalues are given by the discrete Fourier transform. While tight estimates are known for the average decay of continuous Fourier coefficients of convex polytopes, we need estimates on discrete Fourier coefficients, about which much less is known. To bridge this gap, we prove a bound on the convergence rate of discrete Fourier coefficients of convex polytopes to the continuous Fourier coefficients. We also use an averaging argument in order to be able to work with bounds on the average decay of the Fourier spectrum, rather than having to estimate specific Fourier coefficients. These techniques are general, and may be more widely applicable to geometric combinatorial discrepancy questions. Our version of the Fourier method has the curious feature that, even though we average over rotations of the polytope B , in the end the lower bound holds for the set system induced only by translations and dilations of B .

§4. Upper bound for Tusnády's problem. In this section we give the proof of Theorem 1. Let us fix the n -point set $P \subset \mathbb{R}^d$ once and for all. Without loss of generality, assume that each $p \in P$ has a distinct last coordinate, and order the points in P in increasing order of their last coordinate as p_1, \dots, p_n . Write each p_i as $p_i = (q_i, r_i)$, where $q_i \in \mathbb{R}^{d-1}$ and $r_i \in \mathbb{R}$. With this notation, and the ordering we assumed, we have that $r_i < r_j$ whenever $i < j$.

Let $Q := \{q_i : 1 \leq i \leq n\}$. Notice that this is an n -point set in \mathbb{R}^{d-1} . Denote the sets in $\mathcal{A}_{d-1}(Q)$ as A_1, \dots, A_m (in no particular order). Using Lemma 7, there exist vectors u_1, \dots, u_m and v_1, \dots, v_n such that

$$\langle u_i, v_j \rangle = \begin{cases} 1 & q_j \in A_i \\ 0 & q_j \notin A_i, \end{cases} \quad (6)$$

and $\|u_i\|_2 \leq C_d(1 + \log n)^{(d-1)}$, $\|v_j\|_2 \leq 1$ for all i and j . Define the symmetric polytope

$$K := \{x \in \mathbb{R}^m : |\langle u_i, x \rangle| \leq C'_d(1 + \log n)^{d-1/2} \text{ for all } i \in \{1, \dots, m\}\},$$

where $C'_d > C_d$ is a constant large enough that $\mu_m(K) \geq 1 - 1/(2n)$. The fact that such a constant exists follows from standard concentration of measure results in Gaussian space. Indeed, using a Bernstein-type inequality for Gaussian measure, we can show that, for C'_d big enough, $\mu_m(S_i) \geq 1 - 1/(2n^{d+1})$ for all

$i \in [m]$, where $S_i := \{x : |\langle u_i, x \rangle| \leq C'_d(1 + \log n)^{d-1/2}\}$. By the union bound, since $m \leq n^d$, this implies $\mu_m(K) = \mu_m(\bigcap_{i=1}^m S_i) \geq 1 - 1/2n$.

The body K and the vectors v_1, \dots, v_n then satisfy the assumptions of Lemma 9, and, therefore, there exists an assignment of signs $\chi : [n] \rightarrow \{-1, 1\}$ such that, for any k , $1 \leq k \leq n$,

$$\sum_{j=1}^k \chi(j)v_j \in 5K.$$

By the definition of K , this is equivalent to

$$\forall i \in \{1, \dots, m\} : \left| \sum_{j=1}^k \chi(j)\langle u_i, v_j \rangle \right| \leq 5C'_d(1 + \log n)^{d-1/2}. \quad (7)$$

For each i , $1 \leq i \leq m$, let us define $A'_i = \{p_j : q_j \in A_i\}$. We claim that for any $x \in \mathbb{R}^d$, we can write $A(x) \cap P$ as $A'_i \cap \{p_1, \dots, p_k\}$ for some i and k . (Here we assume that $A(x) \cap P$ is non-empty: the other case is irrelevant to the proof.) To see this, let $x = (y, x_d)$, where $y \in \mathbb{R}^{d-1}$ and $x_d \in \mathbb{R}$. Let i be such that $A(y) \cap Q = A_i$, and let k be the largest integer such that $r_k \leq x_d$. Then:

$$A(x) \cap P = \{p_j : q_j \in A(y), j \leq k\} = A'_i \cap \{p_1, \dots, p_k\}.$$

It follows that

$$\begin{aligned} \left| \sum_{j: p_j \in A(x)} \chi(j) \right| &= \left| \sum_{j: q_j \in A_i, j \leq k} \chi(j) \right| \\ &= \left| \sum_{j \leq k} \chi(j)\langle u_i, v_j \rangle \right| \leq 5C'_d(1 + \log n)^{d-1/2}, \end{aligned}$$

where the penultimate equality follows from (6), and the final inequality is (7). Since x was arbitrary, we have shown that $\text{disc } \mathcal{A}_d(P) \leq 5C'_d(1 + \log n)^{d-1/2}$, as was required.

