

Unit and distinct distances in typical norms

Noga Alon*

Matija Bucić[†]

Lisa Sauermann[‡]

Abstract

Erdős' unit distance problem and Erdős' distinct distances problem are among the most classical and well-known open problems in discrete mathematics. They ask for the maximum number of unit distances, or the minimum number of distinct distances, respectively, determined by n points in the Euclidean plane. The question of what happens in these problems if one considers normed spaces other than the Euclidean plane has been raised in the 1980s by Ulam and Erdős and attracted a lot of attention over the years. We give an essentially tight answer to both questions for almost all norms on \mathbb{R}^d , in a certain Baire categorical sense.

For the unit distance problem we prove that for almost all norms $\|\cdot\|$ on \mathbb{R}^d , any set of n points defines at most $\frac{1}{2}d \cdot n \log_2 n$ unit distances according to $\|\cdot\|$. We also show that this is essentially tight, by proving that for *every* norm $\|\cdot\|$ on \mathbb{R}^d , for any large n , we can find n points defining at least $\frac{1}{2}(d-1-o(1)) \cdot n \log_2 n$ unit distances according to $\|\cdot\|$.

For the distinct distances problem, we prove that for almost all norms $\|\cdot\|$ on \mathbb{R}^d any set of n points defines at least $(1-o(1))n$ distinct distances according to $\|\cdot\|$. This is clearly tight up to the $o(1)$ term.

We also answer the famous Hadwiger–Nelson problem for almost all norms on \mathbb{R}^2 , showing that their unit distance graph has chromatic number 4.

Our results settle, in a strong and somewhat surprising form, problems and conjectures of Brass, Matoušek, Brass–Moser–Pach, Chilakamarri, and Robertson. The proofs combine combinatorial and geometric ideas with tools from Linear Algebra, Topology and Algebraic Geometry.

1 Introduction

1.1 Unit distances

Erdős' unit distance problem raised in 1946 in [20] (see also [8, Chapter 5]) is among the best-known open problems in combinatorics. The problem asks about estimating the maximum possible number $U_{\|\cdot\|_2}(n)$ of unit distances determined by n distinct points in the Euclidean plane \mathbb{R}^2 according to the Euclidean norm $\|\cdot\|_2$. The best bounds to date are

$$n^{1+\Omega(1/\log \log n)} \leq U_{\|\cdot\|_2}(n) \leq O(n^{4/3}).$$

The lower bound appeared in the initial paper of Erdős [20], and the upper bound was first proved by Spencer, Szemerédi and Trotter [41] in 1984, see also [45] for a short and elegant proof based on the Crossing Lemma. More on the rich history of this problem can be found in the surveys [8, 46].

An interesting variant of the problem deals with the same question in general real normed spaces. This was first suggested by Ulam and described explicitly by Erdős in the early 1980s [19]. We call a norm $\|\cdot\|$ on \mathbb{R}^d a d -norm and denote by $U_{\|\cdot\|}(n)$ the maximum possible number of unit distances determined by a set of n distinct

*Department of Mathematics, Princeton University, Princeton, USA. Research supported in part by NSF grant DMS-2154082 and by USA-Israel BSF grant 2018267. Email: nalon@math.princeton.edu.

[†]School of Mathematics, Institute for Advanced Study and Department of Mathematics, Princeton University, Princeton, USA. Research supported in part by NSF grants CCF-1900460 and DMS-2349013. Email: matija.bucic@ias.edu.

[‡]Department of Mathematics, Massachusetts Institute of Technology, Cambridge, USA. Supported in part by NSF Award DMS-2100157 and a Sloan Research Fellowship. Email: lsauerma@mit.edu.

points in \mathbb{R}^d . The unit distance graph of a d -norm $\|\cdot\|$ is the graph whose vertices are all points of \mathbb{R}^d , where two points are adjacent if and only if the distance between them is one. Thus, $U_{\|\cdot\|}(n)$ is the maximum number of edges in an n -vertex subgraph of the unit distance graph of $\|\cdot\|$. An initial, simple observation is that if the boundary of the unit ball of a d -norm $\|\cdot\|$ contains a straight-line segment then $U_{\|\cdot\|}(n)$ is quadratic in n , as in this case there are two infinite subsets A, B of \mathbb{R}^d so that the distance between any $\mathbf{a} \in A$ and $\mathbf{b} \in B$ is one. On the other hand, if $\|\cdot\|$ is a 2-norm which is strictly convex (meaning that the boundary of the unit ball contains no straight line segment), then one can extend the known proofs from the Euclidean case to show that $U_{\|\cdot\|}(n) \leq O(n^{4/3})$. Valtr [48] constructed a strictly convex 2-norm in which for every n there exist n -element point sets with at least $\Omega(n^{4/3})$ unit distances. See also [39] for a much more general family of norms with the same property. This shows that the upper bound cannot be improved in general. See [8] for more details about the history of the problem of estimating $U_{\|\cdot\|}(n)$ for general norms.

It is not difficult to see that for any 2-norm $\|\cdot\|$, we have $U_{\|\cdot\|}(n) \geq (\frac{1}{2} - o(1))n \log_2 n$. Indeed, the graph of the k -dimensional hypercube is a subgraph of the unit distance graph of any 2-norm $\|\cdot\|$ as shown by choosing k random unit vectors $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_k \in \mathbb{R}^2$ and defining $\mathbf{v}_S = \sum_{i \in S} \mathbf{u}_i$ for every subset $S \subseteq \{1, 2, \dots, k\}$. If two subsets $S, S' \subseteq \{1, 2, \dots, k\}$ differ in a single element, then the distance according to $\|\cdot\|$ between \mathbf{v}_S and $\mathbf{v}_{S'}$ is one, as desired.

A remarkable, often repeated (see e.g. [8, 30, 33, 38, 43, 46]), result of Matoušek [34] from 2011 shows that for *most* norms, in a Baire Categorical sense (that will be described in detail in Section 3), this is not far from optimal. Namely for a typical 2-norm $\|\cdot\|$, $U_{\|\cdot\|}(n) \leq O(n \log n \log \log n)$. This suggests the obvious problem of deciding whether or not the $\log \log n$ term is necessary. The question of how small $U_{\|\cdot\|}(n)$ can be has already been considered by Brass [6] in 1996, who in particular asked if there exists a 2-norm with $U_{\|\cdot\|}(n) \leq O(n \log n)$. The corresponding problem in higher dimensions has been considered as well. In particular, Brass, Moser and Pach [8, Chapter 5, Problems 4 and 5, p. 195] conjectured that for every $d \geq 3$ and every d -norm $\|\cdot\|$, $U_{\|\cdot\|}(n)$ is asymptotically larger than $\Omega(n \log n)$ and asked whether or not for $d \geq 4$ there is a d -norm $\|\cdot\|$ so that $U_{\|\cdot\|}(n) = o(n^2)$. Note that for the d -dimensional Euclidean norm $\|\cdot\|_2$ it is easy to see that $U_{\|\cdot\|}(n) \geq \Omega(n^2)$ for every $d \geq 4$ (and in fact, the precise constant is known as well, see [18]), showing that for the Euclidean norms, the problem is simple in all dimensions $d \geq 4$, despite being wide open in dimensions 2 and 3.

In the present paper, we settle the above-mentioned questions of Brass, of Matoušek and of Brass, Moser and Pach in all dimensions in the following strong form.

Theorem 1.1. *For any $d \geq 2$, for most d -norms $\|\cdot\|$, we have*

$$U_{\|\cdot\|}(n) \leq \frac{d}{2} \cdot n \log_2 n$$

for all n . More precisely, for all d -norms $\|\cdot\|$ besides a meagre set, the following holds: For every $n \geq 1$ and every set of n points in \mathbb{R}^d , there are at most $\frac{d}{2} \cdot n \log_2 n$ unit distances according to $\|\cdot\|$ among the n points.

Theorem 1.2. *Let $d \geq 2$ be fixed. Then for every d -norm $\|\cdot\|$, we have*

$$U_{\|\cdot\|}(n) \geq \frac{d-1-o(1)}{2} \cdot n \log_2 n.$$

for all large n . That is, for every d -norm $\|\cdot\|$ and every n , there exists a set of n points in \mathbb{R}^d such that the number of unit distances according to $\|\cdot\|$ among the n points is at least $\frac{d-1-o(1)}{2} \cdot n \log_2 n$, where the $o(1)$ -term tends to zero as $n \rightarrow \infty$.

1.2 Distinct distances

The equally famous Erdős distinct distances problem is concerned with estimating the minimum possible number $D_{\|\cdot\|_2}(n)$ of distinct distances determined by n points in the Euclidean plane \mathbb{R}^2 . See for example the book [26] focusing on this question, its history, generalisations and connections to other areas. The Erdős distinct distances problem also traces its origins to the 1946 paper [20] of Erdős. In this paper, he considered the same problem for higher-dimensional Euclidean spaces as well. For the planar case, Erdős proved an upper bound of $O(n/\sqrt{\log n})$ and conjectured this is tight. After a long sequence of improvements of the lower bound, Guth and Katz established in [29] a nearly tight lower bound of $\Omega(n/\log n)$. For higher-dimensional Euclidean spaces even the correct exponent of n is not known, see [40] for the best known bounds.

The Erdős distinct distances problem has a long history in the case of general d -norms as well, see [43] for a recent survey on what is known in this direction. Given a d -norm $\|\cdot\|$ let us denote by $D_{\|\cdot\|}(n)$ the minimum possible number of distinct distances, according to $\|\cdot\|$, determined by n points in \mathbb{R}^d . For every d -norm $\|\cdot\|$ and every n , we clearly have $D_{\|\cdot\|}(n) \leq n - 1$ by considering any set of n points along an arithmetic progression on a line. Brass conjectured that $D_{\|\cdot\|}(n) = o(n)$ for any $d \geq 2$ and any d -norm $\|\cdot\|$ (see [8, Chapter 5.4, Conjecture 5, p. 211]). Here we refute his conjecture in a strong form. For most d -norms $\|\cdot\|$, we show that $D_{\|\cdot\|}(n)$ is not only linear in n , but is in fact of the form $(1 - o(1))n$.

Theorem 1.3. *For any fixed $d \geq 2$, for most d -norms $\|\cdot\|$ we have*

$$D_{\|\cdot\|}(n) = (1 - o(1))n$$

for all n . More precisely, for all d -norms $\|\cdot\|$ besides a meagre set, the following holds: For every n , among any n points in \mathbb{R}^d there are at least $(1 - o(1))n$ distinct distances according to $\|\cdot\|$, where the $o(1)$ -term tends to zero as $n \rightarrow \infty$.

1.3 Hadwiger–Nelson Problem

The question of determining the chromatic number of the unit distance graph of the Euclidean norm in the plane is yet another famous open problem in discrete geometry, known as the Hadwiger–Nelson problem. In other words, this question asks how many colours are needed in order to colour all points in the plane such that any two points with Euclidean distance one receive different colours. This problem dates back to 1950, and it has been known for a long time that the answer is at least 4 and at most 7. In a recent computational breakthrough [11, 22], the lower bound has been improved to 5, sparking a collaborative Polymath project [37] focusing on the problem. See also [8, 43] for more details on the history of the Hadwiger–Nelson problem.

The analogue of this problem for other planar norms was studied by Chilakamarri [9] in 1991, who showed that the unit distance graph of every 2-norm has chromatic number at least 4 and at most 7. In addition, [9, Problem 5], attributed to Robertson, asks for the chromatic number of the unit distance graph of at least one strictly convex 2-norm to be evaluated. We prove that the chromatic number of the unit distance graph of most 2-norms equals 4.

Theorem 1.4. *For all 2-norms $\|\cdot\|$ besides a meagre set, the unit distance graph of $\|\cdot\|$ has chromatic number equal to 4.*

In contrast, as mentioned above, it was recently shown [11, 22] that for the Euclidean 2-norm the chromatic number of the unit distance graph is at least 5. Thus, the behaviour for most 2-norms is different from the Euclidean 2-norm, which in particular disproves a conjecture of Chilakamarri [9, p. 355].

In dimension d , our arguments give an upper bound of 2^d for the chromatic number of the unit distance graph of most d -norms. It is known that the chromatic number of the unit distance graph of any d -norm is at most exponential in d (see [25, 32]), but the exponential base in these known bounds is much larger than 2. See also [43] for a more detailed survey on what is known surrounding this problem. We furthermore note that results of Frankl and Wilson [23] give an exponential lower bound for the chromatic number of the unit distance graph of all d -norms which are invariant under coordinate permutations.

Our upper bound of 2^d for most d -norms actually also holds for the chromatic number of the odd distance graph (i.e. the graph whose edges correspond to pairs of points in \mathbb{R}^d whose distance is an odd integer), which is a stronger result as the unit distance graph is a subgraph of the odd distance graph. This is very different from the Euclidean case, where by a recent result of Davies [10] already in the plane the chromatic number of the odd distance graph is infinite.

2 Proof overview

The proofs of our theorems use arguments from combinatorics, polyhedral and discrete geometry, topology and algebra. In this section, we give a high-level proof overview. We note that this overview is simplifying numerous parts and omitting certain important points, since the goal of this section is to present just the general ideas and how they come together. We begin with Theorem 1.1.

The proof of Theorem 1.1 splits into two parts. The first, more combinatorial part shows that for a norm with a lot of unit distances among some set of n points there must be a lot of rational linear dependencies between the corresponding unit vectors. More precisely, if the bound in Theorem 1.1 fails for some d -norm $\|\cdot\|$, namely if we can find n points with more than $(d/2) \cdot n \log_2 n$ unit distances, then there is some set of unit vectors whose \mathbb{Q} -linear span contains a lot of other unit vectors. The second, more geometric part consists of showing that for most norms, this situation cannot happen. In other words, for most norms, there cannot be a collection of unit vectors from which a lot of other unit vectors can be obtained as rational linear combinations.

In the first part, we fix n points with more than $(d/2) \cdot n \log_2 n$ unit distances according to some d -norm $\|\cdot\|$. Our goal is to show that there must be a lot of rational linear dependencies between the vectors $\mathbf{u}_1, \dots, \mathbf{u}_k$ with $\|\mathbf{u}_1\| = \dots = \|\mathbf{u}_k\| = 1$ describing these unit distances. If there are not so many rational linear dependencies (more precisely, if for every ℓ , the \mathbb{Q} -linear span of any ℓ of the vectors $\mathbf{u}_1, \dots, \mathbf{u}_k$ contains at most $d \cdot \ell$ of the vectors $\mathbf{u}_1, \dots, \mathbf{u}_k$), then we can find a subset of these unit vectors that is linearly independent over \mathbb{Q} and accounts for a relatively large number of unit distances among the n fixed points (namely, for more than $(1/2) \cdot n \log_2 n$ unit distances). On the other hand, for n fixed points, we can analyse for how many pairs of points their difference can be among a given list of \mathbb{Q} -linearly independent vectors, and it turns out that this can happen for at most $(1/2) \cdot n \log_2 n$ pairs of points (this follows from an isoperimetric inequality for grids due to Bollobás and Leader [5], but we also give a self-contained proof). This contradiction means that our original set of unit vectors $\mathbf{u}_1, \dots, \mathbf{u}_k$ must have many rational linear dependencies, as desired.

For the second part, we need to show that there are only “few” special d -norms with the property that for some collection of unit vectors, a lot of other unit vectors can be obtained as rational linear combinations. So we need to show that the set of d -norms with this property can be covered by countably many nowhere-dense subsets in the space of d -norms (with a certain natural topology, see Section 3 for more details). Roughly speaking, our countably many subsets correspond to the different possibilities for the size of the collection of unit vectors and the rational coefficients in the relevant rational linear combinations, as well as a lower bound for the separation angle between any two of these unit vectors. To prove that each of these subsets is nowhere dense in the space of all d -norms, we show that for any d -norm whose unit ball is a convex polytope with sufficiently small facets, there is a slight perturbation (obtained by slightly translating each facet hyperplane)

that does not appear as the unit ball of an accumulation point of the subset. Heuristically speaking, if such a perturbation appears as the unit ball of a norm from this subset, then certain linear equations need to be satisfied for the translation lengths for the facets containing the unit vectors appearing in the above property (as a consequence of the linear relations between the unit vectors themselves). A generic perturbation does not satisfy these linear equations and therefore cannot correspond to an accumulation point of norms from this subset. This establishes that each of these subsets is indeed nowhere dense, and concludes the proof of Theorem 1.1.

The proof of Theorem 1.3 follows a similar overall approach, but various steps become more involved and significant new ideas are required. In the first part, for n given points with few distinct distances according to some d -norm, we work over the field extension of \mathbb{Q} generated by the distinct distances between these n points (so we consider linear relations with coefficients in this extension field). We also need to establish stronger quantitative bounds for the number of unit vectors that we find in the span of some collection of unit vectors. Overall, the first part of the proof of Theorem 1.3 is much more involved than for Theorem 1.1. For example, one cannot use the Bollobás–Leader edge-isoperimetric inequality anymore in this setting, since here we need a more general statement (see Lemma 4.3) compared to the setting of Theorem 1.1 (we prove this more general statement via combinatorial arguments). The second part of the proof of Theorem 1.3 is relatively similar to Theorem 1.1, but requires some basic tools from algebraic geometry instead of the linear algebraic arguments in the proof of Theorem 1.1.

For our result on the Hadwiger–Nelson problem for typical norms in Theorem 1.4, we combine the second part of the proof of Theorem 1.1 with a modified version of the first part. In particular, we make crucial use of the Edmonds Matroid Partitioning Theorem [14].

We note that the general idea of trying to control the number of linear dependencies between unit vectors in a typical norm is due to Matoušek [34], and his proof of the upper bound $U_{\|\cdot\|}(n) \leq O(n \log n \log \log n)$ for most 2-norms $\|\cdot\|$ similarly splits into two parts as our proof of Theorem 1.1. For the second part, showing that for most norms there cannot be many linear dependencies between unit vectors, our approach is inspired by his arguments. However, he uses hands-on geometric arguments in dimension 2 about “bulging” line segments or decomposing into trapezoids, that do not seem to easily extend to higher dimensions. Therefore, to obtain our results in any dimension, we introduced several new ideas, in particular making the proof more algebraic. On the other hand, the first part of our proof is radically different from Matoušek’s. His proof is based on graph theoretic “expansion” arguments via probabilistic methods, leading to his weaker bound. We develop a completely new strategy for the first part of the proof instead, resulting in an essentially tight bound.

To prove our new lower bound on $U_{\|\cdot\|}(n)$ for any d -norm $\|\cdot\|$ in Theorem 1.2, we consider Minkowski sums of certain carefully chosen small point sets, such that in a grid-like fashion we obtain many unit distances. To ensure that we can find suitable small point configurations serving as the base sets of our construction, we use a result from dimension theory: the Hurewicz dimension lowering theorem.

Organisation. The next section contains some background and preliminary lemmas. The first part of the proofs of Theorems 1.1 and 1.3 can be found in Section 4, and the second part in Section 5. Theorem 1.4 is then proved in Section 6, relying on the previous sections. Finally, we prove Theorem 1.2 in Section 7 and finish with some concluding remarks and open problems in Section 8.

3 Geometric preliminaries

3.1 Background

We begin by setting up some notation and introducing the notions we will work with. We note that all our logarithms are in base 2 unless specified otherwise.

For $d \geq 1$, a norm on \mathbb{R}^d is a mapping $\|\cdot\|$ that assigns a non-negative real number $\|\mathbf{x}\|$ to each $\mathbf{x} \in \mathbb{R}^d$ such that the following three conditions hold:

- For every $\mathbf{x} \in \mathbb{R}^d$, we have $\|\mathbf{x}\| = 0$ if and only if $\mathbf{x} = \mathbf{0}$.
- We have $\|\alpha\mathbf{x}\| = |\alpha| \cdot \|\mathbf{x}\|$ for all $\alpha \in \mathbb{R}$ and $\mathbf{x} \in \mathbb{R}^d$.
- The triangle inequality holds, meaning that $\|\mathbf{x} + \mathbf{y}\| \leq \|\mathbf{x}\| + \|\mathbf{y}\|$ for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$.

Each norm $\|\cdot\|$ on \mathbb{R}^d is uniquely specified by its *unit ball*, defined as the set of all $\mathbf{x} \in \mathbb{R}^d$ for which $\|\mathbf{x}\| \leq 1$. A unit ball of any norm is a closed, bounded, $\mathbf{0}$ -symmetric convex body containing $\mathbf{0}$ in its interior. Furthermore, any such body appears as the unit ball of a unique norm. Let \mathcal{B}_d denote the set of all unit balls of norms in \mathbb{R}^d or equivalently the set of all closed, bounded, $\mathbf{0}$ -symmetric convex bodies in \mathbb{R}^d containing $\mathbf{0}$ in the interior. As discussed below, it is known that \mathcal{B}_d endowed with the so-called Hausdorff metric d_H forms a Baire space.

The Hausdorff distance $d_H(A, B)$ of two sets $A, B \subseteq \mathbb{R}^d$ is defined as

$$d_H(A, B) := \max \left\{ \sup_{\mathbf{a} \in A} \inf_{\mathbf{b} \in B} \|\mathbf{a} - \mathbf{b}\|_2, \sup_{\mathbf{b} \in B} \inf_{\mathbf{a} \in A} \|\mathbf{a} - \mathbf{b}\|_2 \right\},$$

where $\|\cdot\|_2$ denotes the Euclidean distance in \mathbb{R}^d . If $A, B \subseteq \mathbb{R}^d$ are closed and bounded, then one can replace the suprema and infima in the above definition with maxima and minima. So in this case the Hausdorff distance $d_H(A, B)$ is simply the “maximum distance” of a point in A from the set B or of a point in B from the set A . Note that the Hausdorff distance satisfies the triangle inequality.

A set S in a metric (or topological) space X is nowhere dense if every non-empty open set $U \subseteq X$ contains a nonempty open set V with $V \cap S = \emptyset$. A meagre set in X is a countable union of nowhere dense sets. Note that a subset of any meagre set is also meagre. The space X is called a Baire space if the complement of each meagre set in X is dense. It is known that \mathcal{B}_d endowed with the Hausdorff metric d_H forms a Baire space (this follows for example from [28, Theorem 6.4] together with the Baire Category Theorem).

The diameter of a non-empty bounded closed subset $S \subseteq \mathbb{R}^d$ (with respect to the Euclidean distance) is defined as

$$\text{diam}(S) = \max_{\mathbf{a}, \mathbf{b} \in S} \|\mathbf{a} - \mathbf{b}\|_2$$

(note that this maximum is indeed well-defined since S is closed and bounded).

A half-space in \mathbb{R}^d is the closed subset of \mathbb{R}^d given by the solutions $\mathbf{x} \in \mathbb{R}^d$ to some linear inequality of the form $\mathbf{a} \cdot \mathbf{x} \leq b$ for some $\mathbf{a} \in \mathbb{R}^d$ and some $b \in \mathbb{R}$ (geometrically, this is the set of points on one side of the affine hyperplane given by $\mathbf{a} \cdot \mathbf{x} = b$, including the hyperplane itself). A polytope $P \subseteq \mathbb{R}^d$ is an intersection of finitely many half-spaces in \mathbb{R}^d . Note that every polytope P is convex. Every bounded polytope $P \subseteq \mathbb{R}^d$ can also be described as a convex hull of finitely many points in \mathbb{R}^d .

