

Bounds on distinct and repeated dot product trees

Aaron Autry, Slade Gunter, Christopher Housholder, and Steven Senger

Abstract We study questions inspired by Erdős' celebrated distance problems with dot products in lieu of distances, and for more than a single pair of points. In particular, we study point configurations present in large finite point sets in the plane that are described by weighted trees. We give new lower bounds on the number of distinct sets of dot products serving as weights for a given type of tree in any large finite point set. We also demonstrate the existence of many repetitions of some special sets of dot products occurring in a given type of tree in different constructions, narrowing gap between the best known upper and lower bounds on these configurations.

1 Introduction

1.1 Background

In 1946, in [10], Paul Erdős posed a problem that continues to inspire mathematicians: In a large finite set of points in the plane, what is the upper bound of the number pairs can determine the same distance? This is often called the unit distance problem. A related problem also described in [10] is the so-called distinct distances problem, which asks for a lower bound on the number of distinct distances that must be determined by any large finite set of points in the plane.

For the unit distance problem, after some initial activity, general progress seems to have halted in 1984, at the upper bound of a constant times $n^{\frac{4}{3}}$, given by Szemerédi and Trotter in [24]. In the case of the distinct distances problem, there was much activity and incremental progress up until 2010, when Guth and Katz essentially resolved the problem in [13]. For a more detailed look, see [12] and [13], and the references contained therein.

Though Erdős originally asked about pairs of points determining a fixed distance, there are a number of well-studied variants involving k -tuples of points determining other values, such as direction, angle, or dot product. Moreover, there are many different settings besides finite point sets in $\mathbb{R}^2$, such as fractal sets, higher dimensions, or even different types of vector spaces and modules over finite fields and rings. For some examples and discussions, see [6, 9, 16, 18, 28] and the references contained therein. Also, we will assume that the dot products discussed below are all nonzero, unless otherwise specified, as zero dot products can exhibit quite distinct geometric structure. See [14, 18, 20] for more on the unique geometric properties of zero dot products.

We highlight for the moment on the specific problem of the number of distinct dot products determined by a large finite point set in the plane. Given a set of n points in the plane, a straightforward argument using the celebrated Szemerédi-Trotter theorem (from [27]) guarantees the existence of at least a constant times $n^{\frac{2}{3}}$ distinct dot products. We prove this below as Corollary 1. This was the best-known bound for many years until the recent developments by Hanson, Roche-Newton, and the fourth listed author in [15], where an exponent of $\frac{2}{3} + \frac{1}{2739}$ was achieved by relying on tools from additive combinatorics.

1.2 Setting

In this note, our primary focus will be trees. We now describe the setup in detail. For a given large finite point set $E \subseteq \mathbb{R}^d$, with $d \geq 2$, we consider G_E the complete edge-weighted graph whose vertices are points from E and whose edge weights are determined by the dot product between the associated vertices. Now, given a tree T , we estimate two things: how many different sequences of edge weights correspond to instances of T as a subgraph in G_E , and how many times a specific sequence of edge weights occurs as the weights for a copy of T in G_E . The former is analogous to the distinct distances problem, while the latter is akin to the unit distance problem.

To this end, we introduce some notation. Given a tree T with vertices $(v_1, \dots, v_{k+1})$ we denote its k edges by $\{\{v_{i_1}, v_{i_2}\}, \{v_{i_3}, v_{i_4}\}, \dots, \{v_{i_{2k-1}}, v_{i_{2k}}\}\}$, where the indices for each pair of vertices corresponding to an edge, $\{v_{i_{2j-1}}, v_{i_{2j}}\}$, are indexed in lexicographical order. Now, given a sequence of real number weights $\vec{w} := (\alpha_1, \dots, \alpha_k)$, we can consider the weighted tree $T_{\vec{w}}$ where the edge connecting $v_{i_{2j-1}}$ to $v_{i_{2j}}$ has weight α_j , for $j = 1, \dots, k$. In what follows, we will usually be concerned with edge weights equal to the dot product determined by the pair of points corresponding to the pair of vertices.

There is an abundance of literature on the case that our tree T is a path of length k , sometimes called a k -chain. See for example [2], [11], and [22] for distance chain bounds, and see [3] and [20] for results related to dot product chains. More general trees have also been widely studied, but we direct special attention to [21], by Ou and Taylor, and later, [23], by Dung, Pham, and the fourth listed author, for motivating the results here.

Throughout this paper, we assume that k is like a constant compared to the size of any subset E . To ease exposition, we use the asymptotic symbols $X \lesssim Y$ if $X = O(Y)$, and $X \approx Y$ when $X = \Theta(Y)$. Moreover, we write $X \gtrsim Y$ when for every $\varepsilon > 0$, there exists a constant $C_\varepsilon > 0$ such that $X \gtrsim C_\varepsilon q^\varepsilon Y$.

1.3 Main results

Our main result follows the ideas in [21]. It gives lower bounds for the number of distinct sets of dot products much arise as edge weights for trees in a large finite point set in the plane. We note that there is a hypothesis about the number of points along a line through the origin, as there is a technical obstruction from certain special cases of point sets which we discuss in further detail below.