§5. Upper bound for an arbitrary polytope. The main ingredient in extending Theorem 1 to arbitrary polytopes is a geometric decomposition given by Matoušek. To describe the decomposition we define the admissible k -composition $\text{AC}_k(\mathcal{F})$ of sets from a (finite) set system \mathcal{F} as follows. For $k = 0$, $\text{AC}_k(\mathcal{F}) = \emptyset$; for an integer $k > 0$, we have

$$\begin{aligned} \text{AC}_k(\mathcal{F}) &:= \{F_1 \cup F_2 : F_1 \in \text{AC}_{k_1}(\mathcal{F}), F_2 \in \text{AC}_{k_2}(\mathcal{F}), F_1 \cap F_2 = \emptyset, k_1 + k_2 = k\} \\ &\quad \cup \{F_1 \setminus F_2 : F_1 \in \text{AC}_{k_1}(\mathcal{F}), F_2 \in \text{AC}_{k_2}(\mathcal{F}), F_2 \subseteq F_1, k_1 + k_2 = k\}. \end{aligned}$$

By an easy induction on k , we see that

$$\text{disc } \text{AC}_k(\mathcal{F}) \leq k \text{disc } \mathcal{F}, \quad (8)$$

$$\gamma_2(\text{AC}_k(\mathcal{F})) \leq k\gamma_2(\mathcal{F}). \quad (9)$$

We also extend the notion of anchored boxes to “corners” whose bounding hyperplanes are not necessarily orthogonal. Let $W = \{w_1, \dots, w_d\}$ be a basis of \mathbb{R}^d . Then we define $\mathcal{A}_W := \{A_W(x) : x \in \mathbb{R}^d\}$, where $A_W(x) = \{y \in \mathbb{R}^d : \langle w_i, y \rangle \leq \langle w_i, x \rangle \text{ for all } i \in [d]\}$.

The following lemma gives the decomposition result we need.

LEMMA 10 [25]. *Let $B \subset \mathbb{R}^d$ be a convex polytope. There exists a constant k depending on d and B , and k bases W_1, \dots, W_k of \mathbb{R}^d such that every $B' \in \mathcal{T}_B$ belongs to $\text{AC}_k(\mathcal{A}_{W_1} \cup \dots \cup \mathcal{A}_{W_k})$. Moreover, $e_1 \in W_1 \cap W_2 \cap \dots \cap W_k$, where e_1 is the first standard basis vector of \mathbb{R}^d .*

Matoušek does not state the condition after “moreover”; nevertheless, it is easy to verify this condition holds for the recursive decomposition in his proof of Lemma 10.

We will also need a bound on $\gamma_2(\mathcal{A}_W(P))$ for a basis W and an n -point set P .

LEMMA 11. *For any $\ell \geq 1$ there exists a constant C_ℓ such that the following holds. For any set W of ℓ linearly independent vectors in \mathbb{R}^d , and any n -point set P , $\gamma_2(\mathcal{A}_W(P)) \leq C_\ell(1 + \log n)^\ell$.*

Proof. Let $W = \{w_1, \dots, w_\ell\}$, and let u be the linear transformation from \mathbb{R}^ℓ to \mathbb{R}^d that sends the i th standard basis vector e_i to w_i for each $i \in [\ell]$. Let $Q = u^*(P) := \{u^*(p) : p \in P\}$, where u^* is the adjoint of u . It is easy to verify that $\mathcal{A}_W(P) = \mathcal{A}_\ell(Q)$, and, therefore,

$$\gamma_2(\mathcal{A}_W(P)) = \gamma_2(\mathcal{A}_\ell(Q)) \leq C_\ell(1 + \log n)^\ell,$$

where the final inequality follows from Lemma 7. □

As a warm-up, let us first prove an upper bound on $\gamma_2(\mathcal{T}_B(P))$.

THEOREM 12. *For any $d \geq 1$ and any closed convex polytope $B \subset \mathbb{R}^d$ there exists a constant $C_{d,B}$ such that for any set P of n points in \mathbb{R}^d ,*

$$\gamma_2(\mathcal{T}_B(P)) \leq C_{d,B}(1 + \log n)^d.$$

Proof. Let W_1, \dots, W_k be as in Lemma 10. Using (3) and Lemma 11,

$$\gamma_2(\mathcal{A}_{W_1}(P) \cup \dots \cup \mathcal{A}_{W_k}(P)) \leq \sqrt{k}C_d(1 + \log n)^d.$$

Using Lemma 10, $\mathcal{T}_B(P) \subseteq \text{AC}_k(\mathcal{A}_{W_1}(P) \cup \dots \cup \mathcal{A}_{W_k}(P))$. Together with (9) and the trivial fact that $\gamma_2(\mathcal{F}') \leq \gamma_2(\mathcal{F})$ whenever $\mathcal{F}' \subseteq \mathcal{F}$, this implies

$$\begin{aligned} \gamma_2(\mathcal{T}_B(P)) &\leq \gamma_2(\text{AC}_k(\mathcal{A}_{W_1}(P) \cup \dots \cup \mathcal{A}_{W_k}(P))) \\ &\leq k\gamma_2(\mathcal{A}_{W_1}(P) \cup \dots \cup \mathcal{A}_{W_k}(P)) \\ &\leq k\sqrt{k}C_d(1 + \log n)^d. \end{aligned}$$

Since k depends on d and B only, this finishes the proof of the theorem. □

Proof of Theorem 2. As in the proof of Theorem 1, we fix the n -point set $P \subset \mathbb{R}^d$ once and for all, and we order P as p_1, \dots, p_n in increasing order of the last coordinate. We write each p_i as $p_i = (q_i, r_i)$ for $q_i \in \mathbb{R}^{d-1}$ and $r_i \in \mathbb{R}$, and define $Q := \{q_i : 1 \leq i \leq n\}$.