Every $\mathbf{0}$ -symmetric polytope P containing $\mathbf{0}$ in its interior can be written in the form

$$P = \{\mathbf{x} \in \mathbb{R}^d : |\mathbf{o}_i \cdot \mathbf{x}| \leq t_i \text{ for } i = 1, \dots, h\}$$

with non-zero vectors $\mathbf{o}_1, \dots, \mathbf{o}_h \in \mathbb{R}^d$ and positive real numbers t_1, \dots, t_h . The facets of such a polytope are the intersections of the form $P \cap H$ for some hyperplane H such that the intersection $P \cap H$ is $(d-1)$ -dimensional

and P is contained in one of the closed half-spaces bounded by H . If P is $\mathbf{0}$ -symmetric, then the facets appear in pairs of opposite facets (which are parallel to each other).

A set $B \subseteq \mathbb{R}^d$ is strictly convex, if for all distinct $\mathbf{a}, \mathbf{b} \in B$ and all $0 < \alpha < 1$ the point $\alpha\mathbf{a} + (1 - \alpha)\mathbf{b}$ is in the interior of B . For a strictly convex set $B \subseteq \mathbb{R}^d$, for every point \mathbf{b} on the boundary of B there exists a hyperplane H with $H \cap B = \{\mathbf{b}\}$ such that B is contained in one of the half-spaces bounded by H . The unit ball B of a norm $\|\cdot\|$ on \mathbb{R}^d is strictly convex if the triangle inequality is a strict inequality $\|\mathbf{x} + \mathbf{y}\| < \|\mathbf{x}\| + \|\mathbf{y}\|$ for all non-zero vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$ that are not multiples of each other. Indeed, in this case for non-zero $\mathbf{a}, \mathbf{b} \in B$ with $\text{span}_{\mathbb{R}}(\mathbf{a}) \neq \text{span}_{\mathbb{R}}(\mathbf{b})$ and $0 < \alpha < 1$, we have $\|\alpha\mathbf{a} + (1 - \alpha)\mathbf{b}\| < \|\alpha\mathbf{a}\| + \|(1 - \alpha)\mathbf{b}\| = \alpha\|\mathbf{a}\| + (1 - \alpha)\|\mathbf{b}\| \leq \alpha + (1 - \alpha) = 1$. For distinct vectors $\mathbf{a}, \mathbf{b} \in B$ with $\text{span}_{\mathbb{R}}(\mathbf{a}) = \text{span}_{\mathbb{R}}(\mathbf{b})$ we always have $\|\alpha\mathbf{a} + (1 - \alpha)\mathbf{b}\| < 1$ since we either have $\max\{\|\mathbf{a}\|, \|\mathbf{b}\|\} < 1$ or $\mathbf{a} = -\mathbf{b}$. And for $\mathbf{a}, \mathbf{b} \in B$ with $\mathbf{a} = \mathbf{0}$ or $\mathbf{b} = \mathbf{0}$ we also trivially have $\|\alpha\mathbf{a} + (1 - \alpha)\mathbf{b}\| < 1$. Thus the triangle inequality for $\|\cdot\|$ being strict for all non-zero vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$ with $\text{span}_{\mathbb{R}}(\mathbf{x}) \neq \text{span}_{\mathbb{R}}(\mathbf{y})$ indeed implies that B is strictly convex.

Finally, we record the following simple algebraic fact, which we will use in our proof of Theorem 1.3.

Fact 3.1. For any positive integer m , given $m+1$ rational functions $f_1, \dots, f_{m+1} \in \mathbb{R}(x_1, \dots, x_m)$ in m variables with real coefficients, there exists a nonzero polynomial $P \in \mathbb{R}[y_1, \dots, y_{m+1}]$ such that $P(f_1, \dots, f_{m+1}) = 0$.

This follows immediately from the well-known fact that the transcendence degree of the function field $\mathbb{R}(x_1, \dots, x_m)$ over \mathbb{R} is equal to m . Indeed, since $m + 1$ is larger than this transcendence degree, there must be an algebraic relationship between f_1, \dots, f_{m+1} .

3.2 Geometric lemmas

This section contains some basic geometric lemmas. Although most of the content of this section is well-known, we include some of the proofs for the reader's convenience.

The first two lemmas below are only needed to prove the third lemma in this section.

Lemma 3.2. *Let $T \subseteq \mathbb{R}^d$ be a bounded subset and let $\varepsilon > 0$. Then there exists a finite subset $S \subseteq T$ such that for every point $\mathbf{t} \in T$ there exists a point $\mathbf{s} \in S$ with $\|\mathbf{s} - \mathbf{t}\|_2 \leq \varepsilon$.*

We remark that such a subset S is called an ε -net of T .

Proof. Since T is bounded, it is contained in the Euclidean ball of radius R around $\mathbf{0}$ for some $R > 0$. Let us consider the family of all subsets $S \subseteq T$ with the property that $\|\mathbf{s} - \mathbf{s}'\|_2 > \varepsilon$ for all distinct $\mathbf{s}, \mathbf{s}' \in S$. We claim that each such subset S has size $|S| \leq (2R/\varepsilon + 1)^d$. Indeed, the Euclidean balls of radius $\varepsilon/2$ around the points in S are mutually disjoint and contained in the Euclidean ball of radius $R + \varepsilon/2$ around $\mathbf{0}$. Hence, for volume reasons, the set S can consist of at most $(2R/\varepsilon + 1)^d$ points.

Note that $S = \emptyset$ vacuously satisfies the property that $\|\mathbf{s} - \mathbf{s}'\|_2 > \varepsilon$ for all distinct $\mathbf{s}, \mathbf{s}' \in S$. Hence, the family of subsets $S \subseteq T$ with this property is non-empty, and as all its members satisfy $|S| \leq (2R/\varepsilon + 1)^d$, there must be a maximal subset $S \subseteq T$ in this family. So let $S \subseteq T$ be a maximal subset with the property that we have $\|\mathbf{s} - \mathbf{s}'\|_2 > \varepsilon$ for all distinct $\mathbf{s}, \mathbf{s}' \in S$.

Now, let us check that for every point $\mathbf{t} \in T$ there exists a point $\mathbf{s} \in S$ with $\|\mathbf{s} - \mathbf{t}\|_2 \leq \varepsilon$. If $\mathbf{t} \in S$, we can choose $\mathbf{s} = \mathbf{t}$. If $\mathbf{t} \notin S$, then by the maximality of our chosen set S we cannot add \mathbf{t} to the set. This means that we must have $\|\mathbf{s} - \mathbf{t}\|_2 \leq \varepsilon$ for some $\mathbf{s} \in S$, as desired. \square

The following lemma states that close to any $B \in \mathcal{B}_d$ we can find some strictly convex B' . This lemma can also be deduced from a classical result of Klee [31] from 1959 which says that almost all norms on \mathbb{R}^d are strictly convex.

Lemma 3.3. *For every $B \in \mathcal{B}_d$ and every $\mu > 0$, there exists a strictly convex $B' \in \mathcal{B}_d$ such that $d_H(B, B') \leq \mu$.*

Proof. Let us denote by $\|\cdot\|$ the norm with unit ball B . Then $B = \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x}\| \leq 1\}$. Since B is bounded, we can choose some $c > 0$ such that $\|\mathbf{x}\|_2 \leq c$ for all $\mathbf{x} \in B$. Let $\varepsilon = \mu/c^2$.

Let us define a norm $\|\cdot\|'$ by $\|\mathbf{x}\|' := \|\mathbf{x}\| + \varepsilon\|\mathbf{x}\|_2$ for all $\mathbf{x} \in \mathbb{R}^d$. To check that $\|\cdot\|'$ is indeed a norm, note that the triangle inequalities for $\|\cdot\|$ and $\|\cdot\|_2$ imply that

$$\|\mathbf{x} + \mathbf{y}\|' = \|\mathbf{x} + \mathbf{y}\| + \varepsilon\|\mathbf{x} + \mathbf{y}\|_2 \leq \|\mathbf{x}\| + \|\mathbf{y}\| + \varepsilon\|\mathbf{x}\|_2 + \varepsilon\|\mathbf{y}\|_2 = \|\mathbf{x}\|' + \|\mathbf{y}\|',$$

for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$, so the triangle inequality also holds for $\|\cdot\|'$. Note in addition that we can have equality only if $\|\mathbf{x} + \mathbf{y}\|_2 = \|\mathbf{x}\|_2 + \|\mathbf{y}\|_2$ and since the triangle inequality is strict for $\|\cdot\|_2$ for any non-zero vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$ which are not multiples of each other, the same also holds for $\|\cdot\|'$. Thus, the unit ball B' of the norm $\|\cdot\|'$ is strictly convex.

Note that $B' \subseteq B$, so for any point $\mathbf{b}' \in B'$ we have a point $\mathbf{b} = \mathbf{b}' \in B$ at Euclidean distance $0 \leq \mu$ from \mathbf{b}' . On the other hand, we claim that for any point $\mathbf{b} \in B$ we can also find a point $\mathbf{b}' \in B'$ with Euclidean distance at most μ from \mathbf{b} . If $\|\mathbf{b}\|_2 \leq \mu$, we can simply take $\mathbf{b}' = \mathbf{0} \in B'$. If $\|\mathbf{b}\|_2 > \mu$, we have

$$\left\| \left(1 - \frac{\mu}{\|\mathbf{b}\|_2}\right) \mathbf{b} \right\|' = \left(1 - \frac{\mu}{\|\mathbf{b}\|_2}\right) \|\mathbf{b}\|' = \left(1 - \frac{\mu}{\|\mathbf{b}\|_2}\right) (\|\mathbf{b}\| + \varepsilon\|\mathbf{b}\|_2) \leq (1 - \varepsilon\|\mathbf{b}\|_2)(1 + \varepsilon\|\mathbf{b}\|_2) < 1,$$

using $\varepsilon = \mu/c^2 \leq \mu/(\|\mathbf{b}\|_2)^2$ and $\|\mathbf{b}\| \leq 1$. So we can take $\mathbf{b}' := \left(1 - \frac{\mu}{\|\mathbf{b}\|_2}\right) \mathbf{b} \in B'$ and have $\|\mathbf{b}' - \mathbf{b}\|_2 = \frac{\mu}{\|\mathbf{b}\|_2} \cdot \|\mathbf{b}\|_2 = \mu$. This shows that $d_H(B, B') \leq \mu$, as desired. \square

We will need the following lemma which tells us that we can approximate any $B \in \mathcal{B}_d$ with a polytope in \mathcal{B}_d with small facets. We note there is plenty of research concerned with similar polytope approximation problems, see e.g. [4] and references therein, mostly concerned with minimising the number of facets needed to obtain a good approximation. Here, we are not concerned with the number of facets of the approximating polytope, but we need all of the facets to have small diameter.

Lemma 3.4. *For every $B \in \mathcal{B}_d$ and every $\mu > 0$, there exists a bounded $\mathbf{0}$ -symmetric polytope $B' \in \mathcal{B}_d$ containing $\mathbf{0}$ in its interior such that $d_H(B, B') < \mu$ and all facets of B' have diameter less than μ (with respect to the Euclidean distance).*

Note that we can endow the set of all $(d-1)$ -dimensional hyperplanes in \mathbb{R}^d with the natural topology induced by \mathbb{R}^{d+1} when identifying a hyperplane described by an equation of the form $\mathbf{a} \cdot \mathbf{x} = b$ with $\|\mathbf{a}\|_2 = 1$ with the point (\mathbf{a}, b) . More precisely, such a hyperplane with an equation of the form $\mathbf{a} \cdot \mathbf{x} = b$ corresponds to a point in the quotient space $(\mathbf{a}, b) \in (S^{d-1} \times \mathbb{R})/\{\pm 1\}$ (here S^{d-1} is the $(d-1)$ -dimensional unit sphere in \mathbb{R}^d and we consider the quotient by the action of $\{\pm 1\}$ on $(\mathbf{a}, b) \in S^{d-1} \times \mathbb{R}$, where -1 acts by sending $(\mathbf{a}, b) \mapsto (-\mathbf{a}, -b)$).

We denote by $\text{dist}_2(\mathbf{p}, H)$ the (Euclidean) distance from a point $\mathbf{p} \in \mathbb{R}^d$ to a hyperplane $H \subseteq \mathbb{R}^d$.

Proof. By Lemma 3.3, there exists a strictly convex $B'' \in \mathcal{B}_d$ with $d_H(B'', B) < \mu/2$. It is now sufficient to show that there exists a bounded $\mathbf{0}$ -symmetric polytope $B' \in \mathcal{B}_d$ containing $\mathbf{0}$ in its interior such that $d_H(B', B'') < \mu/2$ and all facets of B' have diameter less than μ

Let \mathcal{H} be the set of all $(d-1)$ -dimensional hyperplanes H in \mathbb{R}^d such that $H \cap B''$ has diameter at least μ . Note that \mathcal{H} is a closed and bounded subset of the set of all $(d-1)$ -dimensional hyperplanes in \mathbb{R}^d (to see that \mathcal{H} is closed, note that $\text{diam}(H \cap B'')$ is a continuous function on the set of all $(d-1)$ -dimensional hyperplanes $H \subseteq \mathbb{R}^d$). Thus, \mathcal{H} is compact in the induced topology.

Every hyperplane $H \in \mathcal{H}$ cuts the strictly convex set B'' into two parts $B''_1(H)$ and $B''_2(H)$ (strictly speaking these are the intersections of B'' with the two closed half-spaces bounded by H), neither of which is contained in H , by the strict convexity assumption on B'' . For every hyperplane $H \in \mathcal{H}$, define

$$f(H) = \min \left\{ \max_{\mathbf{b} \in B''_1(H)} \text{dist}_2(\mathbf{b}, H), \max_{\mathbf{b} \in B''_2(H)} \text{dist}_2(\mathbf{b}, H) \right\},$$

and note that $f(H) > 0$ for every $H \in \mathcal{H}$. Now f is a continuous function on a compact space and so it attains a minimum. So let $\eta > 0$ be this minimum, then $f(H) \geq \eta$ for all $H \in \mathcal{H}$.

Now, let us apply Lemma 3.2 to the set B'' and $\varepsilon = \min\{\mu, \eta\}/2$, i.e. let us choose an ε -net $S \subseteq B''$ of B'' . By adding up to $d+1$ additional points to S , we may assume that $\mathbf{0}$ is in the interior of the convex hull of S .

We may furthermore assume that S is $\mathbf{0}$ -symmetric, meaning that $-S = S$ (indeed, for every point $\mathbf{s} \in S$ we may add the point $-\mathbf{s}$ to S if it is not already contained in S).

Now, for every hyperplane $H \in \mathcal{H}$ there exist points in S on both sides of H (more precisely, the two open half-spaces bounded by H both contain at least one point in S). Indeed, since $f(H) \geq \eta$, there exist points $\mathbf{b}_1 \in B''_1(H) \subseteq B''$ and $\mathbf{b}_2 \in B''_2(H) \subseteq B''$ with $\text{dist}_2(\mathbf{b}_1, H) \geq \eta$ and $\text{dist}_2(\mathbf{b}_2, H) \geq \eta$. Note that \mathbf{b}_1 and \mathbf{b}_2 are on opposite sides of the hyperplane H (and not on H itself). Now, we can choose points $\mathbf{s}_1, \mathbf{s}_2 \in S$ with $\|\mathbf{s}_1 - \mathbf{b}_1\| \leq \varepsilon \leq \eta/2$ and $\|\mathbf{s}_2 - \mathbf{b}_2\| \leq \varepsilon \leq \eta/2$. Then \mathbf{s}_1 must be on the same side of H as \mathbf{b}_2 (and not on H itself), and similar for \mathbf{s}_2 . Thus, \mathbf{s}_1 and \mathbf{s}_2 lie on opposite sides of H in the two open half-spaces bounded by H . So indeed for every $H \in \mathcal{H}$ the two open half-spaces bounded by H both contain at least one point in S .

Finally, let us define $B' = \text{conv}(S)$ to be the convex hull of S . By our assumptions on S , the set B' is a bounded $\mathbf{0}$ -symmetric polytope containing $\mathbf{0}$ in its interior. In particular, B' is d -dimensional

To check that $d_H(B'', B') \leq \mu/2$, first note that $B' = \text{conv}(S) \subseteq \text{conv}(B'') = B''$ (as B'' is a convex set). So we have $\sup_{\mathbf{b}' \in B'} \inf_{\mathbf{b} \in B''} \|\mathbf{b} - \mathbf{b}'\|_2 = 0 < \mu$. Furthermore, for every $\mathbf{b} \in B''$ there is a point $\mathbf{b}' \in S \subseteq B'$ with $\|\mathbf{b} - \mathbf{b}'\|_2 \leq \varepsilon \leq \mu/2$ and hence $\sup_{\mathbf{b} \in B''} \inf_{\mathbf{b}' \in B'} \|\mathbf{b} - \mathbf{b}'\|_2 \leq \mu/2$. This shows that indeed $d_H(B'', B') \leq \mu/2$.

Finally, it remains to check that the facets of B' all have diameter less than μ . So suppose that B' had a facet of diameter at least μ . Then for the $(d-1)$ -dimensional hyperplane H through this facet, the set $H \cap B''$ (which contains this facet) has diameter at least μ . Hence $H \in \mathcal{H}$, but this means that the two open half-spaces bounded by H both contain at least one point in $S \subseteq B'$. This is a contradiction to the fact that H is a hyperplane through a facet of B' . Hence the facets of B' indeed all have diameter less than μ . \square

Lemma 3.5. *Let $\delta > 0$ and let $B, B' \in \mathcal{B}_d$ be such that $d_H(B, B') \leq \delta$. Then for every point \mathbf{x} on the boundary of B there exists a point \mathbf{y} on the boundary of B' with $\|\mathbf{x} - \mathbf{y}\|_2 \leq \delta$.*

Proof. We distinguish three cases depending on the position of \mathbf{x} in relation to B' . First, if \mathbf{x} is on the boundary of B' , we can take $\mathbf{y} = \mathbf{x}$ to satisfy the statement in the lemma.

Next, assume that $\mathbf{x} \notin B'$. Then, as $d_H(B, B') \leq \delta$, we have $\min_{\mathbf{y}' \in B'} \|\mathbf{x} - \mathbf{y}'\|_2 = \inf_{\mathbf{y}' \in B'} \|\mathbf{x} - \mathbf{y}'\|_2 \leq \delta$, and so there exists a point $\mathbf{y}' \in B'$ with $\|\mathbf{x} - \mathbf{y}'\|_2 \leq \delta$. Now, consider the straight line segment from $\mathbf{y}' \in B'$ to $\mathbf{x} \notin B'$. This straight line segment must contain some point \mathbf{y} on the boundary of B' , and we have $\|\mathbf{x} - \mathbf{y}\|_2 \leq \|\mathbf{x} - \mathbf{y}'\|_2 \leq \delta$.

Finally, assume that \mathbf{x} is in the interior of B' . Recalling that \mathbf{x} is on the boundary of B , let H be a supporting hyperplane of B through \mathbf{x} (this means that B is contained in one of the closed half-spaces bounded by H

and \mathbf{x} is on H). Now consider the ray orthogonal to H , starting at \mathbf{x} , pointing away from B . This ray needs to contain a point \mathbf{y} on the boundary of B' (since the start of the ray at \mathbf{x} is in the interior of B' and B' is bounded). If $\|\mathbf{x} - \mathbf{y}\|_2 > \delta$, then the closed Euclidean ball of radius δ around \mathbf{y} is disjoint from the half-space bounded by H containing B . Hence $\inf_{\mathbf{b} \in B} \|\mathbf{b} - \mathbf{y}\|_2 = \min_{\mathbf{b} \in B} \|\mathbf{b} - \mathbf{y}\|_2 > \delta$, which contradicts our assumption $d_H(B, B') \leq \delta$. So we must have $\|\mathbf{x} - \mathbf{y}\|_2 \leq \delta$. \square

4 Point sets with many special differences

The following two lemmas encapsulate the first part of our proofs of Theorems 1.1 and 1.3, respectively (see also the proof overview in Section 2). The first of these lemmas states, roughly speaking, that for a list of vectors $\mathbf{u}_1, \dots, \mathbf{u}_k$ in some vector space V over \mathbb{Q} and any subset of V where many differences are in the set $\{\pm \mathbf{u}_1, \dots, \pm \mathbf{u}_k\}$, there must be a lot of linear dependencies among the vectors $\mathbf{u}_1, \dots, \mathbf{u}_k$. More precisely, the span of some small subset of the vectors $\mathbf{u}_1, \dots, \mathbf{u}_k$ must contain many other vectors among $\mathbf{u}_1, \dots, \mathbf{u}_k$.

When we apply Lemma 4.1 in the next section to complete the proof of Theorem 1.1, we will take $\mathbf{u}_1, \dots, \mathbf{u}_k$ to be the unit vectors appearing as unit distances according to some norm in a point set in \mathbb{R}^d . If there are many such unit distances, then the lemma implies that the span of some small subset of the vectors $\mathbf{u}_1, \dots, \mathbf{u}_k$ must contain many other vectors among $\mathbf{u}_1, \dots, \mathbf{u}_k$. In Section 5 we show (as the second part of our proof) that this is a ‘‘special’’ property, meaning that most norms cannot have this property (more precisely, the set of norms with this property is a meagre set).

Lemma 4.1. *Let V be a vector space over \mathbb{Q} , and let $\mathbf{u}_1, \dots, \mathbf{u}_k \in V$ be non-zero vectors in V . Let $\mathbf{p}_1, \dots, \mathbf{p}_n \in V$ be distinct vectors, and let us consider the graph with vertex set $\{1, \dots, n\}$, where for any $x, y \in \{1, \dots, n\}$ we draw an edge between the vertices x and y if and only if $\mathbf{p}_x - \mathbf{p}_y \in \{\pm \mathbf{u}_1, \dots, \pm \mathbf{u}_k\}$. For some positive integer d , suppose that this graph has more than $\frac{d}{2} \cdot n \log n$ edges. Then there exists a subset $I \subseteq \{1, \dots, k\}$, such that we have $\mathbf{u}_\ell \in \text{span}_{\mathbb{Q}}(\mathbf{u}_i : i \in I)$ for at least $d \cdot |I| + 1$ indices $\ell \in \{1, \dots, k\}$.*

We will prove Lemma 4.1 later in this section, and then use it in the next section to prove Theorem 1.1. The following lemma will in a similar fashion be used to prove Theorem 1.3. When applying this lemma, we will take F to be the field extension of \mathbb{Q} generated by all the distinct distances appearing in a given set of n points according to some given norm.

Lemma 4.2. *Let $d \geq 1$ be an integer and let $0 < \mu < 1$. Suppose that n is sufficiently large with respect to d and μ . Let $F \subseteq \mathbb{R}$ be a subfield of \mathbb{R} , and let V be a vector space over \mathbb{R} . Let $\mathbf{u}_1, \dots, \mathbf{u}_k \in V$ be non-zero vectors in V , and let $\mathbf{p}_1, \dots, \mathbf{p}_n \in V$ be distinct vectors such that not all of $\mathbf{p}_1, \dots, \mathbf{p}_n$ are lying on a common affine line in V (as a vector space over \mathbb{R}). Suppose that for all $x, y \in \{1, \dots, n\}$ we have $\mathbf{p}_x - \mathbf{p}_y \in \text{span}_F(\mathbf{u}_i)$ for some $i \in \{1, \dots, k\}$. Then there exists a subset $I \subseteq \{1, \dots, k\}$, such that we have $\mathbf{u}_\ell \in \text{span}_F(\mathbf{u}_i : i \in I)$ for at least $d \cdot |I| + (1 - \mu) \cdot n + 1$ indices $\ell \in \{1, \dots, k\}$.*

The proofs of Lemmas 4.1 and 4.2 rely on the following lemma, which we prove first.