Theorem 1. *Given a large finite set of points $E \subseteq \mathbb{R}^2$, with no more than $\gtrsim n^{\frac{2}{3}}$ points on any line through the origin, and a tree T on k edges, the number of distinct k -tuples of dot products, $\vec{w}$, for which the tree $T_{\vec{w}}$ is realized in E is at least $\gtrsim n^{\frac{2k}{3}}$.*

The key to proving this result is a pinned dot product result between two sets, stated precisely below as Theorem 6. While this result does not immediately extend to higher dimensions, we can still obtain a nontrivial pinned estimate.

Theorem 2. *Given a large finite set of n points $E \subseteq \mathbb{R}^d$, we have that there exists a point $x \in E$ for which*

$$|\{x \cdot y : y \in E\}| \gtrsim n^{\frac{2}{2d-1}}.$$

The other result we have gives lower bounds on how often a particular set of dot product edge weights can occur in some specific point set constructions. It is based on two different constructions which have different strengths and weaknesses.

Theorem 3. *Given a tree T on k edges, and a large finite natural number n , there exists a k -tuple of dot products $\vec{w}$ and set of n points in $\mathbb{R}^d$ that exhibits*

$$\gtrsim \max \left\{ n^{1 + \frac{k(d-1)}{(d+1)}}, n^{\lceil \frac{k}{2} \rceil} \right\}$$

copies of $T_{\vec{w}}$.

We can see that depending on k and the ambient dimension, one or the other bound may dominate. To compare this with known results, we can consider a special case. Suppose T is a perfect binary tree of height h . It is known that this tree has $2^{h+1} - 1$ vertices, meaning it has $2^{h+1} - 2$ edges. So if we set $k = 2^{h+1} - 2$, Theorem 3 guarantees the existence of a set of n points and a k -tuple of dot products $\vec{w}$ in the plane that has $\gtrsim n^{2^h - 1}$ copies of $T_{\vec{w}}$. This is within one power of n of the upper bounds on the same, given in [14], as the following result will show. We note here that in [14], Theorems 1.12 and 1.13 are correct for c -ary trees when $c = 2$, but have some superficial typos when $c \neq 2$.

Theorem 4. [Theorem 1.13 from [14]] *Given a large finite set of n points $E \subseteq \mathbb{R}^2$, a complete balanced binary tree T of height h , and a $(2^{h+1} - 2)$ -tuple of dot products $\vec{w}$, there cannot be more than $\lesssim n^{2^h}$ copies of $T_{\vec{w}}$ determined by points in E .*

The structure of the rest of the paper is as follows. We first acquaint the reader with our primary tools, then prove the main technical theorem, which is a pinned estimate. We then show how to use this to prove our first main result, which is a lower bound on distinct k -tuples of dot products determined by trees isomorphic to a given tree. Finally, we conclude with two constructions that together establish our second main result, a lower bound on how often a particular k -tuple of dot products determined by trees isomorphic to a given tree can occur.

2 Proofs

2.1 Tools

Given a point $p = (p_1, p_2) \in \mathbb{R}^2 \setminus \{(0, 0)\}$, the set of points $q \in \mathbb{R}^2$ that determine dot product α with p form a line whose slope is perpendicular to the line through both p and the origin. To see this, we simply recall the definition of dot product and rearrange the equation $p_1 q_1 + p_2 q_2 = \alpha$. We sometimes call this the α -line of p , and denote it $\ell_\alpha(p)$. A direct consequence of this definition is that $p \neq q$ implies $\ell_\alpha(p) \neq \ell_\alpha(q)$. We now recall the Szemerédi-Trotter theorem from [27]. Given a point p and a line ℓ , we call the pair (p, ℓ) an incidence if the point p is on the line ℓ .

Theorem 5. *Given a set of n points and m lines in the plane, the number of incidences between the points and lines is at most*

$$\lesssim (mn)^{\frac{2}{3}} + n + m.$$

We now give the following well-known corollary alluded to above to illustrate some of the basic ideas to come.

Corollary 1. *Given a large finite set of n points $E \subseteq \mathbb{R}^2$, the point pairs determine at least $n^{\frac{2}{3}}$ distinct dot products.*

Proof. Fix any $\alpha \neq 0$, construct the set of α -lines $\ell_\alpha(p)$ for every $p \in E$. Any time a point q is on an α -line $\ell_\alpha(p)$, we have that $p \cdot q = \alpha$, and vice versa. So the number of times that the dot product α occurs between pairs of points in the set E is exactly the number of incidences between points and α -lines. Since there were n points, and each determines a unique α -line, we have n lines. So by Theorem 5, we have that α can occur no more than $\lesssim (n^2)^{\frac{2}{3}} = n^{\frac{4}{3}}$ times. Since this is true for any $\alpha \neq 0$, and there are n^2 distinct pairs of points, we get that the number of distinct dot products must be at least

$$\frac{n^2}{n^{\frac{4}{3}}} \gtrsim n^{\frac{2}{3}}.$$

Corollary 1 shows that there are many distinct dot products, but it does not show that any particular point determines many distinct dot products. For this, we introduce another concept. Given a point p and a set of points E , we define the *pinned* dot product set $\Pi_p(E)$, to be the set of dot products determined by p and the points in the set E . In this case, we call the point p a *pin*. Specifically,

$$\Pi_p(E) := \{p \cdot q : q \in E\}.$$

2.2 The pinned result

We now state a pinned version of Corollary 1.