Let W_1, \dots, W_k be as in Lemma 10, and let $W'_i = W_i \setminus \{e_1\}$ for each i , $1 \leq i \leq k$. Observe that W'_i is a set of $d-1$ linearly independent vectors, and, using (3) and Lemma 11,

$$\gamma_2(\mathcal{A}_{W'_1}(Q) \cup \dots \cup \mathcal{A}_{W'_k}(Q)) \leq \sqrt{k} C_d (1 + \log n)^{d-1}.$$

With an argument using Lemma 9 analogous to the one used in the proof of Theorem 1, we can then show that there exists a constant $C'_{d,B}$ depending on B and d and a coloring $\chi : [n] \rightarrow \{-1, 1\}$, such that for any integer i , $1 \leq i \leq k$, and any $x \in \mathbb{R}^d$,

$$\left| \sum_{j: p_j \in \mathcal{A}_{W'_i}(x)} \chi(j) \right| \leq C'_{d,B} (1 + \log n)^{d-1/2}.$$

Here $C'_{d,B}$ is implicitly assumed to depend on k as well, which depends on B and d . This establishes that $\text{disc } \mathcal{F} \leq C'_{d,B} (1 + \log n)^{d-1/2}$, where $\mathcal{F} = \mathcal{A}_{W_1}(P) \cup \dots \cup \mathcal{A}_{W_k}(P)$. Because, using Lemma 10, $\mathcal{T}_B(P) \subseteq \text{AC}_k(\mathcal{F})$, (8) implies

$$\text{disc } \mathcal{T}_B(P) \leq \text{disc } \text{AC}_k(\mathcal{F}) \leq k \text{disc } \mathcal{F} \leq k C'_{d,B} (1 + \log n)^{d-1/2}.$$

This finishes the proof of the theorem. \square

The same asymptotic bound with \mathcal{T}_B replaced by $\text{POL}(\mathcal{H})$ for a family of hyperplanes \mathcal{H} can be proved by replacing Lemma 10 with an analogous decomposition lemma for $\text{POL}(\mathcal{H})$, also proved in [25].

The upper bound on $\sqrt{t_u t_q}$ in Theorem 5 follows from Theorem 12 and the observation that the upper bound on γ_2 can be achieved by a factorization into matrices with entries in $\{0, 1\}$, which is equivalent to an oblivious data structure with multiplicity 1. The upper bound on error in Theorem 6 follows from Theorems 8 and 12.

§6. Lower bound. In this section we prove lower bounds on $\text{disc}(n, \mathcal{T}_B)$ matching the known lower bounds on $\text{disc}(n, \mathcal{A}_d)$ up to constants when B is a generic convex polytope (“generic” is defined below). In order to use Fourier analytic techniques, it will be convenient to work with a modification of the incidence matrix of $\mathcal{T}_B(P)$. To this end, let us call a function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ periodic if for every $x \in [0, 1)^d$ and every vector y in the integer lattice \mathbb{Z}^d we have $f(x) = f(x+y)$. Let $Q_n := \{i/n\}_{i=0}^{n-1}$. For a periodic function $f : \mathbb{R}^d \rightarrow \mathbb{R}$, define a real $n^d \times n^d$ matrix $M(f, n)$ indexed by Q_n^d :

$$m_{x,y}(f, n) := f(x - y),$$

where, x and y range over Q_n^d . Of special interest to us are periodic functions defined by convex polytopes $B \subseteq [0, 1]^d$. Let us define the function $f_B : \mathbb{R}^d \rightarrow \mathbb{R}$ to be equal to the indicator function 1_B of B on $[0, 1]^d$, extended periodically to the rest of \mathbb{R}^d . For convenience, we use the notation $M(B, n) := M(f_B, n)$.

There are two main observations that motivate studying $M(B, n)$. First, each row of $M(B, n)$ is the indicator vector of the disjoint union of most 2^d sets from $\mathcal{T}_B(Q_n^d) = \mathcal{T}_B(-Q_n^d) = \mathcal{T}_B(Q_n^d)$, so any lower bound on the hereditary discrepancy of $M(B, n)$ implies a lower bound on $\text{disc}(n, \mathcal{T}_B)$. Second, because $M(B, n)$ is the matrix of a convolution operator, it is diagonalized by the (discrete multidimensional) Fourier transform, and we can use known results on the Fourier spectra of convex polytopes to derive bounds on the trace norm of $M(B, n)$, and therefore on $\gamma_2(M(B, n))$.

Before we continue, let us introduce the standard notation for the Fourier coefficients. For the remainder of this section, we use $i = \sqrt{-1}$ to denote the imaginary unit. For any $u \in \mathbb{R}^d$ and a periodic function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ integrable on $[0, 1]^d$ we define the Fourier coefficient

$$\hat{f}(\xi) := \int_{[0,1]^d} f(x) e^{-2\pi i \langle \xi, x \rangle} dx.$$

We also define the discrete Fourier coefficients:

$$\tilde{f}(\xi, n) := \frac{1}{n^d} \sum_{q \in Q_n} f(q) e^{-2\pi i \langle \xi, q \rangle}.$$

It is well known, and easy to verify, that the eigenvalues of $M(B, n)$ are given by $n^d \tilde{f}_B(\xi, n)$. There is quite a bit known about the continuous Fourier coefficients $\hat{f}_B(\xi)$, but comparatively less known about the discrete Fourier coefficients. Intuitively, bounds on the Fourier coefficients are easier to prove in the continuous domain because powerful tools like the divergence theorem are available. In order to use the known bounds on $\hat{f}_B(\xi)$, we estimate the rate of convergence of $\tilde{f}_B(\xi, n)$ to $\hat{f}_B(\xi)$ with n . We follow an approach similar to that used by Epstein [16] in the one-dimensional setting. In order to adapt his results to our setting, we need two additional ingredients: a higher-dimensional analog of Jackson's theorem in approximation theory, and a continuous approximation to the indicator function f_B . Next, we state the higher-dimensional Jackson-type theorem (given by Yudin, also spelled Judin) that we use.