Lemma 4.3. *Let V be a vector space over a field F , and let $\mathbf{u}_1, \dots, \mathbf{u}_k \in V$ be linearly independent vectors in V . For some $n \geq 1$, let $\mathbf{p}_1, \dots, \mathbf{p}_n \in V$ be distinct vectors, and let G be a graph with vertex set $\{1, \dots, n\}$ satisfying the following two conditions*

- *For every edge xy of the graph G , we have $\mathbf{p}_x - \mathbf{p}_y \in \text{span}_F(\mathbf{u}_i)$ for some $i \in \{1, \dots, k\}$.*
- *For each $i \in \{1, \dots, k\}$, the subgraph of G consisting of all edges xy with $\mathbf{p}_x - \mathbf{p}_y \in \text{span}_F(\mathbf{u}_i)$ is a forest.*

Then the graph G has at most $\frac{1}{2} \cdot n \log n$ edges.

In the proof of this lemma, we will use the following simple inequality.

Lemma 4.4. *For any positive integers $n_1 \geq n_2 \geq \dots \geq n_\ell \geq 1$ with $n = n_1 + \dots + n_\ell$ (and $\ell \geq 1$), we have*

$$n_1 - \frac{1}{2} \sum_{i=1}^{\ell} n_i \log n_i \geq n - \frac{1}{2} \cdot n \log n.$$

Proof. Recall that the binary entropy function given by $H(t) = -t \log t - (1-t) \log(1-t)$ for $t \in (0, 1)$ and $H(0) = H(1) = 0$ is a concave function on the interval $[0, 1]$. It intersects the line given by $2 - 2t$ at $t = \frac{1}{2}$ and $t = 1$, and so for $\frac{1}{2} \leq t \leq 1$ we have $H(t) \geq 2 - 2t$ and therefore $t + \frac{1}{2}H(t) \geq 1$.

Let us now prove the desired statement by induction on ℓ . For $\ell = 1$, we have $n_1 = n$ and the statement is trivially true. So let us assume that $\ell \geq 2$ and that we already proved the lemma for $\ell - 1$. Now, setting $n'_1 = n_1 + n_\ell \leq 2n_1$, we have $\frac{1}{2} \leq n_1/n'_1 \leq 1$ and therefore

$$\begin{aligned} n_1 - \frac{1}{2} \cdot n_1 \log n_1 - \frac{1}{2} \cdot n_\ell \log n_\ell &= n_1 - \frac{1}{2} \cdot n'_1 \cdot \log n'_1 - \frac{1}{2} \cdot n_1 \cdot \log(n_1/n'_1) - \frac{1}{2} \cdot n_\ell \cdot \log(n_\ell/n'_1) \\ &= -\frac{1}{2} \cdot n'_1 \cdot \log n'_1 + n'_1 \cdot \left(\frac{n_1}{n'_1} - \frac{1}{2} \cdot \frac{n_1}{n'_1} \cdot \log(n_1/n'_1) - \frac{1}{2} \cdot \frac{n_\ell}{n'_1} \cdot \log(n_\ell/n'_1) \right) \\ &= -\frac{1}{2} \cdot n'_1 \cdot \log n'_1 + n'_1 \cdot \left(\frac{n_1}{n'_1} + \frac{1}{2} \cdot H(n_1/n'_1) \right) \geq -\frac{1}{2} \cdot n'_1 \cdot \log n'_1 + n'_1. \end{aligned}$$

Thus,

$$n_1 - \frac{1}{2} \sum_{i=1}^{\ell} n_i \log n_i \geq n'_1 - \frac{1}{2} \cdot n'_1 \log n'_1 - \sum_{i=2}^{\ell-1} n_i \log n_i \geq n - \frac{1}{2} \cdot n \log n,$$

where the last inequality follows from the induction hypothesis applied to $n'_1 \geq n_2 \geq \dots \geq n_{\ell-1}$ (noting that $n'_1 + n_2 + \dots + n_{\ell-1} = (n_1 + n_\ell) + n_2 + \dots + n_{\ell-1} = n$). \square

Now, we are ready to prove Lemma 4.3.

Proof of Lemma 4.3. We will prove the lemma by induction on n . If $n = 1$, then the graph has $0 = (1/2) \cdot 1 \cdot \log 1$ edges.

So let us now assume that $n \geq 2$, and that we have already proved the lemma for all smaller values of n . Let G be a graph as in the statement of the lemma.

If the graph G is not connected, then we can divide it into two disconnected parts with $n_1 \geq 1$ and $n_2 \geq 1$ vertices (where $n_1 + n_2 = n$). By the induction assumption these parts have at most $(1/2) \cdot n_1 \log n_1$ and $(1/2) \cdot n_2 \log n_2$ edges, respectively. Hence, the total number of edges in G is at most

$$\frac{1}{2} \cdot n_1 \log n_1 + \frac{1}{2} \cdot n_2 \log n_2 \leq \frac{1}{2} \cdot (n_1 + n_2) \log n = \frac{1}{2} \cdot n \log n,$$

as desired. So let us from now on assume that G is connected.

If G has no edges, the desired conclusion trivially holds, so let us assume that G has at least one edge xy . Then there exists some index $i \in \{1, \dots, k\}$ with $\mathbf{p}_x - \mathbf{p}_y \in \text{span}_F(\mathbf{u}_i)$. Upon relabelling, we may assume that $i = k$, i.e. we may assume that there exists at least one edge xy in G with $\mathbf{p}_x - \mathbf{p}_y \in \text{span}_F(\mathbf{u}_k)$.

Let us now colour any edge xy in G red if we have $\mathbf{p}_x - \mathbf{p}_y \in \text{span}_F(\mathbf{u}_k)$. Then there is at least one red edge in G . We remark that by the second condition on the graph G in the statement of Lemma 4.3 the red edges form a forest.

Let us fix an arbitrary vertex $z \in \{1, \dots, n\}$. Since G is connected, every vertex $w \in \{1, \dots, n\}$ can be reached by some path in G starting at z . For every edge xy along this path, we have $\mathbf{p}_x - \mathbf{p}_y \in \text{span}_F(\mathbf{u}_i)$ for some $i \in \{1, \dots, k\}$. Adding up $\mathbf{p}_x - \mathbf{p}_y$ for all edges xy along the path now gives a representation $\mathbf{p}_w = \mathbf{p}_z + a_1\mathbf{u}_1 + \dots + a_k\mathbf{u}_k$ with coefficients $a_1, \dots, a_k \in F$. As the vectors $\mathbf{u}_1, \dots, \mathbf{u}_k \in V$ are linearly independent over F , for every given $w \in \{1, \dots, n\}$ there is a unique such representation. In particular, for every $w \in \{1, \dots, n\}$, the vector \mathbf{p}_w lies in exactly one of the sets $\mathbf{p}_z + a\mathbf{u}_k + F\mathbf{u}_1 + \dots + F\mathbf{u}_{k-1}$ for $a \in F$.

This gives a partition of the vertices $w \in \{1, \dots, n\}$ into subsets W_a for $a \in F$, where for each $a \in F$ the set W_a consists of those vertices $w \in \{1, \dots, n\}$ such that $\mathbf{p}_w \in \mathbf{p}_z + a\mathbf{u}_k + F\mathbf{u}_1 + \dots + F\mathbf{u}_{k-1}$. Note that $\sum_{a \in F} |W_a| = n$.

If xy is an edge in G which is not red, then x and y must belong to the same set W_a . Indeed, we have $\mathbf{p}_x - \mathbf{p}_y \in \text{span}_F(\mathbf{u}_i)$ for some $i \in \{1, \dots, k\}$ and so if $\mathbf{p}_x \in \mathbf{p}_z + a\mathbf{u}_k + F\mathbf{u}_1 + \dots + F\mathbf{u}_{k-1}$, then we have $\mathbf{p}_y \in \mathbf{p}_z + a\mathbf{u}_k + F\mathbf{u}_1 + \dots + F\mathbf{u}_{k-1}$ as well. Hence, every non-red edge of G is inside one of the induced subgraphs $G[W_a]$ for $a \in F$ with $W_a \neq \emptyset$.

Recall that the red edges of G form a forest. We claim that for every $a \in F$ the vertices in the set W_a must all be in distinct connected components of the red forest. Indeed, suppose towards a contradiction that for some $a \in F$ there are two distinct vertices $w, w' \in W_a$ belonging to the same component of the red forest. Then w and w' can be connected by a path of red edges, and for all edges xy on this path we have $\mathbf{p}_x - \mathbf{p}_y \in \text{span}_F(\mathbf{u}_k)$. Adding this up over all edges on the path between w and w' , we conclude that $\mathbf{p}_w - \mathbf{p}_{w'} \in \text{span}_F(\mathbf{u}_k)$. On the other hand, since $w, w' \in W_a$, we have $\mathbf{p}_w, \mathbf{p}_{w'} \in \mathbf{p}_z + a\mathbf{u}_k + F\mathbf{u}_1 + \dots + F\mathbf{u}_{k-1}$ and hence $\mathbf{p}_w - \mathbf{p}_{w'} \in \text{span}_F(\mathbf{u}_1, \dots, \mathbf{u}_{k-1})$. Since the vectors $\mathbf{u}_1, \dots, \mathbf{u}_k \in V$ are linearly independent over F , this implies that $\mathbf{p}_w - \mathbf{p}_{w'} = \mathbf{0}$, which is a contradiction (as the points $\mathbf{p}_1, \dots, \mathbf{p}_n$ are distinct). So indeed, for each $a \in F$, the vertices in W_a belong to distinct connected components of the red forest.

Recalling that there is at least one red edge in G , the two endvertices of this edge must be in different sets W_a and $W_{a'}$. Thus, at least two of the sets W_a for $a \in F$ are non-empty, and since $\sum_{a \in F} |W_a| = n$, we have $|W_a| < n$ for all $a \in F$.

We can now apply the induction hypothesis to each of the graphs $G[W_a]$ for $a \in F$ with $W_a \neq \emptyset$, and obtain that $e(G[W_a]) \leq \frac{1}{2} \cdot |W_a| \cdot \log |W_a|$ whenever $W_a \neq \emptyset$. Hence, the number of non-red edges in G is at most

$$\sum_{\substack{a \in F \\ W_a \neq \emptyset}} \frac{1}{2} \cdot |W_a| \cdot \log |W_a| = \frac{1}{2} \sum_{\substack{a \in F \\ W_a \neq \emptyset}} |W_a| \cdot \log |W_a|.$$

Let $a^* \in F$ be such that $|W_{a^*}|$ is of maximal size among the sets $|W_a|$ for $a \in F$. Then we have $|W_a| \leq |W_{a^*}|$ for all $a \in F$.

We claim that the number of red edges in G is at most $n - |W_{a^*}|$. Indeed, the vertices in W_{a^*} are all in distinct components of the forest formed by the red edges. Hence, this forest has at least $|W_{a^*}|$ components and can therefore have at most $n - |W_{a^*}|$ edges.

Thus, the total number of edges in G (both red edges and non-red edges) is bounded by

$$e(G) \leq n - |W_{a^*}| + \frac{1}{2} \sum_{\substack{a \in F \\ W_a \neq \emptyset}} |W_a| \cdot \log |W_a| \leq n - \left(n - \frac{1}{2} \cdot n \log n \right) = \frac{1}{2} \cdot n \log n,$$

where the second inequality follows from Lemma 4.4. □

Remark. The assertion of Lemma 4.3 is tight for every n which is a power of two, as shown by a generic embedding of the graph of the k -dimensional hypercube. Indeed, for any linearly independent vectors $\mathbf{u}_1, \dots, \mathbf{u}_k$,

we can take $\mathbf{p}_1, \dots, \mathbf{p}_n$ with $n = 2^k$ to be all subset sums of $\{\mathbf{u}_1, \dots, \mathbf{u}_k\}$. Now letting G be the k -dimensional hypercube graph (with $\frac{1}{2} \cdot 2^k k = \frac{1}{2} \cdot n \log n$ edges), all conditions in Lemma 4.3 are satisfied (each of the forests in the second condition is a perfect matching). For the proof of Theorem 1.1 we only need a special case of Lemma 4.3 in which G consists only of edges xy for which $\mathbf{p}_x - \mathbf{p}_y \in \{\pm \mathbf{u}_1, \dots, \pm \mathbf{u}_k\}$ (then each of the forests in the second condition in the lemma is a collection of vertex-disjoint paths). In this case, it is easy to see that the graph G is a subgraph of the infinite graph of all integer sequences in which two sequences are adjacent if and only if they are equal in all coordinates but one, and in this coordinate they differ by one. The precise maximum possible number of edges of a subgraph of n vertices of this graph is known by the isoperimetric inequality of Bollobás and Leader [5, Theorem 15]. It implies the assertion of the lemma for this special case, and supplies also the tight best possible value for all values of n and not only for powers of two.

The second ingredient for the proof of Lemmas 4.1 and 4.2 is the following.

Lemma 4.5. *Let V be a vector space over a field F , and let $\mathbf{u}_1, \dots, \mathbf{u}_k \in V$ be non-zero vectors in V . Let $d \geq 1$ and $m \geq 0$ be integers, and suppose that for every subset $I \subseteq \{1, \dots, k\}$, we have $\mathbf{u}_\ell \in \text{span}_F(\mathbf{u}_i : i \in I)$ for at most $d \cdot |I| + m$ indices $\ell \in \{1, \dots, k\}$. For some real number $M > 0$, let $\lambda_1, \dots, \lambda_k$ be real numbers in the interval $[0, M]$. Then there exists a subset $J \subseteq \{1, \dots, k\}$ such that the vectors \mathbf{u}_j for $j \in J$ are linearly independent (over F) and such that*

$$\sum_{j \in J} \lambda_j \geq \frac{\lambda_1 + \dots + \lambda_k - m \cdot M}{d}.$$

Proof. Upon relabelling, we may assume without loss of generality that $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_k$. Let us now construct a sequence of distinct indices $j_1, \dots, j_r \in \{1, \dots, k\}$ as follows. For any $s \geq 1$, having already chosen the indices j_1, \dots, j_{s-1} , choose j_s to be the minimum index such that $\mathbf{u}_{j_s} \notin \text{span}_F(\mathbf{u}_{j_1}, \dots, \mathbf{u}_{j_{s-1}})$ if such an index exists. If such an index does not exist, i.e. if $\mathbf{u}_1, \dots, \mathbf{u}_k \in \text{span}_F(\mathbf{u}_{j_1}, \dots, \mathbf{u}_{j_{s-1}})$, let us terminate the sequence j_1, \dots, j_{s-1} without choosing any further terms by setting $r = s - 1$. Note that when the sequence terminates we have $\mathbf{u}_1, \dots, \mathbf{u}_k \in \text{span}_F(\mathbf{u}_{j_1}, \dots, \mathbf{u}_{j_r})$. By our assumption applied to $I = \{j_1, \dots, j_r\}$, this implies that $k \leq d \cdot |I| + m = d \cdot r + m$.

Since we have $\mathbf{u}_{j_i} \notin \text{span}_F(\mathbf{u}_{j_1}, \dots, \mathbf{u}_{j_{i-1}})$ for $i = 1, \dots, m$, the vectors $\mathbf{u}_{j_1}, \dots, \mathbf{u}_{j_m}$ are linearly independent.

Now, we claim that for $s = 1, \dots, r$, we have $j_s \leq (s-1)d + m + 1$. Indeed, suppose that for some $1 \leq s \leq m$, we had $j_s > (s-1)d + m + 1$. By the choice of j_s , this would mean that $\mathbf{u}_1, \dots, \mathbf{u}_{(s-1)d+m+1} \in \text{span}_F(\mathbf{u}_{j_1}, \dots, \mathbf{u}_{j_{s-1}})$. But this contradicts our assumption for the set $I = \{j_1, \dots, j_{s-1}\}$. Hence, we must indeed have $j_s \leq (s-1)d + m + 1$ for $s = 1, \dots, r$.

Let us now define $J = \{j_1, \dots, j_r\}$. Then the vectors \mathbf{u}_j for $j \in J$, i.e. the vectors $\mathbf{u}_{j_1}, \dots, \mathbf{u}_{j_r}$, are linearly independent. Furthermore, defining $0 = \lambda_{k+1} = \lambda_{k+2} = \dots$ for notational convenience, we have

$$d \cdot \sum_{j \in J} \lambda_j = d \cdot \sum_{s=1}^r \lambda_{j_s} \geq d \cdot \sum_{s=1}^r \lambda_{(s-1)d+m+1} \geq \sum_{s=1}^r (\lambda_{(s-1)d+1+m} + \dots + \lambda_{sd+m}) = \lambda_{m+1} + \lambda_{m+2} + \dots + \lambda_{rd+m}$$

Here, we used that $\lambda_1 \geq \lambda_2 \geq \dots$ and $j_s \leq (s-1)d + m + 1$ for $s = 1, \dots, r$. Recalling $k \leq rd + m$ and $0 = \lambda_{k+1} = \lambda_{k+2} = \dots$, we can conclude that

$$d \cdot \sum_{j \in J} \lambda_j \geq \lambda_{m+1} + \lambda_{m+2} + \dots + \lambda_{rd+m} = \lambda_1 + \dots + \lambda_k - (\lambda_1 + \dots + \lambda_m) \geq \lambda_1 + \dots + \lambda_k - m \cdot M.$$

Rearranging now gives $\sum_{j \in J} \lambda_j \geq (\lambda_1 + \dots + \lambda_k - mM)/d$, as desired. \square

Remark. The assertion of the Lemma 4.5 can also be established in a shorter way by applying the matroid partition theorem of Edmonds and Fulkerson [15]. To do so, define $d+1$ matroids M_1, \dots, M_{d+1} on the set of

vectors $\{\mathbf{u}_1, \dots, \mathbf{u}_k\}$. Each of the first d matroids M_1, \dots, M_d is the usual linear independence matroid: a set of vectors is independent if and only if it is linearly independent. In the last matroid M_{d+1} , a set is independent if and only if it contains at most m vectors. If every subset I of the vectors spans at most $d \cdot |I| + m$ of the vectors, then for every set of vectors J the sum of ranks of J according to the $d + 1$ matroids M_1, \dots, M_{d+1} is at least $|J|$. By the Edmonds–Fulkerson Theorem this implies that the set $\{\mathbf{u}_1, \dots, \mathbf{u}_k\}$ can be partitioned into $d + 1$ subsets, where one of the subsets is of size at most m and all others are linearly independent sets of vectors. This clearly supplies the conclusion of the lemma.

Now, let us prove Lemma 4.1, using Lemmas 4.3 and 4.5.

Proof of Lemma 4.1. Suppose towards a contradiction that the desired subset $I \subseteq \{1, \dots, k\}$ does not exist. Then for every subset $I \subseteq \{1, \dots, k\}$, we have $\mathbf{u}_\ell \in \text{span}_{\mathbb{Q}}(\mathbf{u}_i : i \in I)$ for at most $d \cdot |I|$ indices $\ell \in \{1, \dots, k\}$. Let G' be the graph with vertex set $\{1, \dots, n\}$ as in the statement of Lemma 4.1 (where for any $x, y \in \{1, \dots, n\}$ we draw an edge between the vertices x and y if and only if $\mathbf{p}_x - \mathbf{p}_y \in \{\pm \mathbf{u}_1, \dots, \pm \mathbf{u}_k\}$). Furthermore, for $i = 1, \dots, k$, let λ_i be the number of edges xy of G' with $\mathbf{p}_x - \mathbf{p}_y \in \{\pm \mathbf{u}_i\}$. Then $\lambda_1 + \dots + \lambda_k \geq e(G') > \frac{d}{2} \cdot n \log n$ by our assumption on the number of edges of G' .

Now, by Lemma 4.5 applied with $F = \mathbb{Q}$ and $m = 0$ there exists a subset $J \subseteq \{1, \dots, k\}$ such that the vectors \mathbf{u}_j for $j \in J$ are linearly independent over \mathbb{Q} and such that $\sum_{j \in J} \lambda_j \geq (\lambda_1 + \dots + \lambda_k)/d$. We may assume without loss of generality that $J = \{1, \dots, k'\}$ for some $k' \in \{1, \dots, k\}$ (otherwise, relabel the indices). Then the vectors $\mathbf{u}_1, \dots, \mathbf{u}_{k'}$ are linearly independent and we have $\lambda_1 + \dots + \lambda_{k'} \geq (\lambda_1 + \dots + \lambda_k)/d > \frac{1}{2} \cdot n \log n$.

Let us now apply Lemma 4.3 to $F = \mathbb{Q}$, the linearly independent vectors $\mathbf{u}_1, \dots, \mathbf{u}_{k'} \in V$, the same $\mathbf{p}_1, \dots, \mathbf{p}_n \in V$ as before, and the graph G obtained from G' by only taking the edges xy with $\mathbf{p}_x - \mathbf{p}_y \in \{\pm \mathbf{u}_1, \dots, \pm \mathbf{u}_{k'}\}$. Note that this graph G satisfies the two conditions in Lemma 4.3. Indeed, for every edge xy of G we have $\mathbf{p}_x - \mathbf{p}_y \in \{\pm \mathbf{u}_i\} \subseteq \text{span}_{\mathbb{Q}}(\mathbf{u}_i)$ for some $i \in \{1, \dots, k'\}$. Furthermore, for every $i \in \{1, \dots, k'\}$, the edges xy with $\mathbf{p}_x - \mathbf{p}_y \in \text{span}_{\mathbb{Q}}(\mathbf{u}_i)$ are precisely the edges with $\mathbf{p}_x - \mathbf{p}_y \in \{\pm \mathbf{u}_i\}$ (as $\mathbf{u}_1, \dots, \mathbf{u}_{k'}$ are linearly independent, we cannot have $\mathbf{p}_x - \mathbf{p}_y \in \text{span}_{\mathbb{Q}}(\mathbf{u}_i)$ if $\mathbf{p}_x - \mathbf{p}_y \in \{\pm \mathbf{u}_j\}$ for $j \neq i$), and these edges form a vertex-disjoint collection of paths and hence a forest. Thus, by Lemma 4.3 the graph G has at most $\frac{1}{2} \cdot n \log n$ edges. On the other hand, the number of edges of G is $\lambda_1 + \dots + \lambda_{k'} > \frac{1}{2} \cdot n \log n$, since for every $i = 1, \dots, k'$ there are λ_i edges xy in G with $\mathbf{p}_x - \mathbf{p}_y \in \{\pm \mathbf{u}_i\}$. This is a contradiction, finishing the proof of Lemma 4.1. \square

For the proof of Lemma 4.2, we also use the following result, proved by Ungar [47] which states that for any n distinct points in \mathbb{R}^t , not all on a common line, the pairs of points define at least $n - 1$ different line directions.

Theorem 4.6 ([47]). *Given n distinct points $\mathbf{p}_1, \dots, \mathbf{p}_n \in \mathbb{R}^t$, not all on one common line, there are at least $n - 1$ different one-dimensional linear subspaces of \mathbb{R}^t that are of the form $\text{span}_{\mathbb{R}}(\mathbf{p}_x - \mathbf{p}_y)$ with $1 \leq x < y \leq n$.*

We remark that Ungar’s result is actually slightly stronger, namely in the case of even n he proved that there are at least n (and not only $n - 1$) different line directions (and in this stronger version, his bounds are tight). Finally, let us prove Lemma 4.2, using Lemmas 4.3 and 4.5 and Theorem 4.6.