Theorem 6. *Given two subsets E and F of $\mathbb{R}^2$ consisting of n points each, where E has no more than $\gtrsim n^{\frac{2}{3}}$ points on any line through the origin, we have that there is a subset $E' \subset E$ of size $|E'| \geq \frac{1}{2}n$ with the property that for any $x \in E'$, the number of distinct dot products determined by x and the points in F is at least $\gtrsim n^{\frac{2}{3}}$. Moreover, if $E = F$, we get the same results without the hypothesis about the number of points on any line through the origin.*

The proof we give essentially comes from Section 9.2 in [12], but Theorem 6 highlights the pinned nature of the result, which is not explicitly stated in the original text. Moreover, we slightly modify the proof to get a result for two sets of points instead of one, which is what we will need later. Unfortunately, because we are specifying which set the good pin must come from, this brings in the possibility of pathological pairs of point sets. For example, if E is a set of n points on a line through the origin and F is a set of n points on a line perpendicular to the line containing E , then we would have that each point $x \in E$ we have $|\Pi_x(F)| = 1$. However, if $E = F$ above, we will see in the proof that we can drop the hypothesis that no line through the origin has more than $\gtrsim n^{\frac{2}{3}}$ points of E .

In order to prove this result, we will rely on a version of the so-called crossing number lemma. In order to state the lemma, we remind the reader of two concepts. First, a multigraph is a generalization of a graph allowing for multiple distinct edges connecting pairs of vertices. We refer to the number of distinct edges connecting a given pair of vertices as the edge multiplicity of that pair of vertices. Second, given a graph G , its crossing number, denoted $cr(G)$ is the minimum number of pairs of edges that cross taken over any drawing of G in the plane. So, a planar graph has crossing number zero, but any graph that is not planar has a positive crossing number. We now state a multigraph version of the crossing number lemma, from Székely in [26].

Lemma 1. *Given a multigraph G with v vertices, e edges, and a maximum edge multiplicity of m , satisfying $e > 5mv$, we have that*

$$cr(G) \gtrsim \frac{e^3}{mv^2}.$$

With this in tow, we prove Theorem 6.

Proof. We will also assume that the origin is not in E or F , as it can only determine the dot product zero anyway. Construct a graph, G , whose vertex set is the point set $E \cup F$. Note, although we defined a graph G_E related to a point set E above, this graph will have very different edges than G_E . For every point $p \in E$, draw the α -lines for each $\alpha \in \Pi_p(F)$. By definition, for each p , the associated lines drawn will be parallel to one another. Now, for every such $\ell_\alpha(p)$ with more than one point on it, we connect vertices by an edge if their corresponding points are consecutive along that line. Notice that we will disregard such lines with exactly one point on them. However, if p and q lie on the same radial line, it is possible that $\ell_\alpha(p) = \ell_\beta(q)$. In this case, the construction of the graph H could lead to multiple edges between some pairs of vertices. Any time that there are multiple edges between a pair of points, r and s , draw each edge as a smooth curve, $\mathcal{C}$, connecting the appropriate points close enough to the line segment, $\mathcal{L}$, connecting those points so that there are no other points from E in the region bounded by $\mathcal{C}$, $\mathcal{L}$, and the points r and s .

We now deal with a possible degenerate case. Suppose there is a point p whose α -lines contribute much fewer than n edges. In this case, this means that for most of the values of $\alpha \in \Pi_p(F)$, we have that $\ell_\alpha(p)$ only has one point. That would imply that p determines $\gtrsim n$ distinct dot products with points in F , and we would be done. So we can assume that it takes about n edges for each point in E . Therefore, there must be about n^2 edges in H .

Define t to be the maximum number of distinct dot products determined by any point in $p \in E$ with points in F . That is, we define

$$t := \max_{p \in E} \{p \cdot q : q \in F\}.$$

The the bulk of rest of the proof is devoted to showing that $t \gtrsim n^{\frac{2}{3}}$. For any fixed radial line, ℓ , the vertices of consecutive points along each of the parallel lines perpendicular to ℓ will be connected by as many edges as there are points on ℓ . Because of our hypothesis on the set E , we know ℓ has no more than $n^{\frac{2}{3}}$ points from E on it. Thus, in our graph, the maximum edge multiplicity will be $m \lesssim n^{\frac{2}{3}}$.

We pause here to note a slight change that applies to the case where $E = F$ to handle the ‘moreover’ part of the theorem. Namely, if we were to have $E = F$, we would not need the hypothesis on the number of points on a radial line, as if there were more than $n^{\frac{2}{3}}$ points on a radial line, any of those points would determine at least as many distinct dot products as are on the radial line, with the points on the radial line in question, so each such point would automatically satisfy the conclusion of the theorem statement.