Definition 1. The *modulus of continuity* of a continuous function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is defined as $\omega(f, \delta) := \sup\{|f(x+t) - f(x)| : x, t \in \mathbb{R}^d, \|t\|_2 \leq \delta\}$.

THEOREM 13 [21]. *There exists a universal constant C such that for any function f and any integer $n \geq 1$ there exists an order n trigonometric polynomial p defined by*

$$p(x) = \sum_{v \in \mathbb{Z}^d : \|v\|_\infty \leq n} c_v e^{2\pi i \langle v, x \rangle}$$

such that

$$\|f - p\|_{\infty} := \sup_{x \in [0,1]^d} |f(x) - p(x)| \leq 4\omega\left(f, \frac{C\sqrt{d}}{n}\right).$$

The next lemma is an analogue of [16, Theorem 3.1].

LEMMA 14. *There exists a universal constant C such that the following holds. Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be a continuous periodic function with modulus of continuity $\omega(f, \delta)$. Then, for any $\xi \in \mathbb{Z}^d$ such that $\|\xi\|_{\infty} \leq n$, we have*

$$|\tilde{f}(\xi, 2n+1) - \hat{f}(\xi)| \leq 8\omega\left(f, \frac{C\sqrt{d}}{n}\right).$$

Proof. Let p be the trigonometric polynomial of order n guaranteed by Theorem 13. Using an elementary calculation, for any $\xi \in \mathbb{Z}^d$ such that $\|\xi\|_{\infty} \leq n$, we have $\tilde{p}(\xi, 2n+1) = \hat{p}(\xi) = c_{\xi}$, where c_{ξ} is the coefficient of $e^{2\pi i \langle \xi, \cdot \rangle}$ in the expansion of p . For any ξ as above, this gives us:

$$\begin{aligned} |\tilde{p}(\xi, 2n+1) - \hat{f}(\xi)| &= |\hat{p}(\xi) - \hat{f}(\xi)| \\ &= \left| \int_{[0,1]^d} (p(x) - f(x)) e^{-2\pi i \langle \xi, x \rangle} dx \right| \leq \|f - p\|_{\infty}, \end{aligned}$$

where we used the trivial case of Hölder's inequality. Similarly,

$$\begin{aligned} |\tilde{f}(\xi, 2n+1) - \tilde{p}(\xi, 2n+1)| &= \frac{1}{(2n+1)^d} \left| \sum_{q \in Q_n} (f(x) - p(x)) e^{-2\pi i \langle \xi, q \rangle} \right| \\ &\leq \|f - p\|_{\infty}. \end{aligned}$$

Combining the two inequalities, and using the bound in Theorem 13, we get

$$\begin{aligned} |\tilde{f}(\xi, 2n+1) - \hat{f}| &\leq |\tilde{f}(\xi, 2n+1) - \tilde{p}(\xi, 2n+1)| + |\tilde{p}(\xi, 2n+1) - \hat{f}(\xi)| \\ &\leq 2\|f - p\|_{\infty} \leq 8\omega\left(f, \frac{C\sqrt{d}}{n}\right). \end{aligned}$$

This completes the proof. \square

In the next lemma we prove an explicit bound on how fast $\tilde{f}_B(\xi, n)$ converges to $\hat{f}_B(\xi)$ with n for a convex polytope B . The bounds are most likely not tight, but sufficient for our purposes, since we only need that the convergence rate is polynomial in $1/n$.

LEMMA 15. *Let $B \subseteq [0, 1]^d$ be a convex polytope with non-empty interior. Then, for any $\xi \in \mathbb{Z}^d$, $\|\xi\|_{\infty} \leq n$, we have*

$$|\tilde{f}_B(\xi, 2n+1) - \hat{f}_B(\xi)| = O_{d,B}\left(\frac{1}{\sqrt{n}}\right).$$

Proof. In order to apply Lemma 14, we need a continuous function f which approximates f_B . We construct this approximation by, roughly, a piecewise linear interpolation along every direction.

Let c lie in the interior of B , and pick a real number $r > 0$ such that $c + rD^d \subseteq B$, where D^d is the unit Euclidean ball in \mathbb{R}^d centered at the origin. Let us define the gauge function g by $g(x) = \inf\{t : x \in (1-t)c + tB\}$ on $[0, 1)^d$. Note that $g(x) \leq (1/r)\|x - c\|_2$, further, that for any $x, t \in \mathbb{R}^d$, $g(x+t) \leq g(x) + g(t+c)$. Using these two observations, we get

$$g(x+t) - g(x) \leq g(t+c) \leq \frac{1}{r}\|t\|_2.$$

By symmetry, we also get that $g(x) - g(x+t) \leq (1/r)\|t\|_2$. Therefore, $\omega(g, \delta) \leq \delta/r$.

We define a periodic function f on $[0, 1)^d$ by

$$f(x) = \begin{cases} 1 & g(x) \leq 1 - \varepsilon \\ \frac{1}{\varepsilon}(1 - g(x)) & 1 - \varepsilon < g(x) < 1 \\ 0 & g(x) \geq 1, \end{cases}$$

for a parameter $0 < \varepsilon < 1$ which depends on n and will be determined later. We then extend f periodically to the rest of \mathbb{R}^d . The function f is defined so that it is continuous and agrees with f_B except for those x for which $1 - \varepsilon < g(x) < 1$. Moreover, observe that

$$\omega(f, \delta) \leq \frac{1}{\varepsilon}\omega(g, \delta) \leq \frac{\delta}{r\varepsilon}.$$


Then, using Lemma 14, for any $\xi \in \mathbb{Z}^d$ such that $\|\xi\|_\infty \leq n$, we have

$$|\tilde{f}(\xi, 2n+1) - \hat{f}(\xi)| \leq \frac{8C\sqrt{d}}{r\varepsilon n}. \quad (10)$$