Proof of Lemma 4.2. Setting $m = \lceil (1 - \mu) \cdot n \rceil \leq n$, we want to prove that there is a subset $I \subseteq \{1, \dots, k\}$ with $\mathbf{u}_\ell \in \text{span}_F(\mathbf{u}_i : i \in I)$ for at least $d \cdot |I| + m + 1$ indices $\ell \in \{1, \dots, k\}$. Suppose towards a contradiction that the desired subset $I \subseteq \{1, \dots, k\}$ does not exist. Then for every subset $I \subseteq \{1, \dots, k\}$, we have $\mathbf{u}_\ell \in \text{span}_F(\mathbf{u}_i : i \in I)$ for at most $d \cdot |I| + m$ indices $\ell \in \{1, \dots, k\}$.

We may assume that for every $i \in \{1, \dots, k\}$ there exist distinct $x, y \in \{1, \dots, n\}$ with $\mathbf{p}_x - \mathbf{p}_y \in \text{span}_F(\mathbf{u}_i)$ since otherwise, we can just omit all indices i for which this is not the case, and relabel the remaining indices.

By Theorem 4.6, there are at least $n - 1$ different line directions in $\text{span}_{\mathbb{R}}(\mathbf{p}_1, \dots, \mathbf{p}_n) \subseteq V$ appearing among the differences $\mathbf{p}_x - \mathbf{p}_y$ with $1 \leq x < y \leq n$. For each of these differences we have $\mathbf{p}_x - \mathbf{p}_y \in \text{span}_F(\mathbf{u}_i)$ for some $i \in \{1, \dots, k\}$. Hence, there must be at least $n - 1$ different vectors \mathbf{u}_i , so $k \geq n - 1$.

Let us now construct a sequence of distinct indices $j_1, \dots, j_r \in \{1, \dots, k\}$ recursively as follows. Let $s \geq 1$, and assume that we have already chosen the indices j_1, \dots, j_{s-1} . Let us now look at the image of the point set $\{\mathbf{p}_1, \dots, \mathbf{p}_n\} \subseteq V$ under the projection map $V \rightarrow V/\text{span}_F(\mathbf{u}_{j_1}, \dots, \mathbf{u}_{j_{s-1}})$. For every point in this image, let us pick a preimage \mathbf{p}_h and let $H_{s-1} \subseteq \{1, \dots, n\}$ be the resulting set of indices $h \in \{1, \dots, n\}$ for these preimages. Then $|H_{s-1}|$ is precisely the size of the image of $\{\mathbf{p}_1, \dots, \mathbf{p}_n\}$ in $V/\text{span}_F(\mathbf{u}_{j_1}, \dots, \mathbf{u}_{j_{s-1}})$ and for every $x \in \{1, \dots, n\}$ there is exactly one $h \in H_{s-1}$ with $\mathbf{p}_x - \mathbf{p}_h \in \text{span}_F(\mathbf{u}_{j_1}, \dots, \mathbf{u}_{j_{s-1}})$. If there exists an index $i \in \{1, \dots, k\}$ for which there is a forest on the vertex set H_{s-1} with at least $\frac{3d}{\mu}$ edges such that for every edge xy of the forest we have $\mathbf{p}_x - \mathbf{p}_y \in \text{span}_F(\mathbf{u}_i)$, then let us pick such an index i and define $j_s = i$. If there is no such index $i \in \{1, \dots, k\}$, let us terminate the sequence j_1, \dots, j_{s-1} without choosing any further terms by setting $r = s - 1$.

Recall that for $s = 0, \dots, r$, we defined $H_s \subseteq \{1, \dots, n\}$ such that $|H_s|$ is the size of the image of the point set $\{\mathbf{p}_1, \dots, \mathbf{p}_n\} \subseteq V$ under the projection map $V \rightarrow V/\text{span}_F(\mathbf{u}_{j_1}, \dots, \mathbf{u}_{j_s})$, and such that for every $x \in \{1, \dots, n\}$ there is exactly one $h \in H_s$ with $\mathbf{p}_x - \mathbf{p}_h \in \text{span}_F(\mathbf{u}_{j_1}, \dots, \mathbf{u}_{j_s})$. In particular, we have $n = |H_0| \geq |H_1| \geq \dots \geq |H_r| \geq 1$. The following claim gives an upper bound on the length of our sequence j_1, \dots, j_r .

Claim 1: We have $r \leq \frac{\mu}{3d} \cdot n$.

Proof. In order to show the claim, it suffices to prove that $|H_s| \leq |H_{s-1}| - \frac{3d}{\mu}$ for $s = 1, \dots, r$. Indeed, then we have $1 \leq |H_r| \leq |H_0| - r \cdot \frac{3d}{\mu} = n - r \cdot \frac{3d}{\mu}$, which implies $r \leq \frac{\mu}{3d} \cdot n$.

So let $s \in \{1, \dots, r\}$, and consider the set $H_{s-1} \subseteq \{1, \dots, n\}$. The image of the point set $\{\mathbf{p}_1, \dots, \mathbf{p}_n\}$ in the quotient space $V' = V/\text{span}_F(\mathbf{u}_{j_1}, \dots, \mathbf{u}_{j_{s-1}})$ is the same as the image of the point set $\{\mathbf{p}_h : h \in H_{s-1}\}$ in this quotient space V' .

Now, $|H_s|$ is the size of this image under the additional projection given by $V' \rightarrow V'/\text{span}_F(\mathbf{u}_{j_s}) = V/\text{span}_F(\mathbf{u}_{j_1}, \dots, \mathbf{u}_{j_s})$. Recall that by the definition of j_s , there is a forest on the vertex set H_{s-1} with at least $3d/\mu$ edges such that for every edge xy of the forest we have $\mathbf{p}_x - \mathbf{p}_y \in \text{span}_F(\mathbf{u}_{j_s})$. For every such edge xy , the points \mathbf{p}_x and \mathbf{p}_y have the same image in $V'/\text{span}_F(\mathbf{u}_{j_s})$. So for every connected component of this forest, the corresponding points \mathbf{p}_x are mapped to the same point in $V'/\text{span}_F(\mathbf{u}_{j_s})$. Since the forest has at least $\frac{3d}{\mu}$ edges, it has at most $|H_{s-1}| - \frac{3d}{\mu}$ connected components. Hence, the image of the point set $\{\mathbf{p}_h : h \in H_{s-1}\}$ in $V'/\text{span}_F(\mathbf{u}_{j_s})$ has size at most $|H_{s-1}| - \frac{3d}{\mu}$, meaning that indeed $|H_s| \leq |H_{s-1}| - \frac{3d}{\mu}$. \square

Claim 2: We have $|H_r| > \frac{\mu}{3d} \cdot n$.

Proof. Suppose towards a contradiction that $|H_r| \leq \frac{\mu}{3d} \cdot n$. Let us fix some index $h^* \in H_r$. Then for every $h \in H_r$ by the assumption in Lemma 4.2 we can choose an index $i(h) \in \{1, \dots, k\}$ with $\mathbf{p}_h - \mathbf{p}_{h^*} \in \text{span}_F(\mathbf{u}_{i(h)})$. Now, defining $I = \{j_1, \dots, j_r\} \cup \{i(h) : h \in H_r\}$, we have $|I| \leq r + |H_r| \leq 2 \cdot \frac{\mu}{3d} \cdot n$ using Claim 1.

Note that now we have $\mathbf{p}_h - \mathbf{p}_{h^*} \in \text{span}_F(\mathbf{u}_i : i \in I)$ for all $h \in H_r$. Furthermore, for every $x \in \{1, \dots, n\}$ there is some $h \in H_r$ with $\mathbf{p}_x - \mathbf{p}_h \in \text{span}_F(\mathbf{u}_{j_1}, \dots, \mathbf{u}_{j_r}) \subseteq \text{span}_F(\mathbf{u}_i : i \in I)$. This implies that $\mathbf{p}_x - \mathbf{p}_{h^*} = \mathbf{p}_x - \mathbf{p}_h + \mathbf{p}_h - \mathbf{p}_{h^*} \in \text{span}_F(\mathbf{u}_i : i \in I)$ for all $x \in \{1, \dots, n\}$. Thus, we obtain $\mathbf{p}_x - \mathbf{p}_y = \mathbf{p}_x - \mathbf{p}_{h^*} - (\mathbf{p}_y - \mathbf{p}_{h^*}) \in \text{span}_F(\mathbf{u}_i : i \in I)$ for all $x, y \in \{1, \dots, n\}$.

Now, we claim that we have $\mathbf{u}_\ell \in \text{span}_F(\mathbf{u}_i : i \in I)$ for all $\ell \in \{1, \dots, k\}$. Indeed, for each $\ell \in \{1, \dots, k\}$, we assumed that there exist two distinct indices $x, y \in \{1, \dots, k\}$ with $\mathbf{p}_x - \mathbf{p}_y \in \text{span}_F(\mathbf{u}_\ell)$. Since $\mathbf{p}_x \neq \mathbf{p}_y$, this means that $\mathbf{p}_x - \mathbf{p}_y = t\mathbf{u}_\ell$ for some $t \in F \setminus \{0\}$. But now we have $t\mathbf{u}_\ell = \mathbf{p}_x - \mathbf{p}_y \in \text{span}_F(\mathbf{u}_i : i \in I)$, implying that $\mathbf{u}_\ell \in \text{span}_F(\mathbf{u}_i : i \in I)$ as claimed.

Thus, the number of indices $\ell \in \{1, \dots, k\}$ with $\mathbf{u}_\ell \in \text{span}_F(\mathbf{u}_i : i \in I)$ is

$$k \geq n - 1 \geq \frac{2\mu}{3} \cdot n + \lceil (1 - \mu) \cdot n \rceil + 1 = d \cdot \frac{2\mu}{3d} \cdot n + m + 1 \geq d \cdot |I| + m + 1,$$

if n is sufficiently large with respect to μ . This contradicts our assumption that such a set I does not exist. \square

Now, for each $i = 1, \dots, k$, let us choose a forest G_i on the vertex set H_r with the maximum possible number of edges such that for every edge xy of the forest we have $\mathbf{p}_x - \mathbf{p}_y \in \text{span}_F(\mathbf{u}_i)$. Let λ_i be the number of edges of this forest, and note that $\lambda_i \leq \frac{3d}{\mu}$ by our construction of the sequence j_1, \dots, j_r (more precisely, by the fact that the sequence terminates at j_r). In particular, the forest G_i has at most $2\lambda_i$ non-isolated vertices.

Recall that for any pair $(x, y) \in H_r \times H_r$ with $x \neq y$, we have $\mathbf{p}_x - \mathbf{p}_y \in \text{span}_F(\mathbf{u}_i)$ for some $i \in \{1, \dots, k\}$. We claim that for each $i \in \{1, \dots, k\}$ there are at most $(2\lambda_i)^2$ pairs $(x, y) \in H_r \times H_r$ with $x \neq y$ such that $\mathbf{p}_x - \mathbf{p}_y \in \text{span}_F(\mathbf{u}_i)$. Indeed, for any such pair (x, y) the vertices x and y must lie in the same component of the forest G_i (since otherwise we could add the edge xy to the forest G_i , which would be a contradiction to our choice of the forest G_i). Since $x \neq y$, this means that x and y are both non-isolated vertices in G_i , meaning that there are at most $2\lambda_i$ choices for x and at most $2\lambda_i$ choices for y . So for each $i \in \{1, \dots, k\}$, there are indeed at most $(2\lambda_i)^2$ pairs $(x, y) \in H_r \times H_r$ with $x \neq y$ satisfying $\mathbf{p}_x - \mathbf{p}_y \in \text{span}_F(\mathbf{u}_i)$.

Since for each of the $|H_r| \cdot (|H_r| - 1)$ pairs $(x, y) \in H_r \times H_r$ with $x \neq y$ there exists an index $i \in \{1, \dots, k\}$ with $\mathbf{p}_x - \mathbf{p}_y \in \text{span}_F(\mathbf{u}_i)$, we can conclude that

$$\frac{|H_r|^2}{2} \leq |H_r| \cdot (|H_r| - 1) \leq \sum_{i=1}^k (2\lambda_i)^2 = 4 \cdot \sum_{i=1}^k \lambda_i^2 \leq \frac{12d}{\mu} \sum_{i=1}^k \lambda_i.$$

Here we used that $|H_r| > \frac{\mu}{3d} \cdot n \geq 2$, which holds by Claim 2, and that $\lambda_i \leq \frac{3d}{\mu}$ for $i = 1, \dots, k$. Rearranging yields $\lambda_1 + \dots + \lambda_k \geq \frac{\mu}{24d} \cdot |H_r|^2$.

Let us now apply Lemma 4.5 with $M = \frac{3d}{\mu}$, recalling our assumption that for every subset $I \subseteq \{1, \dots, k\}$, we have $\mathbf{u}_\ell \in \text{span}_F(\mathbf{u}_i : i \in I)$ for at most $d \cdot |I| + m$ indices $\ell \in \{1, \dots, k\}$. By Lemma 4.5, there is a subset $J \subseteq \{1, \dots, k\}$ such that the vectors \mathbf{u}_j for $j \in J$ are linearly independent over F and such that

$$\sum_{j \in J} \lambda_j \geq \frac{\lambda_1 + \dots + \lambda_k - m \cdot \frac{3d}{\mu}}{d} \geq \frac{\frac{\mu}{24d} \cdot |H_r|^2 - n \cdot \frac{3d}{\mu}}{d}.$$

By Claim 2 we have $n \cdot \frac{3d}{\mu} \leq \left(\frac{3d}{\mu}\right)^2 \cdot |H_r| \leq \frac{\mu}{48d} \cdot |H_r|^2$ if n (and therefore also $|H_r| \geq \frac{\mu}{3d} \cdot n$) is sufficiently large with respect to d and μ . So we can conclude that

$$\sum_{j \in J} \lambda_j \geq \frac{\frac{\mu}{24d} \cdot |H_r|^2 - n \cdot \frac{3d}{\mu}}{d} \geq \frac{\frac{\mu}{48d} \cdot |H_r|^2}{d} = \frac{\mu}{48d^2} \cdot |H_r|^2.$$

Finally, let G be the graph obtained by taking all edges of the forests G_j for $j \in J$. As the vectors \mathbf{u}_j for $j \in J$ are linearly independent over F , the edge sets of these forests are pairwise disjoint, and hence G has $\sum_{j \in J} \lambda_j \geq \frac{\mu}{48d^2} \cdot |H_r|^2$ edges. On the other hand, G satisfies the assumptions in Lemma 4.3, and so Lemma 4.3 implies that G has at most $\frac{1}{2} \cdot |H_r| \log |H_r|$ edges. Thus,

$$\frac{\mu}{48d^2} \cdot |H_r|^2 \leq e(G) \leq \frac{1}{2} \cdot |H_r| \log |H_r|,$$

which is a contradiction if n (and hence also $|H_r| \geq \frac{\mu}{3d} \cdot n$) is sufficiently large with respect to d and μ . \square

5 Proofs of Theorems 1.1 and 1.3: most norms have few special distances

Fix $d \geq 2$. Recall that \mathcal{B}_d is the collection of all closed, bounded, $\mathbf{0}$ -symmetric convex bodies in \mathbb{R}^d with $\mathbf{0}$ in the interior. The sets $B \in \mathcal{B}_d$ are in direct correspondence with the norms $\|\cdot\|$ on \mathbb{R}^d (where to every norm $\|\cdot\|$ on \mathbb{R}^d we associate the set $B \in \mathcal{B}_d$ arising as the unit ball of $\|\cdot\|$).

Let $\mathcal{A} \subseteq \mathcal{B}_d$ be the set of all $B \in \mathcal{B}_d$ arising as the unit ball of a norm $\|\cdot\|$ on \mathbb{R}^d such that for some $n \geq 1$ there exist n points in \mathbb{R}^d with more than $\frac{d}{2} \cdot n \log n$ unit distances according to the norm $\|\cdot\|$. In order to prove Theorem 1.1 we need to show that $\mathcal{A} \subseteq \mathcal{B}_d$ is a meagre set.

For Theorem 1.3, let $\mu > 0$ and let $n_0(d, \mu) > 1/\mu$ be such that the statement in Lemma 4.2 holds for all $n \geq n_0(d, \mu)$. Now, let $\mathcal{A}_\mu^* \subseteq \mathcal{B}_d$ be the set of all $B \in \mathcal{B}_d$ arising as the unit ball of a norm $\|\cdot\|$ on \mathbb{R}^d such that for some $n \geq n_0(d, \mu)$ there exist n points in \mathbb{R}^d with at most $(1 - \mu) \cdot n$ distinct distances according to the norm $\|\cdot\|$. In order to prove Theorem 1.3 it suffices to show that $\mathcal{A}_\mu^* \subseteq \mathcal{B}_d$ is a meagre set. Indeed, then for most norms on \mathbb{R}^d it is true that for all $n \geq n_0(d, \mu)$ any set of n points in \mathbb{R}^d has at least $(1 - \mu) \cdot n$ distinct distances appearing. This means that the number of distinct distances is of the form $(1 - o(1)) \cdot n$.

Let us define $m = 0$ in the setting of Theorem 1.1 and $m = \lceil (1 - \mu)n \rceil$ in the setting of Theorem 1.3.

To show that \mathcal{A} and \mathcal{A}_μ^* are meagre sets, we need to show that \mathcal{A} and \mathcal{A}_μ^* can be covered by a countable union of nowhere dense subsets of \mathcal{B}_d . We will consider suitably defined subsets $\mathcal{A}_{A,\eta} \subseteq \mathcal{B}_d$, indexed by a $(d\ell + m + 1) \times \ell$ matrix A (for a positive integer ℓ) and some rational number $0 < \eta < \pi/2$. In the setting of Theorem 1.1, i.e. to cover the set \mathcal{A} , we will consider rational matrices $A \in \mathbb{Q}^{(d\ell+m+1) \times \ell}$. In the setting of Theorem 1.3, i.e. to cover the set \mathcal{A}_μ^* , our matrices A will have entries in the function field $\mathbb{Q}(x_1, \dots, x_m)$ (i.e. in the field of rational functions in m variables with rational coefficients). In either case, the entries of A are chosen from a countable field, and so there are only countably many choices for such A and η . Thus, it suffices to show that $\mathcal{A} \subseteq \bigcup_{A,\eta} \mathcal{A}_{A,\eta}$ or $\mathcal{A}_\mu^* \subseteq \bigcup_{A,\eta} \mathcal{A}_{A,\eta}$, respectively, and that each of the sets $\mathcal{A}_{A,\eta} \subseteq \mathcal{B}_d$ is nowhere dense in \mathcal{B}_d .

Let us now define our sets $\mathcal{A}_{A,\eta} \subseteq \mathcal{B}_d$, starting with the setting of Theorem 1.1 (recall that then $m = 0$). Given a rational $(d\ell + 1) \times \ell$ matrix $A = (a_{ji})$ and given a rational $0 < \eta < \pi/2$, let $\mathcal{A}_{A,\eta} \subseteq \mathcal{B}_d$ consist of the unit balls $B \in \mathcal{B}_d$ of all norms $\|\cdot\|$ on \mathbb{R}^d for which there exist unit vectors $\mathbf{u}_1, \dots, \mathbf{u}_{d\ell+1} \in \mathbb{R}^d$ (i.e. vectors with endpoints on the boundary of B) satisfying the following two conditions:

- $\mathbf{u}_j = \sum_{i=1}^{\ell} a_{ji} \mathbf{u}_i$ for $j = 1, \dots, d\ell + 1$.
- For all $1 \leq i < j \leq d\ell + 1$, the angle between the two lines $\text{span}_{\mathbb{R}}(\mathbf{u}_i)$ and $\text{span}_{\mathbb{R}}(\mathbf{u}_j)$ is larger than η .

In the setting of Theorem 1.3, given a $(d\ell + m + 1) \times \ell$ matrix $A = (a_{ji})$ with entries in $\mathbb{Q}(x_1, \dots, x_m)$ and given a rational $0 < \eta < \pi/2$, let $\mathcal{A}_{A,\eta} \subseteq \mathcal{B}_d$ consist of the unit balls $B \in \mathcal{B}_d$ of all norms $\|\cdot\|$ on \mathbb{R}^d for which there exist $z_1, \dots, z_m \in \mathbb{R}$ and unit vectors $\mathbf{u}_1, \dots, \mathbf{u}_{d\ell+m+1} \in \mathbb{R}^d$ (i.e. vectors with endpoints on the boundary of B) satisfying the following three conditions:

- For every entry $a_{ji} \in \mathbb{Q}(x_1, \dots, x_m)$ of the matrix A , the evaluation $a_{ji}(z_1, \dots, z_m) \in \mathbb{R}$ is well-defined (i.e. the polynomial in the denominator of a_{ji} does not evaluate to zero at (z_1, \dots, z_m)).
- $\mathbf{u}_j = \sum_{i=1}^{\ell} a_{ji}(z_1, \dots, z_m) \mathbf{u}_i$ for $j = 1, \dots, d\ell + m + 1$.
- For all $1 \leq i < j \leq d\ell + m + 1$, the angle between the two lines $\text{span}_{\mathbb{R}}(\mathbf{u}_i)$ and $\text{span}_{\mathbb{R}}(\mathbf{u}_j)$ is larger than η .

In order to prove Theorems 1.1 and 1.3, it suffices to show that each of the sets $\mathcal{A}_{A,\eta}$ is nowhere dense and that $\mathcal{A} \subseteq \bigcup_{A,\eta} \mathcal{A}_{A,\eta}$ and $\mathcal{A}_\mu^* \subseteq \bigcup_{A,\eta} \mathcal{A}_{A,\eta}$, respectively. These statements will be the content of the following lemmas.

Lemma 5.1. *In the setting of Theorem 1.1, we have $\mathcal{A} \subseteq \bigcup_{A,\eta} \mathcal{A}_{A,\eta}$, where the union is over all rational $(d\ell + 1) \times \ell$ matrices $A \in \mathbb{Q}^{(d\ell+1) \times \ell}$ for all positive integers ℓ , and all rational numbers $0 < \eta < \pi/2$.*

Proof. Let $B \in \mathcal{A}$. Then B is the unit ball of a norm $\|\cdot\|$ on \mathbb{R}^d with the property that there exist distinct points $\mathbf{p}_1, \dots, \mathbf{p}_n \in \mathbb{R}^d$ such that there are more than $\frac{d}{2} \cdot n \log n$ unit distances according to the norm $\|\cdot\|$ between the points $\mathbf{p}_1, \dots, \mathbf{p}_n$. Now, let $\mathbf{u}_1, \dots, \mathbf{u}_k \in \mathbb{R}^d$ with $\|\mathbf{u}_1\| = \dots = \|\mathbf{u}_k\| = 1$ be the unit vectors (signed arbitrarily) corresponding to the unit distances between the points $\mathbf{p}_1, \dots, \mathbf{p}_n$. Note that $\mathbf{u}_1, \dots, \mathbf{u}_k$ are non-zero vectors in distinct directions (i.e. the lines $\text{span}_{\mathbb{R}}(\mathbf{u}_i)$ for $i = 1, \dots, k$ are all distinct).