Now that we know the number of vertices, edges, and maximum edge multiplicity, we can try to apply Lemma 1. Recall that we need $e > 5mv$ to apply the lemma. If this condition were to fail, we would have $e \leq 5mv$, but plugging in our values

of e, m , and v would give us $n^2 \lesssim 5n^{\frac{2}{3}}n$ which is a contradiction. Therefore, we can proceed assuming that this condition holds, and by plugging in values of e, m , and v , we are left with

$$cr(G) \gtrsim \frac{e^3}{mv^2} \gtrsim \frac{n^6}{n^{\frac{2}{3}}n^2} \gtrsim n^{\frac{10}{3}}.$$

Next, we will pair this with an upper bound for the crossing number of G . Any crossing between edges only occurs when a line perpendicular to the radial line of one point, say p , crosses a line perpendicular to the radial line of another point, say q . Since each point has no more than t such associated parallel lines (its α -lines), each pair of points can contribute at most t^2 crossings. Because we know that there are about n^2 different pairs of points, we are guaranteed that the total number of crossings is bounded above by $\lesssim n^2 t^2$. Now, while we don't know what the minimum number of crossings in any drawing of G in the plane must be, we can be sure that it is no more than $n^2 t^2$, because we have found an explicit drawing with no more than $n^2 t^2$ crossings. Putting the upper and lower bounds for the crossing number together:

$$n^{\frac{10}{3}} \lesssim cr(G) \lesssim n^2 t^2.$$

This tells us that $t \gtrsim n^{\frac{2}{3}}$, as claimed. Now, we have shown the existence of one point $x \in E$ that determines $\gtrsim n^{\frac{2}{3}}$ distinct dot products with points in F . We call this point a *good pin*. To finish the proof, we just need to show that there are many good pins. To see this, define the set $E_1 := E \setminus \{x\}$, and set F_1 to be any subset of F with $n - 1$ points. We can run the same argument to get that there is a good pin $x_1 \in E_1$, that determines many distinct dot products with the points in F_1 . Repeat this process until we have at least $m = \lceil n/2 \rceil$ such good pins, and obtain the set E_m . To finish, set $E' = E_m$, which is of size $\geq n/2$, and consists entirely of good pins.

2.3 Proof of Theorem 2

With a simple pigeonholing argument, we can use Theorem 6 to get a pinned result in higher dimensions in the case that $E = F$.

Proof. Suppose t_1 is the maximum number of parallel $(d - 1)$ -hyperplanes supported by points in E that are the level sets of dot products determined by points in E . That is, these are the higher dimensional analogs of α -lines. Let x be such a point determining t_1 distinct dot products with points in E . Pick one of these parallel $(d - 1)$ -hyperplanes, with as many or more points than any of the other parallel $(d - 1)$ -hyperplanes and call it P_1 . By the pigeonhole principle, we can see that it has at least $\gtrsim n/t_1$ points of E . Now we repeat this process within P_1 . Specifically, we select $x_2 \in P_1$ be a point determining t_2 distinct dot products, so that no other points in P_1 determine more dot products with points in P_1 . As before, pick a $(d - 2)$ -hyperplane P_2 , that contains greater than $(n/t_1)/t_2$ points of E . We continue doing this, getting successive hyperplanes of one dimension lower until P_{d-2} is a plane

with $(n)/(t_1 t_2 \dots t_{d-2})$ points. By appealing to Theorem 6, we see that this plane has a point x_{d-2} that determines $(n/(t_1 t_2 \dots t_{d-2}))^{\frac{2}{3}}$ distinct dot products.

Now we collect what we know about each of the points $x_1, \dots, x_{d-1}$. Namely, x_1 determines t_1 distinct dot products with the points of E , x_2 determines t_2 distinct dot products with the points of E , and for all other $j < d-1$, we have that x_j determines t_j distinct dot products with the points of E . Finally, x_{d-1} determines is at least $\gtrsim ((n)/(t_1 t_2 \dots t_{d-1}))^{\frac{2}{3}}$ distinct dot products with the points in E . Putting all of this together, we have

$$\max_{x \in E} |\Pi_x(E)| \gtrsim \max \left\{ t_1, t_2, \dots, t_{d-2}, ((n)/(t_1 t_2 \dots t_{d-2}))^{\frac{2}{3}} \right\}. \quad (1)$$

Now we consider the case when all of the components are approximately equal, namely if

$$t_1 \approx t_2 \approx \dots \approx ((n)/(t_1 t_2 \dots t_{d-2}))^{\frac{2}{3}}.$$

In this case, we get that

$$t_1 \approx \left(\frac{n}{t_1^{d-2}} \right)^{\frac{2}{3}},$$

which gives us

$$\max_{x \in E} |\Pi_x(E)| \approx t_1 \gtrsim n^{\frac{2}{2d-1}}.$$

Now we deal with the case that one of the t_j is significantly bigger one of the others. It is straightforward to see that this would only increase the lower bound for $|\Pi_x(E)|$. The same is true if the final term, $((n)/(t_1 t_2 \dots t_{d-2}))^{\frac{2}{3}}$, is significantly larger than the t_j . To prove this rigorously, one can set up a Lagrange multiplier argument, or take logarithms of each term in base n and use linear programming.

To summarize, we are able to pigeonhole our previous argument onto higher dimensions by running the algorithm on each hyperplane determined by our points. In general, our lower bound for higher dimensions is:

$$\max_{x \in E} |\Pi_x(E)| \gtrsim n^{\frac{2}{2d-1}}.$$

2.4 Proof of Theorem 1

This argument closely follows the arguments in [23], by Dung, Pham, and the fourth listed author, which was inspired by the work of Ou and Taylor in [21]. The proof of Theorem 1 will follow by applying this more technical result. For brevity, we often conflate a point in $\mathbb{R}^2$ with its associated vertex in a graph. Given a tree T with a vertex v , we will often refer to the pair of them by (T, v) .