It remains to bound $|\hat{f}_B(\xi) - \hat{f}(\xi)|$ and $|\tilde{f}_B(\xi, 2n+1) - \tilde{f}(\xi, 2n+1)|$. Let $S = \{x \in [0, 1)^d : 1 - \varepsilon < g(x) < 1\}$ be the subset of $[0, 1)^d$ on which f and f_B disagree. Notice that the closure of S is $B \setminus (\varepsilon c + (1 - \varepsilon)B) = c + (B - c) \setminus (1 - \varepsilon)(B - c)$ (see Figure 1), so we have

$$\lambda_d(S) = \lambda(B \setminus (1 - \varepsilon)B) = (1 - (1 - \varepsilon)^d)\lambda_d(B) \leq d\varepsilon,$$

where, in the final inequality, we used the assumption $B \subseteq [0, 1)^d$, which implies $\lambda_d(B) \leq 1$. (Recall that we use λ_d for the Lebesgue measure in \mathbb{R}^d .) We can now bound the first of our error terms using Hölder's inequality: for any $\xi \in \mathbb{Z}^d$

Figure 1: The set S on which f and f_B disagree.

$$\begin{aligned} |\hat{f}_B(\xi) - \hat{f}(\xi)| &= \left| \int_{[0,1]^d} (f_B(x) - f(x)) e^{-2\pi i \langle \xi, x \rangle} dx \right| \\ &= \left| \int_S (f_B(x) - f(x)) e^{-2\pi i \langle \xi, x \rangle} dx \right| \leq \lambda_d(S) \leq d\varepsilon. \end{aligned} \quad (11)$$

A similar calculation for the discrete Fourier coefficients gives us

$$\begin{aligned} |\tilde{f}_B(\xi, 2n+1) - \tilde{f}(\xi, 2n+1)| &= \frac{1}{(2n+1)^d} \sum_{q \in Q_n} (f_B - f(x)) e^{-2\pi i \langle \xi, q \rangle} \\ &= \frac{1}{(2n+1)^d} \sum_{q \in Q_n \cap S} (f_B - f(x)) e^{-2\pi i \langle \xi, q \rangle} \\ &\leq \frac{|S \cap Q_n|}{(2n+1)^d} \end{aligned}$$

for any $\xi \in \mathbb{Z}^d$. By a standard volume argument,

$$|S \cap Q_n| \leq \frac{\lambda_d(S + (1/2n)D^d)}{\lambda_d((1/2n)D^d)} = (2n)^d \frac{\lambda_d(S + (1/2n)D^d)}{\lambda_d(D^d)}.$$

Because $(1/2n)D^d \subseteq (1/2rn)(B - c)$, we have $S + (1/2n)D^d \subseteq c + (1 + 1/2rn)(B - c) \setminus (1 - \varepsilon - 1/2rn)(B - c)$, and, therefore,

$$\lambda_d\left(S + \frac{1}{2n}D^d\right) \leq \left(\left(1 + \frac{1}{2rn}\right)^d - \left(1 - \varepsilon - \frac{1}{2rn}\right)^d\right) \lambda_d(B) \leq \varepsilon d + \frac{3d}{2rn},$$

where the final inequality holds for $n \geq d/2r$. Putting the estimates together, we have

$$|\tilde{f}_B(\xi, 2n+1) - \tilde{f}(\xi, 2n+1)| \leq C_d \varepsilon + \frac{C_d}{rn}, \quad (12)$$

for a constant C_d depending on d and all large enough n .

Combining (10), (11), and (12), for all large enough n and any $\xi \in \mathbb{Z}^d$ such that $\|\xi\|_\infty \leq n$ we get

$$|\tilde{f}_B(\xi, 2n+1) - \hat{f}_B(\xi, 2n+1)| \leq \frac{8C\sqrt{d}}{r\varepsilon n} + (C_d + d)\varepsilon + \frac{C_d}{rn}.$$

By setting $\varepsilon = n^{-1/2}$, we get that the right-hand side is in $O_{d,B}(n^{-1/2})$, as required. \square

Inspecting the proof of Lemma 15, we see that the only dependence on B is via the radius r of the Euclidean ball contained in B : all other constants depend on the dimension only. By John's theorem, we can apply an affine transformation to B so that it is contained in $[0, 1]^d$, and in fact in a Euclidean ball of unit radius, and contains a ball of radius $1/d$. However, the estimates we use below on the Fourier coefficients of f_B do depend on B , so we do not pursue this idea further.

We require a technical definition.

Definition 2. We will say that a polytope B in \mathbb{R}^d is *generic* if there exists a sequence of faces $F_1 \subseteq F_3 \subseteq \dots \subseteq F_{d-1}$ of B such that F_j is a face of dimension j , for every $j \leq d-2$ the face F_j is not parallel to any other face of F_{j+1} , and the facet F_{d-1} is not parallel to any other facet of B .

It is easy to see that B is generic for example if the facets of B have normal vectors u_1, \dots, u_k that are in *general position*, in the sense that for every $J \subseteq [k]$ of size at most d the set of vectors $\{u_j : j \in J\}$ spans a subspace of dimension $|J|$. However, the condition of being generic appears to be much weaker. We use the following estimate on the Fourier coefficients of such a polytope.