Let us now consider \mathbb{R}^d as an (infinite-dimensional) vector space over \mathbb{Q} , and let us apply Lemma 4.1. As in the lemma statement, consider the graph with vertex set $\{1, \dots, n\}$, where for any $x, y \in \{1, \dots, n\}$ we draw an edge between the vertices x and y if and only if $\mathbf{p}_x - \mathbf{p}_y \in \{\pm \mathbf{u}_1, \dots, \pm \mathbf{u}_k\}$. Then the edges correspond precisely to the unit distances according to $\|\cdot\|$ among the points $\mathbf{p}_1, \dots, \mathbf{p}_n$, and so the graph has more than $\frac{d}{2} \cdot n \log n$ edges. Thus, by Lemma 4.1, there exists a subset $I \subseteq \{1, \dots, k\}$, such that we have $\mathbf{u}_j \in \text{span}_{\mathbb{Q}}(\mathbf{u}_i : i \in I)$ for at least $d \cdot |I| + 1$ indices $j \in \{1, \dots, k\}$. Note that we must have $I \neq \emptyset$, since $\mathbf{u}_1, \dots, \mathbf{u}_k$ are non-zero vectors.

Upon relabelling the indices, we may assume that $I = \{1, \dots, \ell\}$ for some integer $\ell \geq 1$, and that $\mathbf{u}_1, \dots, \mathbf{u}_{d\ell+1} \in \text{span}_{\mathbb{Q}}(\mathbf{u}_i : i \in I) = \text{span}_{\mathbb{Q}}(\mathbf{u}_1, \dots, \mathbf{u}_{\ell})$. Then we can find coefficients $a_{ji} \in \mathbb{Q}$ for $j = 1, \dots, d\ell+1$ and $i = 1, \dots, \ell$ such that $\mathbf{u}_j = \sum_{i=1}^{\ell} a_{ji} \mathbf{u}_i$ for all $j = 1, \dots, d\ell+1$.

We can now take A to be the $(d\ell+1) \times \ell$ matrix with entries a_{ji} for $j = 1, \dots, d\ell+1$ and $i = 1, \dots, \ell$. Furthermore, consider the $d\ell+1$ lines $\text{span}_{\mathbb{R}}(\mathbf{u}_i)$ for $i = 1, \dots, d\ell+1$ and choose $0 < \eta < \pi/2$ to be rational and smaller than the angle between any two of these lines. Then $\mathbf{u}_1, \dots, \mathbf{u}_{d\ell+1}$ are unit vectors according to the norm $\|\cdot\|$, such that $\mathbf{u}_j = \sum_{i=1}^{\ell} a_{ji} \mathbf{u}_i$ for all $j = 1, \dots, d\ell+1$ and such that for all $1 \leq i < j \leq d\ell+1$ the angle between the lines $\text{span}_{\mathbb{R}}(\mathbf{u}_i)$ and $\text{span}_{\mathbb{R}}(\mathbf{u}_j)$ is larger than η . This means that the unit ball B of the norm $\|\cdot\|$ belongs to $\mathcal{A}_{A,\eta}$, as desired. \square

Lemma 5.2. *In the setting of Theorem 1.3, we have $\mathcal{A}_{\mu}^* \subseteq \bigcup_{A,\eta} \mathcal{A}_{A,\eta}$, where the union is over all $(d\ell+m+1) \times \ell$ matrices $A \in \mathbb{Q}(x_1, \dots, x_m)^{(d\ell+m+1) \times \ell}$ for all positive integers ℓ , and all rational numbers $0 < \eta < \pi/2$.*

Proof. Let $B \in \mathcal{A}_{\mu}^*$. Then B is the unit ball of a norm $\|\cdot\|$ on \mathbb{R}^d with the property that for some $n \geq n_0(d, \mu)$ there exist distinct points $\mathbf{p}_1, \dots, \mathbf{p}_n \in \mathbb{R}^d$ such that there are at most $(1-\mu) \cdot n \leq m$ distinct distances according to the norm $\|\cdot\|$ between the points $\mathbf{p}_1, \dots, \mathbf{p}_n$. Now, let $\mathbf{u}_1, \dots, \mathbf{u}_k \in \mathbb{R}^d$ with $\|\mathbf{u}_1\| = \dots = \|\mathbf{u}_k\| = 1$ be the unit vectors (signed arbitrarily) in the directions of all the differences between the points $\mathbf{p}_1, \dots, \mathbf{p}_n$, and let $z_1, \dots, z_m > 0$ be positive real numbers such that $\|\mathbf{p}_x - \mathbf{p}_y\| \in \{z_1, \dots, z_m\}$ for all distinct $x, y \in \{1, \dots, n\}$. Note that then for all distinct $x, y \in \{1, \dots, n\}$ we have $\mathbf{p}_x - \mathbf{p}_y = \pm z_j \mathbf{u}_i$ for some $j \in \{1, \dots, m\}$ and $i \in \{1, \dots, k\}$. Defining $F = \mathbb{Q}(z_1, \dots, z_m) \subseteq \mathbb{R}$ to be the field extension of \mathbb{Q} generated by z_1, \dots, z_m , we obtain that for all distinct $x, y \in \{1, \dots, n\}$ we have $\mathbf{p}_x - \mathbf{p}_y \in \text{span}_F(\mathbf{u}_i)$ for some $i \in \{1, \dots, k\}$. Furthermore, note that $\mathbf{u}_1, \dots, \mathbf{u}_k$ are non-zero vectors in distinct directions (i.e. the lines through $\mathbf{u}_1, \dots, \mathbf{u}_k$ are distinct).

If all of the points $\mathbf{p}_1, \dots, \mathbf{p}_n$ are on a common (affine) line in \mathbb{R}^d , then there are at least $n-1$ distinct distances among these points according to the norm $\|\cdot\|$. However we have $n-1 > (1-\mu) \cdot n$ (since $n \geq n_0(d, \mu) > 1/\mu$), so this would be a contradiction. Hence, the points $\mathbf{p}_1, \dots, \mathbf{p}_n$ do not all lie on a common line in \mathbb{R}^d .

By Lemma 4.2 applied to $V = \mathbb{R}^d$ and $F = \mathbb{Q}(z_1, \dots, z_m) \subseteq \mathbb{R}$, there exists a subset $I \subseteq \{1, \dots, k\}$, such that we have $\mathbf{u}_j \in \text{span}_F(\mathbf{u}_i : i \in I)$ for at least $d \cdot |I| + (1-\mu) \cdot n + 1$ indices $j \in \{1, \dots, k\}$. As the number of such indices is an integer, and $m = \lceil (1-\mu)n \rceil$, we have $\mathbf{u}_j \in \text{span}_F(\mathbf{u}_i : i \in I)$ for at least $d \cdot |I| + m + 1$ indices $j \in \{1, \dots, k\}$. Note that we must have $I \neq \emptyset$, since $\mathbf{u}_1, \dots, \mathbf{u}_k$ are non-zero vectors.

Upon relabelling the indices if necessary, we may assume that $I = \{1, \dots, \ell\}$ for some integer $\ell \geq 1$, and that $\mathbf{u}_1, \dots, \mathbf{u}_{d\ell+m+1} \in \text{span}_F(\mathbf{u}_i : i \in I) = \text{span}_F(\mathbf{u}_1, \dots, \mathbf{u}_{\ell})$. Then we can find coefficients $a_{ji}^* \in F = \mathbb{Q}(z_1, \dots, z_m)$ for $j = 1, \dots, d\ell+m+1$ and $i = 1, \dots, \ell$ such that $\mathbf{u}_j = \sum_{i=1}^{\ell} a_{ji}^* \mathbf{u}_i$ for all $j = 1, \dots, d\ell+m+1$. For each of these coefficients $a_{ji}^* \in \mathbb{Q}(z_1, \dots, z_m)$, we can choose a rational function $a_{ji} \in \mathbb{Q}(x_1, \dots, x_m)$ such that $a_{ji}(z_1, \dots, z_m) = a_{ji}^*$ (i.e. a_{ji} evaluates to a_{ji}^* when plugging in z_1, \dots, z_m for the abstract variables

x_1, \dots, x_m in the function field $\mathbb{Q}(x_1, \dots, x_m)$). Note that in particular, all the evaluations $a_{ji}(z_1, \dots, z_m)$ for $j = 1, \dots, dl + m + 1$ and $i = 1, \dots, \ell$ are well-defined.

We can now take A to be the $(dl + m + 1) \times \ell$ matrix with entries a_{ji} for $j = 1, \dots, dl + m + 1$ and $i = 1, \dots, \ell$. Furthermore, consider the $dl + m + 1$ lines $\text{span}_{\mathbb{R}}(\mathbf{u}_i)$ for $i = 1, \dots, dl + m + 1$ and choose $0 < \eta < \pi/2$ to be rational and smaller than the angle between any two of these lines. Then $\mathbf{u}_1, \dots, \mathbf{u}_{dl+m+1}$ are unit vectors according to the norm $\|\cdot\|$, such that $\mathbf{u}_j = \sum_{i=1}^{\ell} a_{ji}^* \mathbf{u}_i = \sum_{i=1}^{\ell} a_{ji}(z_1, \dots, z_m) \mathbf{u}_i$ for all $j = 1, \dots, dl + m + 1$ and such that for all $1 \leq i < j \leq dl + m + 1$ the angle between the lines $\text{span}_{\mathbb{R}}(\mathbf{u}_i)$ and $\text{span}_{\mathbb{R}}(\mathbf{u}_j)$ is larger than η . This means that the unit ball B of the norm $\|\cdot\|$ belongs to $\mathcal{A}_{A,\eta}$, as desired. \square

It remains to show that each of the sets $\mathcal{A}_{A,\eta}$ is nowhere dense in \mathcal{B}_d . This is the content of the following lemma.

Lemma 5.3. *Let $\ell > 0$ be an integer, let $0 < \eta < \pi/2$ be a rational number, and let A be a $(dl + m + 1) \times \ell$ matrix with entries in \mathbb{Q} or in $\mathbb{Q}(x_1, \dots, x_m)$. Then $\mathcal{A}_{A,\eta}$ is nowhere dense in \mathcal{B}_d .*

Proof. As in the lemma statement, let the matrix $A = (a_{ji})$ and $0 < \eta < \pi/2$ be fixed. To show that $\mathcal{A}_{A,\eta}$ is nowhere dense in \mathcal{B}_d , we need to show that for every $B_0 \in \mathcal{B}_d$ and every $\varepsilon > 0$ there exist $B \in \mathcal{B}_d$ and $\delta > 0$ with $d_H(B, B_0) < \varepsilon$ such that no $B' \in \mathcal{B}_d$ with $d_H(B, B') < \delta$ belongs to $\mathcal{A}_{A,\eta}$. So let us fix $B_0 \in \mathcal{B}_d$. It suffices to show the desired statement for small $\varepsilon > 0$ (note that if the statement holds for some $\varepsilon > 0$, then it holds for any larger ε as well).

Recall that $B_0 \subseteq \mathbb{R}^d$ is a closed, bounded, $\mathbf{0}$ -symmetric convex body containing $\mathbf{0}$ in its interior. Due to the latter condition, there exists some $s > 0$ such that $\|\mathbf{b}\|_2 \geq 2s$ for all \mathbf{b} on the boundary of B_0 . By making ε smaller if needed, we may assume that $\varepsilon < s$.

By Lemma 3.4, applied with $\mu = \min\{\varepsilon/2, s \cdot \sin(\eta/3)\}$, we can approximate B_0 with a bounded $\mathbf{0}$ -symmetric polytope B_1 containing $\mathbf{0}$ in its interior such that $d_H(B_0, B_1) < \varepsilon/2$ and all facets of B_1 have diameter at most $s \cdot \sin(\eta/3)$ (with respect to the Euclidean distance). By Lemma 3.5 every point \mathbf{b} on the boundary of B_1 has Euclidean distance at most $\varepsilon/2 < s$ to some point on the boundary of B_0 . In particular, we have $\|\mathbf{b}\|_2 \geq s$ for all \mathbf{b} on the boundary of B_1 (so B_1 contains the Euclidean ball with radius s around the origin).

Let us now choose a system of linear inequalities of the form $|\mathbf{o}_i \cdot \mathbf{x}| \leq s_i$ for $i = 1, \dots, h$ describing the polytope B_1 (with non-zero vectors $\mathbf{o}_1, \dots, \mathbf{o}_h \in \mathbb{R}^d$ and real numbers $s_1, \dots, s_h > 0$). More formally, B_1 is the set of all $\mathbf{x} \in \mathbb{R}^d$ such that $|\mathbf{o}_i \cdot \mathbf{x}| \leq s_i$ for $i = 1, \dots, h$. By rescaling the inequalities, we may choose the vectors $\mathbf{o}_1, \dots, \mathbf{o}_h \in \mathbb{R}^d$ such that $\|\mathbf{o}_1\|_2 = \dots = \|\mathbf{o}_h\|_2 = 1$. As B_1 contains the Euclidean ball of radius s around the origin, we have $s_i \geq s$ for $i = 1, \dots, h$.

Note that geometrically, one can think of $\mathbf{o}_1, \dots, \mathbf{o}_h$ as the (Euclidean) unit vectors orthogonal to the (parallel pairs of) facets of B_1 . Furthermore, note that every point \mathbf{b} on the boundary of B_1 satisfies $|\mathbf{o}_i \cdot \mathbf{b}| = s_i$ for some $i \in \{1, \dots, h\}$.

In the setting of Theorem 1.1 (where $m = 0$ and A has entries in \mathbb{Q}), let us say that an h -tuple $(t_1, \dots, t_h) \in \mathbb{R}^h$ is *achievable* if there exist distinct indices $\varphi(1), \dots, \varphi(dl + 1) \in \{1, \dots, h\}$ and vectors $\mathbf{u}_1, \dots, \mathbf{u}_{dl+1} \in \mathbb{R}^d$ satisfying the following two conditions:

- (a) $\mathbf{u}_j = \sum_{i=1}^{\ell} a_{ji} \mathbf{u}_i$ for $j = 1, \dots, dl + 1$, and
- (b) $|\mathbf{o}_{\varphi(j)} \cdot \mathbf{u}_j| = t_{\varphi(j)}$ for $j = 1, \dots, dl + 1$.

Intuitively, an h -tuple (t_1, \dots, t_h) is achievable if there is a solution $(\mathbf{u}_1, \dots, \mathbf{u}_{dl+1})$ to the linear equations prescribed by the matrix A , such that when plugging in the $dl + 1$ vectors $\mathbf{u}_1, \dots, \mathbf{u}_{dl+1}$ into some choice of $dl + 1$ different expressions $|\mathbf{o}_i \cdot \mathbf{x}|$ for $i \in \{1, \dots, h\}$, one obtains the corresponding entries of (t_1, \dots, t_h) . This

will be relevant in our arguments in the following way: We will consider the polytope $\{\mathbf{x} \in \mathbb{R}^d : |\mathbf{o}_i \cdot \mathbf{x}| \leq t_i \text{ for } i = 1, \dots, h\}$ (note that for (t_1, \dots, t_h) close to (s_1, \dots, s_h) , this can be viewed as a perturbation of the polytope $B_1 = \{\mathbf{x} \in \mathbb{R}^d : |\mathbf{o}_i \cdot \mathbf{x}| \leq s_i \text{ for } i = 1, \dots, h\}$). Whenever there is a solution $(\mathbf{u}_1, \dots, \mathbf{u}_{d\ell+1})$ to the linear equations prescribed by A such that $\mathbf{u}_1, \dots, \mathbf{u}_{d\ell+1}$ lie on distinct facets of this polytope, then the h -tuple (t_1, \dots, t_h) must be achievable (indeed, each \mathbf{u}_j satisfies $|\mathbf{o}_i \cdot \mathbf{u}_j| = t_i$ for the index i corresponding to the facet containing \mathbf{u}_j). Heuristically speaking, we will show that there are only few achievable h -tuples and that therefore this situation can be avoided.

In the setting of Theorem 1.3 (where $m = \lceil (1 - \mu)n \rceil$ and A has entries in $\mathbb{Q}(x_1, \dots, x_m)$), let us say that an h -tuple $(t_1, \dots, t_h) \in \mathbb{R}^h$ is *achievable* if there exist distinct indices $\varphi(1), \dots, \varphi(d\ell + m + 1) \in \{1, \dots, h\}$, real numbers $z_1, \dots, z_m \in \mathbb{R}$ and vectors $\mathbf{u}_1, \dots, \mathbf{u}_{d\ell+m+1} \in \mathbb{R}^d$ satisfying the following three conditions:

- (o) For every entry $a_{ji} \in \mathbb{Q}(x_1, \dots, x_m)$ of the matrix A , the evaluation $a_{ji}(z_1, \dots, z_m) \in \mathbb{R}$ is well-defined.
- (a) $\mathbf{u}_j = \sum_{i=1}^{\ell} a_{ji}(z_1, \dots, z_m) \mathbf{u}_i$ for $j = 1, \dots, d\ell + m + 1$, and
- (b) $|\mathbf{o}_{\varphi(j)} \cdot \mathbf{u}_j| = t_{\varphi(j)}$ for $j = 1, \dots, d\ell + m + 1$.

Again, intuitively, an h -tuple (t_1, \dots, t_h) is achievable if there is a solution $(\mathbf{u}_1, \dots, \mathbf{u}_{d\ell+m+1})$ to the linear equations prescribed by the matrix A with the entries evaluated at (z_1, \dots, z_m) , such that the vectors $\mathbf{u}_1, \dots, \mathbf{u}_{d\ell+m+1}$ satisfy $|\mathbf{o}_i \cdot \mathbf{u}_j| = t_i$ for distinct choices of indices i . This will be relevant in the same way as indicated above in the setting of Theorem 1.1.

For an injective function $\varphi : \{1, \dots, d\ell + m + 1\} \rightarrow \{1, \dots, h\}$, let us say that $(t_1, \dots, t_h) \in \mathbb{R}^h$ is φ -*achievable* if there are vectors $\mathbf{u}_1, \dots, \mathbf{u}_{d\ell+1} \in \mathbb{R}^d$ (and real numbers $z_1, \dots, z_m \in \mathbb{R}$) satisfying the conditions above (more precisely, satisfying the conditions (a) to (b) in the setting of Theorem 1.1 and satisfying the conditions (o) to (b) in the setting of Theorem 1.3). Then (t_1, \dots, t_h) is achievable if and only if it is φ -achievable for some injective function $\varphi : \{1, \dots, d\ell + m + 1\} \rightarrow \{1, \dots, h\}$.

Intuitively speaking, the following two claims show that in both of our two settings only very few h -tuples $(t_1, \dots, t_h) \in \mathbb{R}^h$ are achievable.

Claim 1: In the setting of Theorem 1.1, the set of all achievable h -tuples $(t_1, \dots, t_h) \in \mathbb{R}^h$ can be covered by finitely many $(h - 1)$ -dimensional linear hyperplanes in \mathbb{R}^h .

Proof. Since there are only finitely many possibilities to choose an injective function $\varphi : \{1, \dots, d\ell + 1\} \rightarrow \{1, \dots, h\}$, it suffices to show that for every such function φ the set of all φ -achievable h -tuples $(t_1, \dots, t_h) \in \mathbb{R}^h$ can be covered by finitely many $(h - 1)$ -dimensional linear hyperplanes in \mathbb{R}^h . So let us fix some injective function $\varphi : \{1, \dots, d\ell + 1\} \rightarrow \{1, \dots, h\}$.

If $(t_1, \dots, t_h) \in \mathbb{R}^h$ is φ -achievable, there exist vectors $\mathbf{u}_1, \dots, \mathbf{u}_{d\ell+1} \in \mathbb{R}^d$ satisfying conditions (a) and (b) above. By (b), for every $j = 1, \dots, d\ell + 1$, we have $\mathbf{o}_{\varphi(j)} \cdot \mathbf{u}_j = t_{\varphi(j)}$ or $\mathbf{o}_{\varphi(j)} \cdot \mathbf{u}_j = -t_{\varphi(j)}$. Note that there are $2^{d\ell+1}$ possibilities to make these sign choices. Fixing the sign choices, we can express each $t_{\varphi(j)}$ for $j = 1, \dots, d\ell + 1$ as a linear function of the entries of \mathbf{u}_j (with coefficients given by the entries of the fixed vector $\mathbf{o}_{\varphi(j)}$). However, by (a) the entries of each vector \mathbf{u}_j for $j = 1, \dots, d\ell + 1$ can be expressed as linear functions of the entries of $\mathbf{u}_1, \dots, \mathbf{u}_{\ell}$ (with coefficients given by the entries of the fixed matrix $A = (a_{ji})$). Hence we can express each $t_{\varphi(j)}$ for $j = 1, \dots, d\ell + 1$ as a linear function of the entries of $\mathbf{u}_1, \dots, \mathbf{u}_{\ell} \in \mathbb{R}^d$. These ℓ vectors have $d\ell$ entries in total, and so each $t_{\varphi(j)}$ for $j = 1, \dots, d\ell + 1$ is a linear function of the same $d\ell$ real variables. Since the space of all linear functions in $d\ell$ variables only has dimension $d\ell$, there must be a linear dependence between the $d\ell + 1$ linear functions expressing $t_{\varphi(j)}$ for $j = 1, \dots, d\ell + 1$. In other words, we obtain a linear relationship between the values of $t_{\varphi(j)}$ for $j = 1, \dots, d\ell + 1$ for each of the $2^{d\ell+1}$ sign choices in condition (b). Each such linear relationship gives rise to an $(h - 1)$ -dimensional linear hyperplane in \mathbb{R}^h . Hence the set of φ -achievable h -tuples $(t_1, \dots, t_h) \in \mathbb{R}^h$ can be covered by $2^{d\ell+1}$ linear hyperplanes in \mathbb{R}^h . \square

Claim 2: In the setting of Theorem 1.3, the set of all achievable h -tuples $(t_1, \dots, t_h) \in \mathbb{R}^h$ can be covered by the vanishing sets $\{(y_1, \dots, y_h) \in \mathbb{R}^h : P(y_1, \dots, y_h) = 0\}$ of finitely many non-zero polynomials $P \in \mathbb{R}[y_1, \dots, y_h]$.

Proof. Again, there are only finitely many possibilities to choose an injective function $\varphi : \{1, \dots, dl + m + 1\} \rightarrow \{1, \dots, h\}$, and so it suffices to show that for every such injective function the set of all φ -achievable h -tuples $(t_1, \dots, t_h) \in \mathbb{R}^h$ can be covered by the vanishing sets of finitely many non-zero polynomials. So let us fix some injective function $\varphi : \{1, \dots, dl + m + 1\} \rightarrow \{1, \dots, h\}$.