In essence, what the theorem says is that if we fix a tree T and isolate a vertex v , then there are many points x that serve as the vertex v in many trees T' , isomorphic

to T , and that these isomorphic copies T' have a wide variation of dot product edge weights.

Theorem 7. *Given two subsets E and F of $\mathbb{R}^2$ consisting of n points each, a tree T with k edges, and an arbitrary vertex v of T , we have that there is a subset $E' \subset E$ of size $|E'| \geq 2^{-k}n - o(n)$ with the property that for any $x \in E'$, the number of distinct edge weight k -tuples, $\vec{w}$, from trees T'_w with (T, v) isomorphic to (T', x) , and $T'_w \setminus \{x\} \subset F$ is at least $\gtrsim n^{\frac{2k}{3}}$.*

We now prove Theorem 7 by induction on k . The base case is when $k = 1$, which holds by applying Theorem 6. We assume that the conclusion holds for any tree on $k + 1$ vertices, and then show that it also holds for any tree on $k + 1$ vertices. We will do this by splitting into two cases: the case where v is a leaf of T (only adjacent to one other vertex of T) and the case where v is not a leaf of T .

If v is a leaf of T , then it has a unique adjacent vertex, which we call u . Notice that the tree $S := T \setminus v$ has k edges, so we can apply the induction hypothesis to get that there is a subset $E' \subset E$ with size at least $2^{-k}n$ such that for all $x \in E'$, there are $\gtrsim n^{\frac{2k}{3}}$ distinct dot product edge weight k -tuples $\vec{w}$ for trees S'_w with x as a vertex so that (S'_w, x) is isomorphic to (S, u) . Recall that for tree isomorphism here, we are not concerned with the edge weights. We are showing that there are many different trees with the same general shape, but different edge weights. Next, we can apply Theorem 6 to the set E' and any subset of $F' \subseteq F$ of size $|E'|$ to get that there is a subset $E'' \subseteq E'$ of size at least $|E'|/2$, so that every point $x \in E''$ determines at least $\gtrsim n^{\frac{2}{3}}$ distinct dot products with points in F' . Putting this all together, we have that there is a set, E'' of size at least $|E'|/2 \geq 2^{-(k+1)}n$ so that any point $x \in E''$ satisfies the conclusion of the theorem. This handles the case that v is a leaf.

We now turn our attention to the case that v is not a leaf. In this case, we will break up T into two trees that share only the vertex v . So T_1 is a tree with at least one edge and v as a vertex, T_2 is also a tree with at least one edge with v as a vertex, and the union of the two trees is T . Suppose the number of edges in T_1 is k_1 , and the number of edges in T_2 is k_2 . As k is the total number of edges in T , and they are all accounted for without repeating by considering both T_1 and T_2 , we can see that $k = k_1 + k_2$.




Figure 1: Here we can see the tree T with vertex v decomposed into T_1 and T_2 , both of which also have vertex v .

Next, we arbitrarily partition E into two subsets of equal size (either both $n/2$ or $(n+1)/2$ and $(n-1)/2$), so we have

$$E = E_1 \sqcup E_2.$$

Next, we similarly partition F into two subsets of nearly equal size, to get F_1 and F_2 . We will now perform the following steps for both E_1 and E_2 separately. In what follows, we will often need various sets to be the same size, but they could be off by one. If this is the case, we merely delete a point from the larger set as needed providing a negligible error considering the fact that n is large.

Because $k_1 \leq k$, we can apply the induction hypothesis to E_i and F_1 with (T_1, v) to find a subset $E'_i \subseteq E_i$ of size at least $2^{-k_1}|E_i|$ with the property that for every $x \in E'_i$, there are at least $\gtrsim n^{\frac{2k_1}{3}}$ distinct dot product k_1 -tuples $\vec{w}_1$ corresponding to trees isomorphic to T_1 with x serving as the vertex v and the other vertices coming from points in F_1 . Now, we find an arbitrary subset $F'_2 \subseteq F_2$ of size $|E'_1|$. Because $k_2 \leq k$, we can again appeal to the induction hypothesis for the sets E'_i and F'_2 with the tree T_2 and vertex v . This guarantees the existence of a set $E''_i \subseteq E'_i$ whose size is at least

$$|E_i''| \geq 2^{-k_2} |E_i'| \geq 2^{-k_2-k_1} |E_i| = 2^{-k} |E_i| \geq 2^{-k-1} n - 1, \quad (2)$$

with the property that for every $y \in E_i''$, there are at least $\gtrsim n^{\frac{2k_2}{3}}$ distinct dot product k_2 -tuples $\vec{w}_2$ corresponding to trees isomorphic to T_2 with y serving as the vertex v and the other vertices coming from points in F_2' . Because $E_i'' \subseteq E_i'$, each of these points $y \in E_i''$ also determines at least $\gtrsim n^{\frac{2k_1}{3}}$ distinct dot product k_1 -tuples $\vec{w}_1$ corresponding to trees isomorphic to T_1 with y serving as the vertex v . Because we chose F_1 disjoint from F_2 , we know that the points in F_1 that serve as vertices for isomorphic copies of T_1 are disjoint from the points from F_2 that serve as vertices for the isomorphic copies of T_2 .