THEOREM 16 [10]. *Let B be a generic polytope in \mathbb{R}^d for $d \geq 2$. There exists a constant $c_{d,B}$, possibly depending on d and B , such that for any $\rho \geq 1$*

$$\int_{S^{d-1}} |\hat{f}_B(\rho\xi)| d\sigma_{d-1}(\xi) \geq \frac{c_{d,B} \log^{d-1}(\rho)}{\rho^d},$$

where S^{d-1} is the unit Euclidean sphere in \mathbb{R}^d and σ_{d-1} is the uniform probability measure on S^{d-1} .

In [10] Theorem 16 is stated for a simplex, and after the proof the authors remark that the same proof extends to any polytope which has one face not parallel to any other face. However, in their proof, they induct on lower and lower dimensional faces, and this condition needs to hold for any face they induct on. Our definition of generic appears to be the minimal condition under which their proof goes through. The authors of [10] note that some genericity-like condition is necessary since Theorem 16 does not hold for all values of ρ when B is a cube. Nevertheless, it is likely that a variant of the theorem in which we average over values of ρ in a big enough interval could hold for arbitrary polytopes, and may

yield a lower bound with a constant that only depends on d and not on B . We leave this extension for future work.

Theorem 16, Lemma 15, and an averaging argument yield the following estimate on discrete Fourier coefficients, which is the final step towards our lower bound.

LEMMA 17. *Let $B \subseteq c + D^d$ be a generic convex polytope with non-empty interior, where D^d is the Euclidean unit ball in \mathbb{R}^d centered at the origin and c is the centroid of $[0, 1]^d$. There exists an orthogonal transformation u such that $B_u := c + u(B - c)$ satisfies*

$$\sum_{\xi \in \mathbb{Z}^d : \|\xi\|_\infty \leq n} |\tilde{f}_{B_u}(\xi, 2n+1)| = \Omega_{d,B}(\log^d(n)).$$

Proof. Recall that we use $\mathbf{O}(d)$ to denote the group of orthogonal transformations on \mathbb{R}^d , and θ_d to denote the Haar probability measure on this group. We also use σ_{d-1} for the uniform probability measure on the sphere.

For any $\xi \in \mathbb{Z}^d \setminus \{0\}$ we have

$$\begin{aligned} \int_{\mathbf{O}(d)} |\hat{f}_{B_u}(\xi)| d\theta(u) &= \int_{\mathbf{O}(d)} \left| \int_{c+u(B-c)} e^{-2\pi i \langle \xi, x \rangle} dx \right| d\theta(u) \\ &= \int_{\mathbf{O}(d)} \left| \int_{u(B)} e^{-2\pi i \langle \xi, x \rangle} dx \right| d\theta(u) \\ &= \int_{\mathbf{O}(d)} \left| \int_B e^{-2\pi i \langle \xi, u^*(x) \rangle} dx \right| d\theta(u) \\ &= \int_{\mathbf{O}(d)} \left| \int_B e^{-2\pi i \langle u(\xi), x \rangle} dx \right| d\theta(u) \\ &= \int_{S^{d-1}} \left| \int_B e^{-2\pi i \langle \|\xi\|_2 \zeta, x \rangle} dx \right| d\sigma_{d-1}(\zeta) \\ &\geq \frac{c \log^{d-1}(\|\xi\|_2)}{\|\xi\|_2^d}. \end{aligned}$$

Above, in the penultimate line we used the fact that for any measurable set $Y \subseteq S^{d-1}$, and any $y \in S^{d-1}$, $\sigma_{d-1}(Y) = \theta_d(\{u \in \mathbf{O}(d) : u(y) \in Y\})$. The final inequality is implied by Theorem 16 for an appropriate constant c depending on d and B . Therefore,

$$\int_{\mathbf{O}(d)} \left(\sum_{\xi \in \mathbb{Z}^d : 0 < \|\xi\|_2 = j} |\hat{f}_{B_u}(\xi)| \right) d\theta(u) \geq cm_j \frac{\log^{d-1}(j)}{j^{d/2}},$$

where $m_j = \{\xi \in \mathbb{Z}^d : \|\xi\|_2^2 = j\}$. By Lemma 15, for all sufficiently large n and for all j such that $j \leq (c_1 n)^{1/d}$ for a sufficiently small constant c_1 , we

have $\tilde{f}_{B_u}(\xi, 2n+1) \geq \hat{f}_{B_u}(\xi) - c/2j^{d/2}$ for any $\xi \in \mathbb{N}_0^d$ for which $\|\xi\|_2^2 = j$. Therefore,

$$\int_{\mathbf{O}(d)} \left(\sum_{\xi \in \mathbb{Z}^d: 0 < \|\xi\|_2^2 \leq (c_1 n)^{1/d}} |\tilde{f}_{B_u}(\xi, 2n+1)| \right) d\theta_d(u) \geq c \sum_{j=1}^{\lfloor (c_1 n)^{1/d} \rfloor} m_j \frac{\log^{d-1}(j)}{2j^{d/2}}.$$

There must then exist a choice of $u \in \mathbf{O}(d)$ such that

$$\sum_{\xi \in \mathbb{N}_0^d: \|\xi\|_2^2 \leq (c_1 n)^{1/d}} |\tilde{f}_{B_u}(\xi, 2n+1)| \geq c \sum_{j=1}^{\lfloor (c_1 n)^{1/d} \rfloor} m_j \frac{\log^{d-1}(j)}{2j^{d/2}}. \quad (13)$$