If $(t_1, \dots, t_h) \in \mathbb{R}^h$ is φ -achievable, there exist vectors $\mathbf{u}_1, \dots, \mathbf{u}_{dl+m+1} \in \mathbb{R}^d$ satisfying conditions (o) to (b) above. By (b), for every $j = 1, \dots, dl + m + 1$, we have $\mathbf{o}_{\varphi(j)} \cdot \mathbf{u}_j = t_{\varphi(j)}$ or $\mathbf{o}_{\varphi(j)} \cdot \mathbf{u}_j = -t_{\varphi(j)}$. Note that there are 2^{dl+m+1} possibilities to make these sign choices. Fixing the sign choices, we can express each $t_{\varphi(j)}$ for $j = 1, \dots, dl + m + 1$ as a linear function (with real coefficients) of the entries of \mathbf{u}_j (with coefficients given by the entries of the fixed vector $\mathbf{o}_{\varphi(j)}$). However, by (a) the entries of each vector \mathbf{u}_j for $j = 1, \dots, dl + m + 1$ can be expressed as rational functions of z_1, \dots, z_m and the entries of $\mathbf{u}_1, \dots, \mathbf{u}_\ell$ (where these rational functions are determined by the entries of the fixed matrix $A = (a_{ji})$). Hence we can express each $t_{\varphi(j)}$ for $j = 1, \dots, dl + m + 1$ as a rational function (with real coefficients) of z_1, \dots, z_m and the entries of $\mathbf{u}_1, \dots, \mathbf{u}_\ell \in \mathbb{R}^d$. The variables z_1, \dots, z_m and the dl entries of the ℓ vectors $\mathbf{u}_1, \dots, \mathbf{u}_\ell \in \mathbb{R}^d$ together are $dl + m$ variables in total, and each $t_{\varphi(j)}$ for $j = 1, \dots, dl + m + 1$ is a rational function (with real coefficients) of these $dl + m$ variables (where the coefficients of the rational function are all determined by the fixed vectors $\mathbf{o}_1, \dots, \mathbf{o}_h$ and the fixed matrix $A \in \mathbb{Q}(x_1, \dots, x_m)^{(dl+m+1) \times \ell}$). By Fact 3.1 there exists a non-zero real polynomial P in $dl + m + 1$ variables, such that when plugging the rational functions describing $t_{\varphi(1)}, \dots, t_{\varphi(dl+m+1)}$ into P the resulting expression is zero. This polynomial P is entirely determined by $\mathbf{o}_1, \dots, \mathbf{o}_h$ and the $A \in \mathbb{Q}(x_1, \dots, x_m)^{(dl+m+1) \times \ell}$ (and the sign choices we made above), and we have $P(t_{\varphi(1)}, \dots, t_{\varphi(dl+m+1)}) = 0$. Thus, for each of the 2^{dl+m+1} sign choices in condition (b), we obtain a non-zero polynomial $P \in \mathbb{R}[y_1, \dots, y_h]$ such that $P(t_1, \dots, t_h) = 0$ for every $(t_1, \dots, t_h) \in \mathbb{R}^h$ which is φ -achievable with these sign choices in condition (b). Hence the set of φ -achievable h -tuples $(t_1, \dots, t_h) \in \mathbb{R}^h$ can be covered by the vanishing sets of 2^{dl+m+1} non-zero polynomials $P \in \mathbb{R}[y_1, \dots, y_h]$. \square

Since B_1 is bounded, there exists some $c > 0$ such that $\|\mathbf{b}\|_2 \leq c$ for all $\mathbf{b} \in B_1$. For any $(t_1, \dots, t_h) \in \mathbb{R}_{>0}^h$, we can consider the polytope $\{\mathbf{x} \in \mathbb{R}^d : |\mathbf{o}_i \cdot \mathbf{x}| \leq t_i \text{ for } i = 1, \dots, h\}$. The maximum (Euclidean) diameter of the facets of this polytope depends continuously on (t_1, \dots, t_h) (indeed, for each $j = 1, \dots, h$, the diameter of the subset of those $\mathbf{x} \in \{\mathbf{x} \in \mathbb{R}^d : |\mathbf{o}_i \cdot \mathbf{x}| \leq t_i \text{ for } i = 1, \dots, h\}$ maximizing $\mathbf{o}_j \cdot \mathbf{x}$ depends continuously on (t_1, \dots, t_h)). Note that for $(t_1, \dots, t_h) = (s_1, \dots, s_h)$ this polytope is precisely B_1 and so all of its facets have diameter at most $s \cdot \sin(\eta/3)$ (with respect to the Euclidean distance). Hence we can choose some $0 < \varepsilon' < \frac{\varepsilon s}{2c}$ such that for every $(t_1, \dots, t_h) \in \mathbb{R}_{>0}^h$ with $s_i \leq t_i \leq s_i + \varepsilon'$ for $i = 1, \dots, h$, all facets of the polytope $\{\mathbf{x} \in \mathbb{R}^d : |\mathbf{o}_i \cdot \mathbf{x}| \leq t_i \text{ for } i = 1, \dots, h\}$ have diameter at most $s \cdot \sin(\eta/2)$ (with respect to the Euclidean distance).

By Claim 1 or Claim 2, respectively, the set of achievable h -tuples $(t_1, \dots, t_h) \in \mathbb{R}^h$ can be covered by a finite collection of $(h - 1)$ -dimensional linear hyperplanes in \mathbb{R}^h or by a finite collection of vanishing sets of non-zero polynomials in \mathbb{R}^h . Let H be the union of these hyperplanes or of the vanishing sets of these polynomials. In either case, $H \subseteq \mathbb{R}^d$ is a closed set and H cannot contain the entire box $[s_1, s_1 + \varepsilon'] \times \dots \times [s_h, s_h + \varepsilon']$. Thus, there exists an h -tuple $(t_1, \dots, t_h) \notin H$ with $s_i \leq t_i \leq s_i + \varepsilon'$ for $i = 1, \dots, h$.

Since $(t_1, \dots, t_h) \notin H$ and H is closed, there exists some $0 < \delta' < \varepsilon'/2$ such that we also have $(t'_1, \dots, t'_h) \notin H$ for all $(t'_1, \dots, t'_h) \in \mathbb{R}^h$ satisfying $|t'_i - t_i| \leq \delta'$ for $i = 1, \dots, h$ (geometrically, δ' is the radius of some ball in the ℓ_∞ -norm around (t_1, \dots, t_h) that is disjoint from H). This means that there is no achievable (t'_1, \dots, t'_h) with $|t'_i - t_i| \leq \delta'$ for $i = 1, \dots, h$.

Let us now define

$$B = \{\mathbf{x} \in \mathbb{R}^d : |\mathbf{o}_i \cdot \mathbf{x}| \leq t_i \text{ for } i = 1, \dots, h\}$$

to be the convex $\mathbf{0}$ -symmetric polytope described by the system of linear inequalities $|\mathbf{o}_i \cdot \mathbf{x}| \leq t_i$ for $i = 1, \dots, h$. As $t_i \geq s_i \geq s > 0$ for $i = 1, \dots, h$, the polytope B contains the Euclidean ball of radius s around the origin, and in particular B contains $\mathbf{0}$ in its interior. Furthermore, B is bounded, since it is contained within $\max\{t_1/s_1, \dots, t_h/s_h\}B_1$ (which is bounded). Thus, we have $B \in \mathcal{B}_d$.

Since $s_i \leq t_i \leq s_i + \varepsilon'$ for $i = 1, \dots, h$, by the choice of ε' we know that all facets of the polytope B have diameter at most $s \cdot \sin(\eta/2)$ (with respect to the Euclidean distance).

We claim that $d_H(B, B_1) \leq \varepsilon/2$. First, note that

$$B_1 = \{\mathbf{x} \in \mathbb{R}^d : |\mathbf{o}_i \cdot \mathbf{x}| \leq s_i \text{ for } i = 1, \dots, h\} \subseteq \{\mathbf{x} \in \mathbb{R}^d : |\mathbf{o}_i \cdot \mathbf{x}| \leq t_i \text{ for } i = 1, \dots, h\} = B,$$

since $s_i \leq t_i$ for $i = 1, \dots, t$. Thus, for every $\mathbf{b}' \in B_1$ we have $\inf_{\mathbf{b} \in B} \|\mathbf{b}' - \mathbf{b}\|_2 = 0$. Furthermore, for every $\mathbf{b} \in B$, we have

$$\left| \mathbf{o}_i \cdot \frac{s}{s + \varepsilon'} \mathbf{b} \right| = \frac{s}{s + \varepsilon'} \cdot |\mathbf{o}_i \cdot \mathbf{b}| \leq \frac{s}{s + \varepsilon'} \cdot t_i \leq \frac{s}{s + \varepsilon'} \cdot (s_i + \varepsilon') \leq s_i$$

for $i = 1, \dots, h$. Hence, letting $\mathbf{b}' = \frac{s}{s + \varepsilon'} \mathbf{b}$, we have $\mathbf{b}' \in B_1$ and $\|\mathbf{b} - \mathbf{b}'\|_2 = \frac{\varepsilon'}{s} \cdot \|\mathbf{b}'\|_2 \leq \frac{\varepsilon' c}{s} < \varepsilon/2$, using that $\|\mathbf{b}'\|_2 \leq c$ (as $\mathbf{b}' \in B_1$) and that $\varepsilon' < \frac{\varepsilon s}{2c}$. This shows that $\inf_{\mathbf{b}' \in B_1} \|\mathbf{b}' - \mathbf{b}\|_2 < \varepsilon/2$ for all $\mathbf{b} \in B$. Thus, we indeed have $d_H(B, B_1) \leq \varepsilon/2$.

Now, since $d_H(B, B_1) \leq \varepsilon/2$ and $d_H(B_0, B_1) < \varepsilon/2$, we have $d_H(B, B_0) < \varepsilon$ by the triangle inequality. In order to finish the proof of Lemma 5.3, let us now show that there is some $\delta > 0$ such that no $B' \in \mathcal{B}_d$ with $d_H(B, B') < \delta$ belongs to $\mathcal{A}_{A, \eta}$.

Let us take δ such that $0 < \delta < \delta'$ and small enough such that $(s \cdot \sin(\eta/2) + 2\delta)/(s - \delta) \leq \sin \eta$. We claim that then for any two points $\mathbf{b}, \mathbf{b}' \in \mathbb{R}^d$ with Euclidean distance at most δ from the same facet of B , the angle between the lines $\text{span}_{\mathbb{R}}(\mathbf{b})$ and $\text{span}_{\mathbb{R}}(\mathbf{b}')$ is at most η . Indeed, recalling that each facet of B has diameter at most $s \cdot \sin(\eta/2)$, by the triangle inequality we have $\|\mathbf{b} - \mathbf{b}'\|_2 \leq 2\delta + s \cdot \sin(\eta/2)$. Hence the sine of the angle between the lines $\text{span}_{\mathbb{R}}(\mathbf{b})$ and $\text{span}_{\mathbb{R}}(\mathbf{b}')$ is at most $\|\mathbf{b} - \mathbf{b}'\|_2 / \|\mathbf{b}\|_2 \leq (s \cdot \sin(\eta/2) + 2\delta)/(s - \delta) \leq \sin \eta$ (noting that $\|\mathbf{b}\|_2 \geq s - \delta$, since B contains the Euclidean ball of radius s around the origin and \mathbf{b} has distance at most δ to the boundary of B). So this indeed shows that the angle between the lines $\text{span}_{\mathbb{R}}(\mathbf{b})$ and $\text{span}_{\mathbb{R}}(\mathbf{b}')$ is at most η .

It remains to show that there is no $B' \in \mathcal{B}_d$ with $d_H(B, B') < \delta$ belonging to $\mathcal{A}_{A, \eta}$. Suppose towards a contradiction that there exists some $B' \in \mathcal{A}_{A, \eta}$ such that $d_H(B, B') < \delta$. By the definition of $\mathcal{A}_{A, \eta}$, this means that there are vectors $\mathbf{u}_1, \dots, \mathbf{u}_{d\ell+m+1} \in \mathbb{R}^d$ with endpoints on the boundary of B' , such that for all $1 \leq i < j \leq d\ell+m+1$, the angle between the two lines $\text{span}_{\mathbb{R}}(\mathbf{u}_i)$ and $\text{span}_{\mathbb{R}}(\mathbf{u}_j)$ is larger than η . Furthermore, in the setting of Theorem 1.1 (where $m = 0$), we have $\mathbf{u}_j = \sum_{i=1}^{\ell} a_{ji} \mathbf{u}_i$ for $j = 1, \dots, d\ell + 1$. In the setting of Theorem 1.3 we also have $z_1, \dots, z_m \in \mathbb{R}$ such that $a_{ji}(z_1, \dots, z_m)$ is well-defined for all entries a_{ji} of A and such that $\mathbf{u}_j = \sum_{i=1}^{\ell} a_{ji}(z_1, \dots, z_m) \mathbf{u}_i$ for $j = 1, \dots, d\ell + m + 1$. Note that this means that condition (a) is satisfied in the setting of Theorem 1.1 and conditions (o) and (a) are satisfied in the setting of Theorem 1.3.

Interpreting $\mathbf{u}_1, \dots, \mathbf{u}_{d\ell+m+1}$ as points in \mathbb{R}^d , on the boundary of B' , by Lemma 3.5 there exist points $\mathbf{b}_1, \dots, \mathbf{b}_{d\ell+m+1} \in \mathbb{R}^d$ on the boundary of B with $\|\mathbf{u}_j - \mathbf{b}_j\|_2 < \delta$ for $j = 1, \dots, d\ell + m + 1$ (recall that $d_H(B', B) < \delta$). We claim that no two of the points $\mathbf{b}_1, \dots, \mathbf{b}_{d\ell+m+1}$ can lie on the same facet or on opposite facets of the polytope B . Indeed, if two of these points \mathbf{b}_j and $\mathbf{b}_{j'}$ with $j \neq j'$ were on the same facet, then by the choice of δ the angle between the lines $\text{span}_{\mathbb{R}}(\mathbf{u}_i)$ and $\text{span}_{\mathbb{R}}(\mathbf{u}_j)$ would be at most η . Similarly, if \mathbf{b}_j and $\mathbf{b}_{j'}$ with $j \neq j'$ were on opposite facets, then \mathbf{b}_j and $-\mathbf{b}_{j'}$ would be on the same facet and so the angle between the lines $\text{span}_{\mathbb{R}}(\mathbf{u}_i)$ and $\text{span}_{\mathbb{R}}(-\mathbf{u}_j) = \text{span}_{\mathbb{R}}(\mathbf{u}_j)$ would also be at most η . In either case, this is a

contradiction to the conditions for the vectors $\mathbf{u}_1, \dots, \mathbf{u}_{dl+m+1}$. Thus, no two of the points $\mathbf{b}_1, \dots, \mathbf{b}_{dl+m+1}$ can lie on the same facet or on opposite facets of the polytope B .

Since $\mathbf{b}_1, \dots, \mathbf{b}_{dl+m+1}$ are on the boundary of $B = \{\mathbf{x} \in \mathbb{R}^d : |\mathbf{o}_i \cdot \mathbf{x}| \leq t_i \text{ for } i = 1, \dots, h\}$, for every $j = 1, \dots, dl+m+1$ we can find some $\varphi(j) \in \{1, \dots, h\}$ such that $|\mathbf{o}_{\varphi(j)} \cdot \mathbf{b}_j| = t_{\varphi(j)}$. Since no two of the points $\mathbf{b}_1, \dots, \mathbf{b}_{dl+m+1}$ are on the same facet or on opposite facets of B , the indices $\varphi(1), \dots, \varphi(dl+m+1) \in \{1, \dots, h\}$ must be distinct, and so they define an injective function $\varphi : \{1, \dots, dl+m+1\} \rightarrow \{1, \dots, h\}$.

For $j = 1, \dots, dl+m+1$, let us define $t'_{\varphi(j)} = |\mathbf{o}_{\varphi(j)} \cdot \mathbf{u}_j|$. Since $\|\mathbf{u}_j - \mathbf{b}_j\|_2 < \delta < \delta'$ and $\|\mathbf{o}_{\varphi(j)}\|_2 = 1$, we have $|t'_{\varphi(j)} - t_{\varphi(j)}| \leq |\mathbf{o}_{\varphi(j)} \cdot \mathbf{u}_j - \mathbf{o}_{\varphi(j)} \cdot \mathbf{b}_j| \leq \|\mathbf{o}_{\varphi(j)}\|_2 \cdot \|\mathbf{u}_j - \mathbf{b}_j\|_2 < \delta'$ for $j = 1, \dots, dl+m+1$. For every $i \in \{1, \dots, h\} \setminus \{\varphi(1), \dots, \varphi(dl+m+1)\}$, let us define $t'_i = t_i$. Then $(t'_1, \dots, t'_h) \in \mathbb{R}^h$ satisfies $|t'_i - t_i| < \delta'$ for $i = 1, \dots, h$, and so by our choice of δ' , the h -tuple (t'_1, \dots, t'_h) is not achievable.

On the other hand, the vectors $\mathbf{u}_1, \dots, \mathbf{u}_{dl+m+1} \in \mathbb{R}^d$ satisfy $|\mathbf{o}_{\varphi(j)} \cdot \mathbf{u}_j| = t'_{\varphi(j)}$ for $j = 1, \dots, dl+m+1$, meaning that condition (b) is satisfied in both the setting of Theorem 1.1 and the setting of Theorem 1.3. Since we already saw that (a) is satisfied (and also (o) in the setting of Theorem 1.3), this means that (t'_1, \dots, t'_h) is φ -achievable and in particular achievable. This contradiction shows that indeed there is no $B' \in \mathcal{B}_d$ with $d_H(B, B') < \delta$ belonging to $\mathcal{A}_{A,\eta}$. This finishes the proof of Lemma 5.3, and hence of Theorems 1.1 and 1.3. \square

6 Proof of Theorem 1.4: upper bounding the chromatic number

In this section, we prove Theorem 1.4. As mentioned in the introduction, we will actually prove a more general result giving an upper bound for the chromatic number of the odd distance graph of most d -norms, in any dimension d . The odd distance graph of a d -norm $\|\cdot\|$ is the graph whose vertices are the points in \mathbb{R}^d , where two points are adjacent if and only if the distance between them according to the norm $\|\cdot\|$ is an odd integer.

Theorem 6.1. *For any $d \geq 1$, for all d -norms $\|\cdot\|$ besides a meagre set, the odd distance graph of $\|\cdot\|$ has chromatic number at most 2^d .*

Since the unit distance graph is a subgraph of the odd distance graph, Theorem 6.1 implies that for all 2-norms besides a meagre set, the unit distance graph has chromatic number at most 4. Combined with the fact that this chromatic number is at least 4 for any 2-norm [9, Theorem 3], this shows Theorem 1.4.

Proof of Theorem 6.1. Our goal is to show that every d -norm $\|\cdot\|$ for which the odd distance graph has chromatic number larger than 2^d has the property that there exist vectors $\mathbf{u}_1, \dots, \mathbf{u}_k \in \mathbb{R}^d$ in distinct directions with $\|\mathbf{u}_1\| = \dots = \|\mathbf{u}_k\| = 1$ and a subset $I \subseteq \{1, \dots, k\}$, such that we have $\mathbf{u}_j \in \text{span}_{\mathbb{Q}}(\mathbf{u}_i : i \in I)$ for at least $d \cdot |I| + 1$ indices $j \in \{1, \dots, k\}$. This would be sufficient, as the arguments in the previous section show that the set of d -norms with this property is a meagre set (indeed, by the second half of the proof of Lemma 5.1 the unit ball of the norm $\|\cdot\|$ is contained in the union $\bigcup_{A,\eta} \mathcal{A}_{A,\eta}$ considered in the statement of Lemma 5.1, and this union is a meagre set, as each of the sets $\mathcal{A}_{A,\eta}$ is nowhere dense by Lemma 5.3).

So let $\|\cdot\|$ be a d -norm whose odd distance graph has chromatic number larger than 2^d , and suppose towards a contradiction that it does not satisfy the property above. By the De Bruijn–Erdős Theorem, the odd distance graph of $\|\cdot\|$ must contain a finite subgraph G with chromatic number $\chi(G) > 2^d$. Let $\mathbf{p}_1, \dots, \mathbf{p}_n \in \mathbb{R}^d$ correspond to the vertices of the subgraph G . Furthermore, let $\mathbf{u}_1, \dots, \mathbf{u}_k \in \mathbb{R}^d$ be the unit vectors (signed arbitrarily) whose odd multiples give rise to edges of G . More precisely, $\mathbf{u}_1, \dots, \mathbf{u}_k \in \mathbb{R}^d$ is a list of non-zero vectors in distinct directions (i.e. the lines $\text{span}_{\mathbb{R}}(\mathbf{u}_i)$ for $i = 1, \dots, k$ are all distinct) with $\|\mathbf{u}_1\| = \dots = \|\mathbf{u}_k\| =$

1 such that for every edge $\mathbf{p}_x\mathbf{p}_y$ of G (with $x, y \in \{1, \dots, n\}$), there exists an index $i \in \{1, \dots, k\}$ and an odd integer m with $\mathbf{p}_x - \mathbf{p}_y = m\mathbf{u}_i$.

Recall that by our assumption about the norm $\|\cdot\|$, for any subset $I \subseteq \{1, \dots, k\}$ we have $\mathbf{u}_j \in \text{span}_{\mathbb{Q}}(\mathbf{u}_i : i \in I)$ for at most $d \cdot |I|$ indices $j \in \{1, \dots, k\}$. By the Edmonds Matroid Decomposition Theorem [14], this implies that there is a partition $\{1, \dots, k\} = J_1 \cup \dots \cup J_d$ such that for each $\ell = 1, \dots, d$ the set vectors $\{\mathbf{u}_i : i \in J_\ell\}$ is linearly independent over \mathbb{Q} . This gives rise to a partition of the edges of the graph G into subgraphs G_1, \dots, G_d , where for any $\ell = 1, \dots, d$, an edge $\mathbf{p}_x\mathbf{p}_y$ of G belongs to the subgraph G_ℓ if we have $\mathbf{p}_x - \mathbf{p}_y = m\mathbf{u}_i$ for some index $i \in J_\ell$ and an odd integer m .

We claim that each of the subgraphs G_ℓ for $\ell = 1, \dots, d$ is bipartite. Indeed, suppose that for some $\ell \in \{1, \dots, d\}$ the subgraph G_ℓ contains an odd cycle. Walking along the cycle and summing up the corresponding terms $\mathbf{p}_x - \mathbf{p}_y = m\mathbf{u}_i$ for the edges along the cycle gives an integer linear relation between the vectors \mathbf{u}_i for $i \in J_\ell$. As the cycle is odd, there must be an index $i \in J_\ell$ such that an odd number of edges on the cycle contribute a term of the form $\mathbf{p}_x - \mathbf{p}_y = m\mathbf{u}_i$ (with an odd integer m). For this index i , the resulting coefficient of \mathbf{u}_i in this integer linear relation is odd, and is therefore in particular non-zero. Hence we obtained a non-trivial integer linear relation between the vectors \mathbf{u}_i for $i \in J_\ell$, contradicting the fact that these vectors are linearly independent over \mathbb{Q} .

Thus, for $\ell = 1, \dots, d$, the graph G_ℓ must indeed be bipartite, meaning that $\chi(G_\ell) \leq 2$. This implies $\chi(G) = \chi(G_1 \cup \dots \cup G_d) \leq \chi(G_1) \cdots \chi(G_d) \leq 2^d$, giving the desired final contradiction. \square

7 Proof of Theorem 1.2: point sets with many unit distances

This section is concerned with lower bounds for $U_{\|\cdot\|}(n)$. Before turning to the proofs, let us introduce a standard piece of notation that we will use throughout the section. Given a collection of sets $S_1, \dots, S_k \in \mathbb{R}^d$ their *Minkowski sum* $S_1 + \dots + S_k$ is defined as the set of all points of the form $\mathbf{x}_1 + \dots + \mathbf{x}_k$ with $\mathbf{x}_i \in S_i$ for $i = 1, \dots, k$.

7.1 Planar norms

Before proving our lower bound for $U_{\|\cdot\|}(n)$ in Theorem 1.2 in arbitrary dimension, we begin by showing a slightly stronger lower bound in dimension 2. We note that, as we will explain in Section 8, it is possible to obtain a further slight improvement.