We now generalize a term from before. Here, a *good pin* is a choice of v determining sufficiently many distinct dot product tuples from their respective dot product trees, in analogy to the same term applied to single dot products in the proof of Theorem 6. Recall that we run the above argument for both E_1 and E_2 , which were disjoint. So we can add the good pins v coming from E_1 to the good pins v coming from E_2 and appeal to (2) to get a lower bound on the number good pins to be

$$|E'| \geq |E_1'' \sqcup E_2''| \geq \left(2^{-k-1} n - 1\right) + \left(2^{-k-1} n - 1\right) = 2^{-k} n - 2 = 2^{-k} n - o(1),$$

as claimed.

To finish, notice that by definition, the dot product k -tuple $\vec{w}$ is what we get by combining $\vec{w}_1$ and $\vec{w}_2$. Moreover, F_1 and F_2 were disjoint, so the total number of distinct dot product k -tuples $\vec{w}$ determined by isomorphic copies of T with v coming from E and the other vertices coming from F is at least the product of the number of dot product k_1 -tuples $\vec{w}_1$ associated to isomorphic copies of T_1 from F_1' times the number of dot product k_2 -tuples $\vec{w}_2$ associated to isomorphic copies of T_2 coming from points in F_2' .

$$\gtrsim n^{\frac{2k_1}{3}} \cdot n^{\frac{2k_2}{3}} = n^{\frac{2(k_1+k_2)}{3}} = n^{\frac{2k}{3}}, \quad (3)$$

finishing the induction step, and therefore the proof of Theorem 7.

2.5 Proof of Theorem 3

Theorem 3 will follow from the following two constructions, which we state separately as propositions. The first one is adapted from a construction given in [20].

Proposition 1. *Given a tree T on k edges, there exists a set of n points in $\mathbb{R}^d$ and a k -tuple of dot product edge weights $\vec{w}$ so that there are $\gtrsim n^{\lfloor k/2 \rfloor}$ isomorphic copies of T in the set whose edge weights are $\vec{w}$.*

The second proposition is adapted from a construction given in [19].

Proposition 2. *Given a tree T on k edges, there exists a set of n points in $\mathbb{R}^d$ so that there are $\gtrsim n^{1+k(d-1)/(d+1)}$ isomorphic copies of T in the set whose edge weights are 1.*

Combining these two yields Theorem 3. We now prove each of the propositions.

2.5.1 Proof of Proposition 1

This construction is a modification of a chains construction in [20]. It can be easily modified any number of ways, but we restrict to one way that is relatively easy to communicate, while still showing the essential components that one might want to modify. We begin by 2-coloring the vertices tree T , which we can do because the chromatic number of any tree is at most 2. Call the two bipartition sets U and V , and suppose without loss of generality that $|U| \geq |V|$. Now let k_1 denote the size of U and let k_2 denote the size of V . So $k_1 + k_2 = k + 1$. Recall that by definition, we know $k_1 = |U| \geq \lceil k/2 \rceil$. Now we proceed with the following algorithm. Pick any vertex in U and call it u_1 .

We will now put points on the standard xy -coordinate plane. Place $\lfloor (n - k_2)/k_1 \rfloor$ points along the line $x = 1$ in the plane. Now suppose that u_1 has a_1 neighbors in T called $v_{1,1}$ through v_{1,a_1} . For each such neighbor $v_{1,j}$, place a point at $(1 + j, 0)$. So any of the points along the line $x = 1$ have dot product $1 + j$ with the point at $(1 + j, 0)$. Now, call the b neighbors of $v_{1,1}$ (other than u_1) $u_2, u_3, \dots, u_b$. For each such neighbor u_j , place $\lfloor (n - k_2)k_1 \rfloor$ points on the line $x = 1 + a_1 + j$. Continue by placing single points along the x -axis for vertices from V , and columns of $\lfloor (n - k_2)/k_1 \rfloor$ points for each vertex from U , keeping track of the relative dot products. After all of the vertices in U and V have been processed, place any other remaining points in such a way that they do not overlap with previously placed points, until you have a total of n points, and the construction is finished.




Figure 2: Here we can see the tree T with vertices broken up into U and V , along with a set of points in $\mathbb{R}^2$ that correspond to it. Note for example, the point at $(6, 0)$ will always be a choice for $v_{2,1}$, and its dot product with any point in the column of points with $x = 8$, each of which is a possible choice for u_4 , will have dot product 48, so we have the dot product edge weight of 48 on the edge connecting $v_{2,1}$ and u_4 .

2.5.2 Proof of Proposition 2

This construction is essentially the discrete part of the construction in [19]. We describe this construction in full detail for a self-contained exposition. The new piece here is how to embed the trees. Also, in Proposition 2, we stress that the dot products are always 1, as opposed to the previous construction, for Proposition 1, where our choices of dot product edge weights might have to be different. The basic idea is to construct two sets, E and F , where every point in either set makes the dot product 1 with many points in the other set. The set E will be a kind of lattice, so it will have many points on many hyperplanes. The set F will be the points that make dot product 1 with many of the hyperplanes. We now make this precise.