Let us fix such a u for the rest of the proof. We proceed to estimate the right-hand side of (13). Define $\ell = \lfloor (c_1 n)^{1/d} \rfloor$ to be the upper bound of the summation, and let $k = \lfloor \sqrt{\ell} \rfloor$. Let $m_{\leq j} = m_1 + \dots + m_j$. Using summation by parts, we have

$$\begin{aligned} \sum_{j=1}^{\ell} m_j \frac{\log^{d-1}(j)}{j^{d/2}} &\geq \log^{d-1}(k) \sum_{j=k}^{\ell} \frac{m_j}{j^{d/2}} \\ &= \log^{d-1}(k) \frac{m_{\leq \ell}}{\ell^{d/2}} \\ &\quad + \log^{d-1}(k) \sum_{j=k}^{\ell-1} m_{\leq j} \left(\frac{1}{j^{d/2}} - \frac{1}{(j+1)^{d/2}} \right). \end{aligned}$$

By standard estimates (e.g. Minkowski's first theorem), there exists a constant c_2 depending on the dimension d such that $m_{\leq j} \geq c_2 j^{d/2}$. By convexity, $1/j^{d/2} - 1/(j+1)^{d/2} \geq (d/2)/(j+1)^{(d+2)/2}$. Plugging these inequalities into the bound above, we get that

$$\sum_{j=1}^{\ell} m_j \frac{\log^{d-1}(j)}{j^{d/2}} \geq c_2 \log^{d-1}(k) \sum_{j=k}^{\ell} \frac{j^{d/2}}{(j+1)^{(d+2)/2}} = \Omega_d(\log^d n).$$

Together with (13), this completes the proof of the lemma. \square

Proof of Theorem 3. By scaling we can assume that $B \subseteq c + D^d$, where D^d is the Euclidean unit ball in \mathbb{R}^d centered at the origin and c is the centroid of $[0, 1)^d$. Let u be the orthogonal transformation given by Lemma 17, and let $M := M(B_u, 2n+1)$. It is easy to verify that M is diagonalized by the collection of orthogonal eigenvectors $\{(e^{2\pi i \langle \xi, x \rangle})_{x \in Q_{2n+1}^d} : \xi \in \mathbb{Z}^d, \|\xi\|_\infty \leq n\}$, and the eigenvalue associated with the eigenvector $(e^{2\pi i \langle \xi, x \rangle})_{x \in Q_{2n+1}^d}$ is $(2n+1)^d \tilde{f}_{B_u}(\xi, 2n+1)$. Since M is a normal matrix, i.e. $M^\top M = M M^\top$, or, equivalently, since it is diagonalized by a system of orthogonal eigenvectors, its singular values are equal to the absolute values of its eigenvalues. Using Lemma 17 we have

the estimate

$$\begin{aligned}\gamma_2(M) &\geq \frac{1}{(2n+1)^d} \|M\|_{\text{tr}} \\ &= \sum_{\xi \in \mathbb{Z}^d: \|\xi\|_\infty \leq n} |\tilde{f}_{B_u}(\xi, 2n+1)| = \Omega_{d,B}(\log^d(n)).\end{aligned}\quad (14)$$

Let A be the incidence matrix of $\mathcal{T}_{B_u}(-Q_{2n+1})$. Notice that every row of M can be represented as the disjoint sum of at most 2^d rows of A . Therefore, we can write $M = A_1 + \cdots + A_k$, where $k \leq 2^d$, and each row of each matrix A_j is also a row of A . Since duplication and rearrangement of rows preserve γ_2 , and dropping rows does not increase it, by the triangle inequality for γ_2 we have

$$\gamma_2(M) \leq \sum_{j=1}^k \gamma_2(A_j) \leq 2^d \gamma_2(A) = 2^d \gamma_2(\mathcal{T}_{B_u}(-Q_{2n+1})).$$

Define the pointset $P = \{u^*(x-c) + c : x \in -Q_{2n+1}\}$, and notice that $\mathcal{T}_B(P) = \mathcal{T}_{B_u}(-Q_{2n+1})$, so, using the inequality above and (14), we have

$$\gamma_2(\mathcal{T}_B(P)) = \Omega_{d,B}(\log^d n). \quad (15)$$

Equations (2) and (15) imply

$$\text{disc}(n^d, \mathcal{T}_B) \geq \text{herdisc } \mathcal{T}_B(P) = \Omega_{d,B}(\log^{d-1} n).$$

Therefore, $\text{disc}(n, \mathcal{T}_B) = \Omega_{d,B}(\log^{d-1} n)$, as was to be proved. \square

Equation (15) implies the lower bound on $\sqrt{t_u t_q}$ in Theorem 5. The lower bound on error in Theorem 6 follows from equation (15) and Theorem 8.

Acknowledgements. This research was partially supported by NSERC. The author would like to thank the organizers of the 2016 discrepancy theory workshop in Varenna, dedicated to the memory of Klaus Roth and Jiří Matoušek, where some of the initial ideas in this paper were conceived. The author would also like to thank József Beck for suggesting the problem of extending the lower bound for Tusnády's problem to arbitrary polytopes, and Kunal Talwar for some initial discussions of this problem.

References

1. C. Aistleitner, D. Bilyk and A. Nikolov, Tusnády's problem, the transference principle, and non-uniform QMC sampling. *CoRR Preprint*, 2017, [arXiv:1703.06127](https://arxiv.org/abs/1703.06127).
2. R. Alexander, Principles of a new method in the study of irregularities of distribution. *Invent. Math.* **103**(2) (1991), 279–296.
3. W. Banaszczyk, Balancing vectors and Gaussian measures of n -dimensional convex bodies. *Random Structures Algorithms* **12**(4) (1998), 351–360.
4. W. Banaszczyk, On series of signed vectors and their rearrangements. *Random Structures Algorithms* **40**(3) (2012), 301–316.