Proposition 7.1. *For every 2-norm $\|\cdot\|$, we have*

$$U_{\|\cdot\|}(n) \geq \left(\frac{1}{\log 3} - o(1) \right) \cdot n \log n$$

for all n . In other words, for every 2-norm $\|\cdot\|$ and every n , there exists a set of n points in \mathbb{R}^2 such that the number of unit distances according to $\|\cdot\|$ among the n points is at least $\left(\frac{1}{\log 3} - o(1) \right) n \log n$, where the $o(1)$ -term tends to zero as $n \rightarrow \infty$.

If the norm $\|\cdot\|$ is not strictly convex (i.e. if the boundary of the unit ball of the norm contains a line segment), then we can find n points with at least $\lfloor n^2/4 \rfloor$ unit distances. Indeed, in this case, the corresponding unit distance graph contains the complete bipartite graph $K_{m,m}$ for every m . So to prove Proposition 7.1, we may assume that the norm is strictly convex.

The main ingredient in the proof is to show that any power of a triangle is a subgraph of the unit distance graph of any strictly convex norm on \mathbb{R}^2 . More concretely, for any k , we find a collection of k triangles with all sides having unit length according to the given norm, and take the Minkowski sum of the vertex sets of these triangles. This yields 3^k points and the corresponding unit distance graph is a k -fold product of the triangle K_3 (in the usual graph theoretic sense of graph products). To find these triangles, we prove the following simple lemma.

Lemma 7.2. *Let $\|\cdot\|$ be a strictly convex 2-norm. Then for any integer k there exist subsets $S_1, S_2, \dots, S_k \subseteq \mathbb{R}^2$ of size $|S_1| = \dots = |S_k| = 3$ with $|S_1 + \dots + S_k| = 3^k$ and such that for each $i = 1, \dots, k$ any two distinct points in S_i have unit distance according to $\|\cdot\|$.*

Note that the condition $|S_1 + \dots + S_k| = 3^k$ means precisely that all the 3^k points of the form $\mathbf{x}_1 + \mathbf{x}_2 + \dots + \mathbf{x}_k$ with $\mathbf{x}_i \in S_i$ for $i = 1, \dots, k$ are distinct. Furthermore, note that the second condition means that for each $i = 1, \dots, k$, the three points in S_i determine a triangle with all three side lengths equal to one according to the norm $\|\cdot\|$.

Proof. We construct the sets S_1, S_2, \dots one by one, maintaining the property that after S_1, S_2, \dots, S_i have been determined, we have $|S_1 + \dots + S_i| = 3^i$. We need the simple fact (see, for example, [6, Lemma 2]) that for any two distinct points \mathbf{a}_1 and \mathbf{a}_2 in a strictly convex 2-dimensional real normed space $(\mathbb{R}^2, \|\cdot\|)$ there are at most two points $\mathbf{b} \in \mathbb{R}^2$ satisfying both $\|\mathbf{b} - \mathbf{a}_1\| = 1$ and $\|\mathbf{b} - \mathbf{a}_2\| = 1$ (i.e. having unit distance from both \mathbf{a}_1 and \mathbf{a}_2).

Assume S_1, S_2, \dots, S_{i-1} have already been constructed, and let us now choose S_i . For any point $\mathbf{x} \in \mathbb{R}^2$ with $\|\mathbf{x}\| = 1$ (i.e. for any point \mathbf{x} on the boundary of the unit ball of the norm $\|\cdot\|$), let us imagine that we move a point \mathbf{y} along the boundary of the unit ball of $\|\cdot\|$ from \mathbf{x} to the antipodal point $-\mathbf{x}$. By continuity of the distance $\|\mathbf{y} - \mathbf{x}\|$ as \mathbf{y} moves, at some point in this process we must have $\|\mathbf{y} - \mathbf{x}\| = 1$. In other words, for any point $\mathbf{x} \in \mathbb{R}^2$ with $\|\mathbf{x}\| = 1$, there must be a point $\mathbf{y}(\mathbf{x}) \in \mathbb{R}^2$ with $\|\mathbf{y}(\mathbf{x})\| = 1$ and $\|\mathbf{y}(\mathbf{x}) - \mathbf{x}\| = 1$. We will choose the set S_i to consist of $\mathbf{0}$, \mathbf{x} and $\mathbf{y} = \mathbf{y}(\mathbf{x})$, for an appropriate choice of \mathbf{x} with $\|\mathbf{x}\| = 1$. Note that then any two distinct points in S_i have unit distance according to $\|\cdot\|$.

In order to ensure that $|S_1 + \dots + S_i| = 3^i$, we need to ensure that the sets $\{\mathbf{u} + \mathbf{0}, \mathbf{u} + \mathbf{x}, \mathbf{u} + \mathbf{y}\}$ are pairwise disjoint for all $\mathbf{u} \in S_1 + S_2 + \dots + S_{i-1}$. If the sets $\{\mathbf{u} + \mathbf{0}, \mathbf{u} + \mathbf{x}, \mathbf{u} + \mathbf{y}\}$ and $\{\mathbf{v} + \mathbf{0}, \mathbf{v} + \mathbf{x}, \mathbf{v} + \mathbf{y}\}$ intersect for two distinct elements $\mathbf{u}, \mathbf{v} \in S_1 + S_2 + \dots + S_{i-1}$, then one of the differences $\mathbf{x} - \mathbf{0}, \mathbf{y} - \mathbf{0}, \mathbf{x} - \mathbf{y}$ must be of the form $\pm(\mathbf{u} - \mathbf{v})$. We claim that for any two distinct $\mathbf{u}, \mathbf{v} \in S_1 + S_2 + \dots + S_{i-1}$, there are at most six choices for \mathbf{x} (with $\|\mathbf{x}\| = 1$) such that this happens. Indeed, there are clearly at most two choices for \mathbf{x} with $\mathbf{x} = \mathbf{x} - \mathbf{0} \in \{\pm(\mathbf{u} - \mathbf{v})\}$. Note that \mathbf{x} has unit distance from both \mathbf{y} and $\mathbf{x} - \mathbf{y}$ (since $\|\mathbf{y} - \mathbf{x}\| = 1$ and $\|\mathbf{y}\| = 1$). So if $\mathbf{y} = \mathbf{y} - \mathbf{0} \in \{\pm(\mathbf{u} - \mathbf{v})\}$ or $\mathbf{x} - \mathbf{y} \in \{\pm(\mathbf{u} - \mathbf{v})\}$, then \mathbf{x} must have unit distance from $\mathbf{u} - \mathbf{v}$ or from $-(\mathbf{u} - \mathbf{v})$. But by the above-mentioned fact, there can be at most two choices for \mathbf{x} with $\|\mathbf{x}\| = 1$ and $\|\mathbf{x} - (\mathbf{u} - \mathbf{v})\| = 1$ and at most two choices for \mathbf{x} with $\|\mathbf{x}\| = 1$ and $\|\mathbf{x} + (\mathbf{u} - \mathbf{v})\| = 1$. Thus, there are indeed at most six choices for \mathbf{x} (with $\|\mathbf{x}\| = 1$) such that the sets $\{\mathbf{u} + \mathbf{0}, \mathbf{u} + \mathbf{x}, \mathbf{u} + \mathbf{y}\}$ and $\{\mathbf{v} + \mathbf{0}, \mathbf{v} + \mathbf{x}, \mathbf{v} + \mathbf{y}\}$ intersect.

Since there are only finitely many choices for distinct $\mathbf{u}, \mathbf{v} \in S_1 + S_2 + \dots + S_{i-1}$, this means that there are only finitely many choices for \mathbf{x} for which the desired condition $|S_1 + \dots + S_i| = 3^i$ fails. Since there are infinitely many points $\mathbf{x} \in \mathbb{R}^2$ with $\|\mathbf{x}\| = 1$, we can indeed take $S_i = \{\mathbf{0}, \mathbf{x}, \mathbf{y}(\mathbf{x})\}$ for some suitably chosen $\mathbf{x} \in \mathbb{R}^2$ (with $\|\mathbf{x}\| = 1$) such that $|S_1 + \dots + S_i| = 3^i$. \square

Proposition 7.1 is an easy consequence of the last lemma.

Proof of Proposition 7.1. As discussed previously we may assume that the norm $\|\cdot\|$ is strictly convex. Now for any k , we can take sets S_1, \dots, S_k as in Lemma 7.2 and consider the Minkowski sum $S = S_1 + S_2 + \dots + S_k$. This gives a set $S \subseteq \mathbb{R}^2$ of size $|S| = 3^k$, such that every point in S has unit distance according to $\|\cdot\|$ from at least $2k$ other points in S (namely, all the points in S obtained by changing exactly one summand in the representation $\mathbf{x}_1 + \dots + \mathbf{x}_k$ with $\mathbf{x}_i \in S_i$ for $i = 1, \dots, k$). This gives us a set of 3^k points in \mathbb{R}^2 with $k \cdot 3^k$ unit distances according to $\|\cdot\|$, which in particular shows that the claimed bound holds without the $o(1)$ -term when n is a power of three.

To obtain the asymptotic bound for every n , let us write n in base three, so $n = a_k \cdot 3^k + a_{k-1} \cdot 3^{k-1} + \dots + a_0$, where $k = \lfloor \log_3 n \rfloor$ and $a_k, \dots, a_0 \in \{0, 1, 2\}$. For each $i = 0, \dots, k$, we now use the construction above to find a_i sets of size 3^i , translated if necessary to ensure that all the points in all of these sets are distinct. In total we obtain $a_k \cdot 3^k + a_{k-1} \cdot 3^{k-1} + \dots + a_0 = n$ points, and among these n points there are at least

$$a_k \cdot k \cdot 3^k + a_{k-1} \cdot (k-1) \cdot 3^{k-1} + \dots + a_1 \cdot 1 \cdot 3 = kn - 3^k \cdot \sum_{i=1}^k a_{k-i} \cdot i \cdot 3^{-i} \geq kn - 2 \cdot 3^k \cdot \sum_{i=1}^k i \cdot 3^{-i} \geq kn - \frac{3}{2}n$$

unit distances. Since $k - \frac{3}{2} = \left(\frac{1}{\log 3} - o(1)\right) \log n$, this completes the proof. \square

7.2 Higher dimension

In this section, we prove Theorem 1.2. We start by showing that the unit distance graph of any d -norm contains a complete bipartite graph with one vertex class of size $d-1$ and the other vertex class of infinite size. To do so, we need the following result, known as the Hurewicz Dimension Lowering Theorem (see [17, Theorem 3.3.10, p. 200]). The notion of dimension in this theorem is the usual Lebesgue covering dimension, as described in [17].

Theorem 7.3 (Hurewicz Dimension Lowering Theorem). *Let X and Y be metric spaces, and assume that X is compact. Let $f : X \rightarrow Y$ be a continuous map, such that all fibers $f^{-1}(y)$ for $y \in Y$ have dimension at most k . Then we have $\dim(X) \leq k + \dim(Y)$.*

The following lemma is the key ingredient for our lower bound result in Theorem 1.2.

Lemma 7.4. *Let $\|\cdot\|$ be a d -norm, and let $\partial B = \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x}\| = 1\}$ be the boundary of the corresponding unit ball. Furthermore, consider distinct points $\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_{d-1} \in \partial B$, and let $\varepsilon > 0$ be a real number. Then there exist points $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{d-1} \in \mathbb{R}^d$ satisfying $\|\mathbf{x}_i - \mathbf{y}_i\|_2 \leq \varepsilon$ for $i = 1, \dots, d-1$, such that the intersection $\bigcap_{i=1}^{d-1} (\mathbf{x}_i + \partial B)$ is infinite.*

Proof. Define a subset U of $\mathbb{R}^d \times (\mathbb{R}^d)^{d-1}$ by

$$U := \{(\mathbf{x}, \mathbf{x}_1, \dots, \mathbf{x}_{d-1}) : \mathbf{x}_i - \mathbf{x} \in \partial B \text{ for } i = 1, \dots, d-1\}.$$

Note that U is homeomorphic to $\mathbb{R}^d \times (\partial B)^{d-1}$, with a homeomorphism $u : U \rightarrow \mathbb{R}^d \times (\partial B)^{d-1}$ given by $u(\mathbf{x}, \mathbf{x}_1, \dots, \mathbf{x}_{d-1}) = (\mathbf{x}, \mathbf{x}_1 - \mathbf{x}, \dots, \mathbf{x}_{d-1} - \mathbf{x})$. Now, let us define

$$U' := \{(\mathbf{x}, \mathbf{x}_1, \dots, \mathbf{x}_{d-1}) \in U : \|\mathbf{x}_i - \mathbf{y}_i\|_2 \leq \varepsilon \text{ for } i = 1, \dots, d-1\}.$$

Observe that the image of U' under the homeomorphism $u : U \rightarrow \mathbb{R}^d \times (\partial B)^{d-1}$ defined above contains the closed set

$$\{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x}\|_2 \leq \varepsilon/2\} \times \{\mathbf{z}_1 \in \partial B : \|\mathbf{z}_1 - \mathbf{y}_1\|_2 \leq \varepsilon/2\} \times \dots \times \{\mathbf{z}_{d-1} \in \partial B : \|\mathbf{z}_{d-1} - \mathbf{y}_{d-1}\|_2 \leq \varepsilon/2\}, \quad (1)$$

Indeed, given $(\mathbf{x}, \mathbf{z}_1, \dots, \mathbf{z}_{d-1})$ in this set, we have $u(\mathbf{x}, \mathbf{x}_1, \dots, \mathbf{x}_{d-1}) = (\mathbf{x}, \mathbf{z}_1, \dots, \mathbf{z}_{d-1})$ for $(\mathbf{x}, \mathbf{x}_1, \dots, \mathbf{x}_{d-1}) \in U'$ given by $\mathbf{x}_i = \mathbf{z}_i + \mathbf{x}$ for $i = 1, \dots, d-1$ (then $\mathbf{x}_i - \mathbf{x} \in \partial B$ and $\|\mathbf{x}_i - \mathbf{y}_i\|_2 = \|\mathbf{x} + \mathbf{z}_i - \mathbf{y}_i\|_2 \leq \|\mathbf{x}\|_2 + \|\mathbf{z}_i - \mathbf{y}_i\|_2 \leq \varepsilon$).

For $i = 1, \dots, d-1$, we have $\mathbf{y}_i \in \partial B$ and hence the set $\{\mathbf{z}_i \in \partial B : \|\mathbf{z}_i - \mathbf{y}_i\|_2 \leq \varepsilon/2\}$ contains a closed subset that is homeomorphic to $[0, 1]^{d-1}$ (indeed, the map $\mathbf{z} \mapsto \mathbf{z}/\|\mathbf{z}\|_2$ defines a homeomorphism from this set to a subset of the $(d-1)$ -dimensional unit sphere in \mathbb{R}^d containing some spherical cap of positive radius). Thus, the set in (1) has a closed subset that is homeomorphic to $[0, 1]^{d+(d-1)^2}$. Therefore the dimension of the closed set in (1) is at least $d + (d-1)^2$, and consequently $\dim(U') \geq d + (d-1)^2$.

Furthermore, the set U' is closed and bounded (for boundedness, note that $\|\mathbf{x}\| \leq \|\mathbf{x}_1\| + 1$ and $\|\mathbf{x}_i - \mathbf{y}_i\|_2 \leq \varepsilon$ for $i = 1, \dots, d-1$ for all $(\mathbf{x}, \mathbf{x}_1, \dots, \mathbf{x}_{d-1}) \in U'$). Thus, U' is compact. Now, let us consider the continuous projection map $f : U' \rightarrow (\mathbb{R}^d)^{d-1}$ given by $f(\mathbf{x}, \mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{d-1}) = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{d-1})$.

Since $\dim((\mathbb{R}^d)^{d-1}) = d(d-1) < \dim(U')$, by the Hurewicz Dimension Lowering Theorem applied to the projection map $f : U' \rightarrow (\mathbb{R}^d)^{d-1}$, there exist points $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{d-1} \in \mathbb{R}^d$ such that the fiber $f^{-1}(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{d-1})$ has dimension at least one. Then this fiber must be infinite, so there are infinitely many points $\mathbf{x} \in \mathbb{R}^d$ with $(\mathbf{x}, \mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{d-1}) \in U'$. By the definition of U' , this implies $\|\mathbf{x}_i - \mathbf{y}_i\|_2 \leq \varepsilon$ for $i = 1, \dots, d-1$. Furthermore, each of the infinitely many points \mathbf{x} with $(\mathbf{x}, \mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{d-1}) \in U' \subseteq U$ satisfies $\mathbf{x}_i - \mathbf{x} \in \partial B$ for $i = 1, \dots, d-1$, meaning that $\mathbf{x} - \mathbf{x}_i \in \partial B$ and $\mathbf{x} \in \mathbf{x}_i + \partial B$. Thus, each of these points \mathbf{x} lies in the intersection $\bigcap_{i=1}^{d-1} (\mathbf{x}_i + \partial B)$, so this intersection is infinite. \square

An easy consequence of the last lemma is the following.

Lemma 7.5. *Let k be a positive integer and let $\|\cdot\|$ be a d -norm. Then there exist points $\mathbf{x}_1, \dots, \mathbf{x}_{d-1} \in \mathbb{R}^d$ with $|\{\mathbf{x}_1, 2\mathbf{x}_1, \dots, k\mathbf{x}_1\} + \dots + \{\mathbf{x}_{d-1}, 2\mathbf{x}_{d-1}, \dots, k\mathbf{x}_{d-1}\}| = k^{d-1}$ and an infinite set $S \subseteq \mathbb{R}^d$ such that every point $\mathbf{z} \in S$ satisfies $\|\mathbf{z} - \mathbf{x}_i\| = 1$ for $i = 1, \dots, d-1$.*

Proof. Let ∂B be the boundary of the unit ball of the norm $\|\cdot\|$. Let us choose points $\mathbf{y}_1, \dots, \mathbf{y}_{d-1} \in \partial B$ such that $|\{\mathbf{y}_1, 2\mathbf{y}_1, \dots, k\mathbf{y}_1\} + \dots + \{\mathbf{y}_{d-1}, 2\mathbf{y}_{d-1}, \dots, k\mathbf{y}_{d-1}\}| = k^{d-1}$ (we can choose the points $\mathbf{y}_1, \dots, \mathbf{y}_{d-1}$ one at a time, maintaining the condition $|\{\mathbf{y}_1, 2\mathbf{y}_1, \dots, k\mathbf{y}_1\} + \dots + \{\mathbf{y}_i, 2\mathbf{y}_i, \dots, k\mathbf{y}_i\}| = k^i$ for $i = 1, \dots, d-1$, by observing that after choosing $\mathbf{y}_1, \dots, \mathbf{y}_{i-1}$ each potential equality between two different sums of the form $a_1\mathbf{y}_1 + \dots + a_i\mathbf{y}_i$ with $a_1, \dots, a_i \in \{1, \dots, k\}$ forbids only a single choice for \mathbf{y}_i). Let $\varepsilon > 0$ be sufficiently small, such that any two distinct points in $\{\mathbf{y}_1, 2\mathbf{y}_1, \dots, k\mathbf{y}_1\} + \dots + \{\mathbf{y}_{d-1}, 2\mathbf{y}_{d-1}, \dots, k\mathbf{y}_{d-1}\}$ have Euclidean distance larger than $2k(d-1)\varepsilon$. Now, by applying Lemma 7.4 with $\mathbf{y}_1, \dots, \mathbf{y}_{d-1}$ and ε , we obtain points $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{d-1}$ with $\|\mathbf{x}_i - \mathbf{y}_i\|_2 \leq \varepsilon$ for $i = 1, \dots, d-1$, such that the intersection $\bigcap_{i=1}^{d-1} (\mathbf{x}_i + \partial B)$ is infinite. Now, each point in $\{\mathbf{x}_1, 2\mathbf{x}_1, \dots, k\mathbf{x}_1\} + \dots + \{\mathbf{x}_{d-1}, 2\mathbf{x}_{d-1}, \dots, k\mathbf{x}_{d-1}\}$ is of the form $a_1\mathbf{x}_1 + \dots + a_{d-1}\mathbf{x}_{d-1}$ with $a_1, \dots, a_{d-1} \in \{1, \dots, k\}$ and has distance at most $k(d-1)\varepsilon$ from the corresponding point $a_1\mathbf{y}_1 + \dots + a_{d-1}\mathbf{y}_{d-1}$. By our choice of ε , this means that the points $a_1\mathbf{x}_1 + \dots + a_{d-1}\mathbf{x}_{d-1}$ with $a_1, \dots, a_{d-1} \in \{1, \dots, k\}$ must all be distinct and hence $|\{\mathbf{x}_1, 2\mathbf{x}_1, \dots, k\mathbf{x}_1\} + \dots + \{\mathbf{x}_{d-1}, 2\mathbf{x}_{d-1}, \dots, k\mathbf{x}_{d-1}\}| = k^{d-1}$. Furthermore, defining $S = \bigcap_{i=1}^{d-1} (\mathbf{x}_i + \partial B)$, the set S is infinite and every point $\mathbf{z} \in S$ satisfies $\|\mathbf{z} - \mathbf{x}_i\| = 1$ for $i = 1, \dots, d-1$. \square

We are now ready to prove Theorem 1.2.

Proof of Theorem 1.2. Throughout the argument, we treat d as a fixed constant and all the asymptotics are with respect to large n . Let us choose $m := \left\lceil \log \left(\frac{n}{\log n} \right) \right\rceil = o(n)$ and $k = \lfloor (n/2^m)^{1/(d-1)} \rfloor = \Theta((\log n)^{1/(d-1)})$, so that when writing $n' = k^{d-1} \cdot 2^m$, we have $n \geq n' \geq k^{d-1}/(k+1)^{d-1} \cdot n = (1 - o(1))n$. Let $\|\cdot\|$ be a d -norm.

By Lemma 7.5 there exist points $\mathbf{x}_1, \dots, \mathbf{x}_{d-1} \in \mathbb{R}^d$ with $|\{\mathbf{x}_1, 2\mathbf{x}_1, \dots, k\mathbf{x}_1\} + \dots + \{\mathbf{x}_{d-1}, 2\mathbf{x}_{d-1}, \dots, k\mathbf{x}_{d-1}\}| = k^{d-1}$ and an infinite subset $S \subseteq \mathbb{R}^d$ such that $\|\mathbf{z} - \mathbf{x}_i\| = 1$ for all $\mathbf{z} \in S$ and $i = 1, \dots, d-1$. Now, let us choose

points $\mathbf{z}_1, \dots, \mathbf{z}_m \in S$ such that the set

$$\{\mathbf{x}_1, 2\mathbf{x}_1, \dots, k\mathbf{x}_1\} + \dots + \{\mathbf{x}_{d-1}, 2\mathbf{x}_{d-1}, \dots, k\mathbf{x}_{d-1}\} + \{0, \mathbf{z}_1\} + \dots + \{0, \mathbf{z}_m\} \quad (2)$$

has size $k^{d-1} \cdot 2^m = n'$ (this is possible by once again choosing the points $\mathbf{z}_1, \dots, \mathbf{z}_m \in S$ one by one, observing that at each step every potential equality between two different sums of the form $a_1\mathbf{x}_1 + \dots + a_{d-1}\mathbf{x}_{d-1} + b_1\mathbf{z}_1 + \dots + b_j\mathbf{z}_j$ with $a_1, \dots, a_{d-1} \in \{1, \dots, k\}$ and $b_1, \dots, b_j \in \{0, 1\}$ forbids at most one choice for \mathbf{z}_j).