For a large natural number n , fix $q \approx n^{\frac{1}{d+1}}$. Now define two sets of numbers:

$$A := \left\{ \frac{q+i}{2q} : i = 1, \dots, q \right\},$$

and

$$B := \left\{ \frac{q^2+i}{2q^2} : i = 1, \dots, q^2 \right\}.$$

To build our lattice, E , we will take the $(d-1)$ -fold Cartesian product of A with itself, and then take the Cartesian product of this set with B . By construction, we can see that $|E| = q^{d-1} \cdot q^2 \approx n$. Next, we focus on the point set F , which will also be built using the sets of numbers A and B . We use $(-A)$ to denote the negatives of each element of A . We write $(-A) := \{-a : a \in A\}$. With this in tow, we define the point set F :

$$F := \left\{ \left(\frac{-c_1}{b}, \frac{-c_2}{b}, \dots, \frac{-c_{d-1}}{b}, \frac{1}{b} \right) : c_j \in (-A), b \in B \right\}.$$

Now we want to get a handle on point pairs whose dot product is 1. Dot products are constant on hyperplanes, so given a d -tuple, $(c_1, c_2, \dots, c_{d-1}, b)$, we can define the hyperplane

$$h(c_1, c_2, \dots, c_{d-1}, b) := \left\{ (x_1, x_2, \dots, x_d) \in \mathbb{R}^d : x_d = \left(\sum_{j=1}^{d-1} c_j x_j \right) + b \right\}.$$

We will now focus on the following set of hyperplanes

$$H := \{h(c_1, c_2, \dots, c_{d-1}, b) : c_j \in (-A), b \in B\}.$$

These hyperplanes are the sets of points that determine dot product 1 with points in F . To prove this, consider a d -tuple, $(c_1, c_2, \dots, c_{d-1}, b) \in (-A)^{d-1} \times B$, and compute the dot product of the associated element of F and any point on the hyperplane associated to the same d -tuple. To be precise, consider the point $f \in F$ given by

$$f = \left(\frac{-c_1}{b}, \frac{-c_2}{b}, \dots, \frac{-c_{d-1}}{b}, \frac{1}{b} \right) \text{ to be associated to } h_f := h(c_1, \dots, c_{d-1}, b),$$

a hyperplane in H , and compute its dot product with an arbitrary $x \in h_f$. Their dot product will be

$$\begin{aligned} f \cdot x &= \left(\frac{-c_1}{b}, \frac{-c_2}{b}, \dots, \frac{-c_{d-1}}{b}, \frac{1}{b} \right) \cdot \left(x_1, x_2, \dots, x_{d-1}, \left(\sum_{j=1}^{d-1} c_j x_j \right) + b \right) \\ &= \left(\sum_{j=1}^{d-1} \frac{-c_j x_j}{b} \right) + \left(\sum_{j=1}^{d-1} \frac{c_j x_j}{b} \right) + \frac{b}{b} = 1. \end{aligned}$$

We now consider any $f \in F$. The associated hyperplane, $h_f \in H$, consists of points that have dot product 1 with f . By construction, for most¹ $f \in F$, the associated hyperplane h_f will have $\approx q^{d-1}$ points from E . Moreover, the dual relation also holds. That is, for most $e \in E$, there are about q^{d-1} points $f \in F$ such that $e \cdot f = 1$.

Now we embed our trees as in the proof of the previous result. For any tree T , we 2-color the vertices to obtain disjoint vertex sets U and V . Fix some vertex from U to be u_1 . Let any point $f_1 \in F$ be a choice for u_1 . We consider the set of q^{d-1} choices of $e \in h_{f_1}$ that have dot product 1 with f_1 . Suppose that u_1 has a_1 neighbors in T , and designate these by $v_{1,1}, v_{1,2}, \dots, v_{1,a_1}$. For each of these, we select a different point in h_{f_1} . Notice that there are q^{d-1} choices for each of them. Now consider the neighbors of $v_{1,1}$ in T , and for each representative point $e_{1,1} \in h_{f_1}$ of each of them, choose from the q^{d-1} points of F that make dot product 1 with $e_{1,1}$. Continuing in this way, each new vertex that is processed in the tree has $\approx q^{d-1}$ possible choices. Since there were $\approx n$ choices for u_1 , and each of the k subsequent points had $q^{d-1} \approx n^{\frac{d-1}{d+1}}$ choices, we end up with a total of $\gtrsim n^{1+k(d-1)/(d+1)}$ isomorphic copies of T in the set whose edge weights are 1, as claimed.