5. N. Bansal, D. Dadush and S. Garg, An algorithm for Komlós conjecture matching Banaszczyk's bound. In *IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016, 9–11 October 2016, Hyatt Regency, New Brunswick, New Jersey, U.S.A.* (ed. I. Dinur), IEEE Computer Society (Washington, DC, 2016), 788–799.
6. N. Bansal and S. Garg, Algorithmic discrepancy beyond partial coloring. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing (STOC 2017)*, ACM (New York, NY, 2017).
7. J. Beck, Balanced two-colorings of finite sets in the square. I. *Combinatorica* **1** (1981), 327–335.
8. J. Beck and W. W. L. Chen, *Irregularities of Distribution (Cambridge Tracts in Mathematics 89)*, Cambridge University Press (Cambridge, 2008). Reprint of the 1987 original [MR 0903025].
9. D. Bilyk, M. T. Lacey and A. Vagharshakyan, On the small ball inequality in all dimensions. *J. Funct. Anal.* **254**(9) (2008), 2470–2502.
10. L. Brandolini, L. Colzani and G. Travaglini, Average decay of Fourier transforms and integer points in polyhedra. *Ark. Mat.* **35**(2) (1997), 253–275.
11. M. Charikar, A. Newman and A. Nikolov, Tight hardness results for minimizing discrepancy. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA, January 23–25, 2011* (ed. D. Randall), SIAM (Philadelphia, PA, 2011), 1607–1614.
12. B. Chazelle, Randomness and complexity. In *The Discrepancy Method*, Cambridge University Press (Cambridge, 2000).
13. M. Drmota, Irregularities of distributions with respect to polytopes. *Mathematika* **43**(1) (1996), 108–119.
14. C. Dwork, F. Mcsherry, K. Nissim and A. Smith, Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography (Lecture Notes in Computer Science 3876)*, Springer (Berlin, 2006), 265–284.
15. C. Dwork and A. Roth, The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.* **9**(3–4) (2014), 211–407.
16. C. L. Epstein, How well does the finite Fourier transform approximate the Fourier transform? *Comm. Pure Appl. Math.* **58**(10) (2005), 1421–1435.
17. M. L. Fredman, The complexity of maintaining an array and computing its partial sums. *J. ACM* **29**(1) (1982), 250–260.
18. M. Götz, Discrepancy and the error in integration. *Monatsh. Math.* **136**(2) (2002), 99–121.
19. J. H. Halton, On the efficiency of certain quasi-random sequences of points in evaluating multi-dimensional integrals. *Numer. Math.* **2** (1960), 84–90.
20. J. M. Hammersley, Monte Carlo methods for solving multivariable problems. *Ann. New York Acad. Sci.* **86** (1960), 844–874.
21. V. A. Judin, A multidimensional Jackson theorem. *Mat. Zametki* **20**(3) (1976), 439–444. Presented at the International Conference on the Theory of Approximation of Functions, held at Kaluga, July 24–28, 1975.
22. K. G. Larsen, On range searching in the group model and combinatorial discrepancy. *SIAM J. Comput.* **43**(2) (2014), 673–686.
23. T. Lee, A. Shraibman and R. Špalek, A direct product theorem for discrepancy. In *Proceedings of the 23rd Annual IEEE Conference on Computational Complexity, CCC 2008, 23–26 June 2008, College Park, Maryland, U.S.A.*, IEEE Computer Society (Washington, DC, 2008), 71–80.
24. N. Linial, S. Mendelson, G. Schechtman and A. Shraibman, Complexity measures of sign matrices. *Combinatorica* **27**(4) (2007), 439–463.
25. J. Matoušek, On the discrepancy for boxes and polytopes. *Monatsh. Math.* **127**(4) (1999), 325–336.
26. J. Matoušek, *Geometric Discrepancy (Algorithms and Combinatorics 18)*, Springer (Berlin, 2010). An illustrated guide, Revised paperback reprint of the 1999 original.
27. J. Matoušek and A. Nikolov, Combinatorial discrepancy for boxes via the γ_2 norm. In *31st International Symposium on Computational Geometry (SoCG 2015) (LIPIcs, Leibniz International Proceedings in Informatics 34)*, Schloss Dagstuhl–Leibniz Center for Informatics (2015), 1–15.
28. J. Matoušek, A. Nikolov and K. Talwar, Factorization norms and hereditary discrepancy. *CoRR Preprint*, 2015, [arXiv:1408.1376v2](https://arxiv.org/abs/1408.1376v2).
29. A. Nikolov, New computational aspects of discrepancy theory. *PhD Thesis*, Rutgers, The State University of New Jersey, 2014.
30. A. Nikolov and K. Talwar, Approximating hereditary discrepancy via small width ellipsoids. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, U.S.A., January 4–6, 2015* (ed. P. Indyk), SIAM (Philadelphia, PA, 2015), 324–336.

31. A. Nikolov, K. Talwar and L. Zhang, The geometry of differential privacy: the small database and approximate cases. *SIAM J. Comput.* **45**(2) (2016), 575–616.
32. K. F. Roth, On irregularities of distribution. *Mathematika* **1** (1954), 73–79.
33. K. F. Roth, Remark concerning integer sequences. *Acta Arith.* **9** (1964), 257–260.
34. W. M. Schmidt, Irregularities of distribution. VII. *Acta Arith.* **21** (1972), 45–50.
35. N. Tomczak-Jaegermann, *Banach-Mazur Distances and Finite-Dimensional Operator Ideals* (Pitman Monographs and Surveys in Pure and Applied Mathematics **38**), Wiley (New York, 1989).

Aleksandar Nikolov,
Department of Computer Science,
University of Toronto,
Toronto,
Canada
E-mail: anikolov@cs.toronto.edu