Now, note that every vector of the form $\mathbf{z}_j - \mathbf{x}_i$ with $i \in \{1, \dots, d-1\}$ and $j \in \{1, \dots, m\}$ occurs as the difference of at least $(k-1) \cdot k^{d-2} \cdot 2^{m-1}$ pairs of points in the set in (2). Since $\|\mathbf{z}_j - \mathbf{x}_i\| = 1$ for all $i \in \{1, \dots, d-1\}$ and $j \in \{1, \dots, m\}$, this means that among the n' points in the set in (2), there are at least

$$(d-1)m \cdot (k-1)k^{d-2} \cdot 2^{m-1} = \frac{d-1}{2} \cdot \left(1 - \frac{1}{k}\right) \cdot (\log n' - (d-1) \log k) \cdot n' = \frac{d-1-o(1)}{2} \cdot n' \log n'$$

unit distances. Since $n' \log n' = (1 - o(1))n \log n$ and $n' \leq n$, this shows there exists a set of at most n points in \mathbb{R}^d with at least $\frac{d-1-o(1)}{2} \cdot n \log n$ unit distances according to the norm $\|\cdot\|$ (by adding more points we can form such a set with exactly n points). \square

8 Concluding remarks and open problems

We have shown in Theorem 1.2 that for every d -norm $\|\cdot\|$, we have $U_{\|\cdot\|}(n) \geq \frac{1}{2}(d-1-o(1)) \cdot n \log n$. In other words, \mathbb{R}^d contains n points that determine at least $\frac{1}{2}(d-1-o(1)) \cdot n \log n$ unit distances according to $\|\cdot\|$, where the $o(1)$ -term tends to zero as $n \rightarrow \infty$. This is nearly tight, as we have proved in Theorem 1.1 that there are d -norms in which no set of n points determines more than $\frac{1}{2}d \cdot n \log n$ unit distances. In fact, almost all d -norms satisfy this property. For $d = 2$ this settles a problem raised by Brass in [6] and by Matoušek in [34], shaving a $\log \log n$ factor from Matoušek's upper bound. For general dimension $d \geq 2$, this provides an essentially tight estimate, up to a $(1 - 1/d - o(1))$ constant factor, and settles, in a strong and somewhat surprising form, Problems 4 and 5 in [8, p. 195].

For $d = 2$ we have shown in Proposition 7.1 that $U_{\|\cdot\|}(n) \geq (\frac{1}{\log 3} - o(1)) \cdot n \log n$ for every 2-norm $\|\cdot\|$, improving upon a well-known and often repeated lower bound of $\frac{1}{2} \cdot n \log n$ coming from the embedding of the hypercube described in the introduction. Our upper bound in Theorem 1.1 for most 2-norms $\|\cdot\|$ is $U_{\|\cdot\|}(n) \leq n \log n$, so there is still a gap between the constant factors in the lower and upper bound. In fact, the constant factor $\frac{1}{\log 3}$ in our new lower bound can be slightly improved. To do so, note that the construction in the proof of Proposition 7.1 is a Minkowski sum of sets $S_i = \{0, \mathbf{x}_i, \mathbf{y}_i\}$. When choosing these sets, the choices of the points \mathbf{x}_i are essentially arbitrary, and it is thus possible to choose them so that for every $i \geq 1$, the points $\mathbf{x}_{4i-3} + \mathbf{x}_{4i-2}$ and $\mathbf{x}_{4i-1} + \mathbf{x}_{4i}$ have unit distance according to $\|\cdot\|$. Similarly, an additional (tiny) improvement can be obtained by repeating this argument recursively. As this still leaves a gap between the upper and lower bounds, we omit the detailed computation. It may be interesting to close the gap between our upper and lower bounds for all dimensions, although our bounds in Theorems 1.1 and 1.2 are already quite close (in particular, the bounds get closer as the dimension increases). It seems that this will require some new ideas.

The proof of Theorem 1.1 applies the fact, established in Section 5, that for most d -norms $\|\cdot\|$ there cannot be too many linear dependencies over the rationals between the unit vectors (more precisely, for a given number of unit vectors, their linear span over the rationals cannot contain too many other unit vectors). It is worth noting, however, that some such linear dependencies always exist. In particular, the triangles constructed in the proof of Lemma 7.2 show that in every strictly convex 2-norm there are infinitely many triples of unit vectors whose sum is 0. Similarly, a simple application of the Hurewicz Dimension Lowering Theorem (Theorem 7.3) implies that for every d -norm $\|\cdot\|$, there are infinitely many $(d-1)$ -tuples $(\mathbf{x}_1, \dots, \mathbf{x}_{d-1})$ of points of unit norm

according to $\|\cdot\|$ such that the $(d-2)$ -tuple of differences $(\mathbf{x}_2 - \mathbf{x}_1, \dots, \mathbf{x}_{d-1} - \mathbf{x}_1)$ is the same for all these $(d-1)$ -tuples $(\mathbf{x}_1, \dots, \mathbf{x}_{d-1})$. In particular, this means that for any $\ell \geq d-1$, we can find a set of ℓ unit vectors in \mathbb{R}^d whose linear span over the rationals contains at least $(d-1) \cdot (\ell - d + 2)$ unit vectors according to $\|\cdot\|$ (in comparison, the arguments in Section 5 show that for a typical d -norm such a span can contain at most $d \cdot \ell$ unit vectors).

Theorem 1.3 provides nearly tight bounds for the largest possible value of $D_{\|\cdot\|}(n)$ for a d -norm $\|\cdot\|$ when n is sufficiently large as a function of d . As far as we know it may be true that an even stronger bound holds, namely that for every fixed d , if n is sufficiently large as a function of d , then there is a d -norm $\|\cdot\|$ so that $D_{\|\cdot\|}(n) = n - 1$. This might even hold for most norms. A proof of this, if true, would require some additional arguments. A careful inspection of the computation in our existing proof shows that it implies that for any d and n , we have $D_{\|\cdot\|}(n) \geq n - O(dn^{3/4})$ for most d -norms $\|\cdot\|$. When n is not large as a function of d , then the determination of the largest possible value of $D_{\|\cdot\|}(n)$ among all d -norms $\|\cdot\|$ is a difficult problem. In particular, the determination of the largest value of n so that $D_{\|\cdot\|}(n) = 1$ for all d -norms $\|\cdot\|$ is equivalent to a well-known conjecture of Petty, as we describe next.

The number of vertices in the largest clique that can be embedded in the unit distance graph of $\|\cdot\|$ is called the equilateral number of $\|\cdot\|$, which we denote here by $e_{\|\cdot\|}$ (equilateral numbers have also been considered in other types of spaces, see e.g. [16]). A conjecture of Petty raised in 1971 ([36], see also [8]) asserts that $e_{\|\cdot\|} \geq d + 1$ for every d -norm $\|\cdot\|$. Note that an equivalent formulation of this conjecture is that $D_{\|\cdot\|}(n) = 1$ for every d -norm $\|\cdot\|$ and every $n \leq d + 1$ (or equivalently, for $n = d + 1$).

This is known in dimension $d \leq 3$ ([36]), and is also known for norms that are sufficiently close to the d -dimensional Euclidean norm ([7], [12]), or more generally to the d -dimensional ℓ_p -norm ℓ_p^d for any $1 < p \leq \infty$ ([44]). Combining this with a result of [2] asserting that every d -dimensional normed space contains a subspace of dimension $r = e^{\Omega(\sqrt{\log d})}$ which is either close to ℓ_2^r or to ℓ_∞^r , it was proved by Swanepoel and Villa [44] that the equilateral number $e_{\|\cdot\|}$ of any d -norm $\|\cdot\|$ is at least $e^{b\sqrt{\log d}}$ for some absolute constant $b > 0$. It is interesting to note that the exact value of the equilateral number $e_{\|\cdot\|}$ is not known even for some simple d -norms like ℓ_1^d , where it is conjectured to be $2d$. See [3], [42], [43] for more information.

Petty's problem as well as the constructions in Section 7 suggest the investigation of graphs that appear as subgraphs in the unit distance graph of any d -norm for some given $d \geq 2$. Any hypercube graph (of any dimension) is an example of such a graph, and a clique of size $e^{b\sqrt{\log d}}$ is another such example (for the absolute constant $b > 0$ in the above-mentioned result of Swanepoel and Villa [44]). The complete bipartite graph $K_{d-1, m}$, for any m , is also an example, as proved in Lemma 7.5. Yet another example that appears as a subgraph of the unit distance graph of any strictly convex 2-norm is the k -th power of a triangle (for any $k \geq 1$), as shown in Lemma 7.2. A characterization of all graphs that are subgraphs of the unit distance graph of any d -norm appears to be difficult.

A related problem is to determine the smallest possible chromatic number of the unit distance graph of a d -norm. In dimension $d = 2$, the answer is 2 and in fact we showed in Theorem 1.4 that for most 2-norms the chromatic number is equal to 4. By the known results about Petty's conjecture, the chromatic number of the unit distance graph of any d -norm is always at least $e^{\Omega(\sqrt{\log d})}$. On the other hand, it is not difficult to show that the chromatic number of the unit distance graph of any d -norm is at most exponential in d (see e.g. [25], [32]). Theorem 6.1 implies an upper bound of 2^d for most d -norms (which gives a better exponential base than the known bounds for all norms). It would be interesting to determine whether the chromatic number of the unit distance graph of a typical d -norm is exponential or sub-exponential in d . By our results here these graphs are rather sparse, hence one may suspect that the chromatic number might be smaller than exponential.

Let us note that the book of Brass, Moser and Pach [8] serves as a remarkable repository of interesting open problems in discrete geometry, many of which have natural extensions to general normed spaces. At least some

of the ideas and tools developed in this paper may be helpful in attacking these questions.

Many of these questions can naturally be posed for typical norms, which in some sense play an analogous role as random graphs do in extremal graph theory. Similarly as with random graphs, our results showcase that considering typical norms can provide answers to extremal questions for general norms. They also show that the Euclidean norm is very special in many regards, and exhibits a different behavior with respect to many of these problems than most other norms. In particular, by the known bounds for $D_{\|\cdot\|_2}(n)$ and $U_{\|\cdot\|_2}(n)$ mentioned in the introduction, our results show that the behaviour for the Euclidean 2-norm with respect to the unit distance problem and the distinct distances problem is very different from the behaviour for a typical 2-norm, and this difference only becomes more pronounced in higher dimensions. While this might be natural in view of the symmetry of the Euclidean norm, we find it surprising that in comparison, for a typical d -norm $\|\cdot\|$, $U_{\|\cdot\|}(n)$ is so small and $D_{\|\cdot\|}(n)$ is so large.

An intriguing open question is to describe explicitly a d -norm $\|\cdot\|$ for which $U_{\|\cdot\|}(n) = O(dn \log n)$ or for which $D_{\|\cdot\|}(n) \geq (1 - o(1))n$ (for large n). Note that a formal statement of this question requires a definition of the notion “explicit” here, a natural one is a norm $\|\cdot\|$ for which there is an efficient algorithm (in any standard model of computation over the reals) for computing the norm $\|\mathbf{x}\|$ of any given vector $\mathbf{x} \in \mathbb{R}^d$. Note also that it is not even obvious that there exists an explicit d -norm as above.

As another application of our methods, we can resolve an analogue of a classical question of Erdős and Moser [21] for typical norms. In 1959, Erdős and Moser asked about the maximum possible number of unit distances among n points in (strictly) *convex* position in the Euclidean plane, and conjectured that the answer is $O(n)$. Füredi [24] proved an upper bound of the form $O(n \log n)$, and the current best bound is still of this form (Aggarwal [1] improved Füredi’s bound by a constant factor). The best-known lower bound of $2n - 7$ is due to Edelsbrunner and Hajnal [13].

Our arguments show that the conjecture of Erdős and Moser is true for most planar norms. More concretely, they imply that for most norms the maximum number of unit distances among n points in strictly convex position in the plane is at most $4n$. Indeed, the arguments in Section 5 imply that for most planar norms, for a list of unit vectors $\mathbf{u}_1, \dots, \mathbf{u}_\ell$, the span $\text{span}_{\mathbb{Q}}(\mathbf{u}_1, \dots, \mathbf{u}_\ell)$ contains a total of at most 2ℓ unit vectors in different directions (where we consider two unit vectors to have the same direction if they agree up to sign). For any n points in strictly convex position in the plane, we can consider the graph corresponding to the unit distances between these points and choose a maximal spanning forest in this graph (i.e. we choose a spanning tree for each component of the graph). Let $\mathbf{u}_1, \dots, \mathbf{u}_\ell$ be the unit vectors corresponding to the edges of this spanning forest, and note that $\ell \leq n - 1$. Now, any unit vector appearing as a distance between two of the n points can be written as a sum of some of the vectors $\mathbf{u}_1, \dots, \mathbf{u}_\ell$ or their negatives (since the two endpoints of the corresponding edge of the graph can be connected by a path in the spanning forest). Hence any unit vector appearing as a distance between the n points must be in $\text{span}_{\mathbb{Q}}(\mathbf{u}_1, \dots, \mathbf{u}_\ell)$, and so for a typical norm there can be at most $2\ell \leq 2n - 2$ different directions of such unit vectors. On the other hand, since the n points are in strictly convex position, for any unit vector \mathbf{u} , there can be at most two pairs of points with distance $\pm\mathbf{u}$. Hence, for a typical planar norm, there can be a total of at most $4n - 4$ unit distances among n points in strictly convex position.

Another interesting open problem we suggest is the possible existence of a zero-one law for typical d -norms: Is it true that for every fixed $d \geq 2$ and every fixed graph H , exactly one of the following two options holds?

- For all d -norms $\|\cdot\|$ besides a meagre set, the unit distance graph of $\|\cdot\|$ contains H as a subgraph.
- For all d -norms $\|\cdot\|$ besides a meagre set, the unit distance graph of $\|\cdot\|$ does not contain H as a subgraph.

Finally, we mention that we have another approach for upper-bounding $U_{\|\cdot\|}(n)$ for typical d -norms $\|\cdot\|$ yielding a somewhat weaker upper bound than in Theorem 1.1, namely of the form $O(d^2 n \log n)$. This approach is more similar to Matoušek’s proof [34] showing $U_{\|\cdot\|}(n) \leq O(n \log n \log \log n)$ for most 2-norms $\|\cdot\|$. In particular,

by a probabilistic argument using a careful multiple exposure process, we manage to improve the graph-theoretic statement in [34, Proposition 2.1], removing the $\log \log n$ factor in this proposition (which causes the $\log \log n$ factor in Matoušek’s result). While the resulting bound $O(d^2 n \log n)$ is weaker (in terms of the d -dependence) than the bound in Theorem 1.1, we plan to write the proof of this improved graph theoretic lemma in a companion note to this paper, since we believe that the argument might be useful in other settings. In particular, a somewhat stronger variant of this lemma would yield an asymptotic answer to the so-called discrete X-Ray reconstruction problem, see [35] for more details.

Remark. While our paper was under review, Greilhuber, Schildkraut, and Tidor [27] found constructions improving the lower bound for the function $U_{\|\cdot\|}(n)$ from Theorem 1.2 to $(d/2 - o(1)) \cdot n \log_2 n$. This shows that our Theorem 1.1 is asymptotically tight.

Acknowledgments. We thank Karim Adiprasito, Pankaj Agarwal, Xiaoyu He, Bo'az Klartag, Mehtaab Sawhney, József Solymosi, Terence Tao and Or Zamir for useful conversations, and János Kollár and Assaf Naor for leading us to reference [17]. The second author gratefully acknowledges the support of the Oswald Veblen fund. We are also very grateful to the anonymous referees for carefully reading the paper and for their helpful comments and suggestions.

References

- [1] A. Aggarwal, *On unit distances in a convex polygon*, Discrete Math. **338**, (2015), no. 3, 88–92.
- [2] N. Alon and V. D. Milman, *Embedding of ℓ_∞^k in finite-dimensional Banach spaces*, Israel J. Math. **45** (1983), no. 4, 265–280.
- [3] N. Alon and P. Pudlák, *Equilateral sets in l_p^n* , Geom. Funct. Anal. **13** (2003), no. 3, 467–482.
- [4] A. Barvinok, *Thrifty approximations of convex bodies by polytopes*, Int. Math. Res. Not. IMRN (2014), no. 16, 4341–4356.
- [5] B. Bollobás and I. Leader, *Edge-isoperimetric inequalities in the grid*, Combinatorica **11** (1991), no. 4, 299–314.
- [6] P. Brass, *Erdős distance problems in normed spaces*, Comput. Geom. **6** (1996), no. 4, 195–214.
- [7] P. Brass, *On equilateral simplices in normed spaces*, Beiträge Algebra Geom. **40** (1999), no. 2, 303–307.
- [8] P. Brass, W. Moser, and J. Pach, *Research problems in discrete geometry*, Springer, New York, 2005.
- [9] K. B. Chilakamarri, *Unit-distance graphs in Minkowski metric spaces*, Geom. Dedicata **37** (1991), no. 3, 345–356.
- [10] J. Davies, *Odd distances in colourings of the plane*, Geom. Funct. Anal. **34** (2024), 19–31.
- [11] A. D. N. J. de Grey, *The chromatic number of the plane is at least 5*, Geombinatorics **28** (2018), no. 1, 18–31.
- [12] B. V. Dekster, *Simplexes with prescribed edge lengths in Minkowski and Banach spaces*, Acta Math. Hungar. **86** (2000), no. 4, 343–358.
- [13] H. Edelsbrunner and P. Hajnal, *A lower bound on the number of unit distances between the vertices of a convex polygon*, J. Combin. Theory Ser. A **56** (1991), no. 2, 312–316.
- [14] J. Edmonds, *Minimum partition of a matroid into independent subsets*, J. Res. Nat. Bur. Standards Sect. B **69B** (1965), 67–72.
- [15] J. Edmonds and D. R. Fulkerson, *Transversals and matroid partition*, J. Res. Nat. Bur. Standards Sect. B **69B** (1965), 147–153.
- [16] C. Elsholtz and W. Klotz, *Maximal dimension of unit simplices*, Discrete Comput. Geom. **34** (2005), no. 1, 167–177.

- [17] R. Engelking, *Theory of dimensions finite and infinite*, Sigma Series in Pure Mathematics, vol. 10, Heldermann Verlag, Lemgo, 1995.
- [18] P. Erdős, *On some applications of graph theory to geometry*, Canadian J. Math. **19** (1967), 968–971.
- [19] P. Erdős, *Problems and results in combinatorial geometry*, Discrete geometry and convexity (New York, 1982), Ann. New York Acad. Sci., vol. 440, New York Acad. Sci., New York, 1985, pp. 1–11.
- [20] P. Erdős, *On sets of distances of n points*, Amer. Math. Monthly **53** (1946), 248–250.
- [21] P. Erdős and L. Moser, *Problem 11*, Canad. Math. Bull. **2** (1959), 43.
- [22] G. Exoo and D. Ismailescu, *The chromatic number of the plane is at least 5: a new proof*, Discrete Comput. Geom. **64** (2020), no. 1, 216–226.
- [23] P. Frankl and R. M. Wilson, *Intersection theorems with geometric consequences*, Combinatorica **1** (1981), no. 4, 357–368.
- [24] Z. Füredi, *The maximum number of unit distances in a convex n -gon*, J. Combin. Theory Ser. A **55** (1990), 316–320.
- [25] Z. Füredi and J.-H. Kang, *Covering the n -space by convex bodies and its chromatic number*, Discrete Math. **308** (2008), no. 19, 4495–4500.
- [26] J. Garibaldi, A. Iosevich, and S. Senger, *The Erdős distance problem*, Student Mathematical Library, vol. 56, American Mathematical Society, Providence, RI, 2011.
- [27] J. Greilhuber, C. Schildkraut, and J. Tidor, *More unit distances in arbitrary norms*, preprint arXiv:2410.07557 (2024).
- [28] P. M. Gruber, *Convex and discrete geometry*, Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol. 336, Springer, Berlin, 2007.
- [29] L. Guth and N. H. Katz, *On the Erdős distinct distances problem in the plane*, Ann. of Math. (2) **181** (2015), no. 1, 155–190.
- [30] G. Kalai, *Some old and new problems in combinatorial geometry I: around Borsuk’s problem*, Surveys in combinatorics, London Math. Soc. Lecture Note Ser., vol. 424, Cambridge Univ. Press, 2015, pp. 147–174.
- [31] V. Klee, *Some new results on smoothness and rotundity in normed linear spaces*, Math. Ann. **139** (1959), 51–63 (1959).
- [32] A. Kupavskiy, *On the chromatic number of \mathbb{R}^n with an arbitrary norm*, Discrete Math. **311** (2011), no. 6, 437–440.
- [33] J. Matoušek, *The dawn of an algebraic era in discrete geometry?*, Proceedings of the 27th European Workshop on Computational Geometry, 2011.
- [34] J. Matoušek, *The number of unit distances is almost linear for most norms*, Adv. Math. **226** (2011), no. 3, 2618–2628.
- [35] J. Matoušek, A. Přívětivý, and P. Škovroň, *How many points can be reconstructed from k projections?*, SIAM J. Discrete Math. **22** (2008), no. 4, 1605–1623.
- [36] C. M. Petty, *Equilateral sets in Minkowski spaces*, Proc. Amer. Math. Soc. **29** (1971), 369–374.
- [37] D. H. J. Polymath, *Polymath16*, web.archive.org/web/20220216001534/https://asone.ai/polymath/index.php?title=Hadwiger-Nelson_problem.
- [38] A. Sheffer, *Polynomial methods and incidence theory*, Cambridge Studies in Advanced Mathematics, vol. 197, Cambridge University Press, Cambridge, 2022.
- [39] J. Solymosi and E. Szabó, *Arrangements of translates of a curve*, preprint arXiv:2208.05525 (2022).
- [40] J. Solymosi and V. H. Vu, *Near optimal bounds for the Erdős distinct distances problem in high dimensions*, Combinatorica **28** (2008), no. 1, 113–125.
- [41] J. Spencer, E. Szemerédi, and W. Trotter, Jr., *Unit distances in the Euclidean plane*, Graph Theory and Combinatorics (Cambridge, 1983), Academic Press, London, 1984, pp. 293–303.
- [42] K. J. Swanepoel, *A problem of Kusner on equilateral sets*, Arch. Math. (Basel) **83** (2004), no. 2, 164–170.

- [43] K. J. Swanepoel, *Combinatorial distance geometry in normed spaces*, New trends in intuitive geometry, Bolyai Soc. Math. Stud., vol. 27, János Bolyai Math. Soc., Budapest, 2018, pp. 407–458.
- [44] K. J. Swanepoel and R. Villa, *A lower bound for the equilateral number of normed spaces*, Proc. Amer. Math. Soc. **136** (2008), no. 1, 127–131.
- [45] L. A. Székely, *Crossing numbers and hard Erdős problems in discrete geometry*, Combin. Probab. Comput. **6** (1997), no. 3, 353–358.
- [46] E. Szemerédi, *Erdős’ unit distance problem*, Open problems in mathematics, Springer, 2016, pp. 459–477.
- [47] P. Ungar, *$2N$ noncollinear points determine at least $2N$ directions*, J. Combin. Theory Ser. A **33** (1982), no. 3, 343–347.
- [48] P. Valtr, *Strictly convex norms allowing many unit distances and related touching questions*, Manuscript (2005).