References

1. D. Barker and S. Senger, *Upper bounds on pairs of dot products*, Journal of Combinatorial Mathematics and Combinatorial Computing, Volume 103, November, 2017, pp. 211–224.
2. M. Bennett, A. Iosevich, and K. Taylor, *Finite chains inside thin subsets of $\mathbb{R}^d$* , Analysis and PDE, volume 9, no. 3, (2016).
3. V. Blevins, D. Crosby, E. Lynch, and S. Senger, *On the Number of Dot Product Chains in Finite Fields and Rings*, Nathanson M. (eds) Combinatorial and Additive Number Theory V. CANT 2021. pp. 1–20. Springer Proceedings in Mathematics & Statistics. Springer, Cham.
4. J. Bourgain, N. Katz, and T. Tao, *A sum-product estimate in finite fields, and applications*, Geom. Funct. Anal. **14** (2004), pp. 27–57.
5. P. Brass, W. Moser, and J. Pach, *Research Problems in Discrete Geometry*, Springer (2000), 499 pp.
6. J. Chapman, B. Erdoğan, D. Hart, A. Iosevich, D. Koh, *Pinned distance sets, k -simplices, Wolff’s exponent in finite fields and sum-product estimates*, Math. Z. 271 (2012), no. 1-2, 63–93.
7. D. Covert and S. Senger, *Pairs of dot products in finite fields and rings*, Nathanson M. (eds) Combinatorial and Additive Number Theory II. CANT 2015, CANT 2016. Springer Proceedings in Mathematics & Statistics, vol 220. Springer, Cham. (Appeared 14 Jan. 2018)
8. D. Covert, A. Iosevich, J. Pakianathan, *Geometric configurations in the ring of integers modulo p^ℓ* , Indiana University Mathematics Journal, 61 (2012), 1949–1969.
9. D. Covert, D. Hart, A. Iosevich, S. Senger, and I. Uriarte-Tuero, *An analog of the Furstenberg-Katznelson-Weiss theorem on triangles in sets of positive density in finite field geometries*, Discrete Math. 311, no. 6, 423–430, (2011).
10. P. Erdős, *On sets of distances of n points*, Amer. Math. Monthly **53** (1946) 248–250.
11. N. Frankl and A. Kupavskii, *Almost sharp bounds on the number of discrete chains in the plane*, Combinatorica 42 (Suppl 1), 1119–1143 (2022).

¹ Here we notice that there could be some members of f whose hyperplanes cut through the edge of E , and therefore hit significantly fewer than $\approx q^{d-1}$ points, but this will happen for $\lesssim |F|$ points. As this is a negligible loss, we ignore it and similar instances of it.

12. J. Garibaldi, A. Iosevich, and S. Senger, *Erdős distance problem*, AMS Student Library Series, 56, (2011).
13. L. Guth and N. H. Katz, *On the Erdős distinct distance problem in the plane*, Annals of Math., Pages 155–190, Volume 181 (2015), Issue 1.
14. S. Gunter, E. Palsson, B. Rhodes, and S. Senger, *Bounds on point configurations determined by distances and dot products*, Nathanson M. (eds) Combinatorial and Additive Number Theory IV. CANT 2019, CANT 2020. Springer Proceedings in Mathematics & Statistics, vol 347. Springer, Cham.
15. B. Hanson, O. Roche-Newton, and S. Senger, *Convexity, superquadratic growth, and dot products*, Journal of the London Mathematical Society, Volume 107, Issue 5, May 2023, pp. 1900–1923.
16. D. Hart, A. Iosevich D. Koh, M. Rudnev, *Averages over hyperplanes, sum-product theory in vector spaces over finite fields and the Erdős-Falconer distance conjecture*, Trans. Amer. Math. Soc. 363 (2011), no. 6, 3255–3275.
17. A. Iosevich and M. Rudnev, *Erdős-Falconer distance problem in vector spaces over finite fields*, Trans. Amer. Math. Soc., 359 (2007), no. 12, 6127–6142.
18. A. Iosevich and S. Senger, *Orthogonal systems in vector spaces over finite fields*, Electronic J. of Combinatorics, Volume 15, December (2008).
19. A. Iosevich and S. Senger, *Falconer-type estimates for dot products*, Bulletin of the Hellenic Mathematical Society **64** 98–110 (2020).
20. S. Kilmer, C. Marshall, and S. Senger, *Dot product chains*, (accepted) arXiv:2006.11467 (2020).
21. Y. Ou and K. Taylor, *Finite point configurations and the regular value theorem in a fractal setting*, Indiana Univ. Math. J. 71 (2022), no. 4, 1707–1761.
22. E. Palsson, S. Senger, and A. Sheffer, *On the number of discrete chains*, (to appear in the Proceedings of the American Mathematical Society).
23. T. Pham, S. Senger, and D. Tran, *Distribution of pinned distance trees in the plane $\mathbb{F}_p^2$* , Discrete Mathematics, Volume 346, Issue 12, December 2023, 113613.
24. J. Spencer, E. Szemerédi, and W. T. Trotter, *Unit distances in the Euclidean plane* B. Bollobás, editor, “Graph Theory and Combinatorics”, pages 293–303, Academic Press, New York, NY, (1984).
25. J. M. Steele, *The Cauchy-Schwarz Master Class ICM Edition: An Introduction to the Art of Mathematical Inequalities*, Cambridge University Press, (2010).
26. L.A. Székely, *Crossing numbers and hard Erdős problems in discrete geometry*, Combin. Probab. Comput. **6** (1997), no. 3, 353–358.
27. E. Szemerédi and W. T. Trotter, *Extremal problems in discrete geometry*, Combinatorica **3** (1983), 381–392.
28. T. V. Pham and L. A. Vinh *Orthogonal Systems in Vector Spaces over Finite Rings*, Electronic J. of Combinatorics, Volume 19, Issue 2 (2012).