Distinct Distances with ℓ_p Metrics

Moaaz AlQady, Samuel Speas, William Dudarov

January 7, 2023

Project Participants

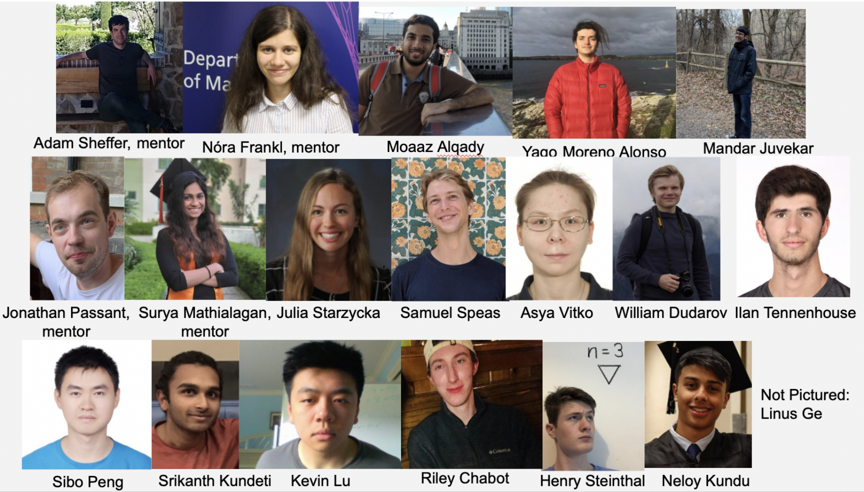

Adam Sheffer, mentor
Nóra Frankl, mentor
Moaaz Alqady
Yago Moreno Alonso
Mandar Juvekar

Jonathan Passant, mentor
Surya Mathialagan, mentor
Julia Starzycka
Samuel Speas
Asya Vitko
William Dudarov
Ilan Tennenhouse

Sibo Peng
Srikanth Kundeti
Kevin Lu
Riley Chabot
Henry Steinthal
Neloy Kundu

Not Pictured:
Linus Ge

Distinct Distances

How many distinct distances are there in the following configurations?

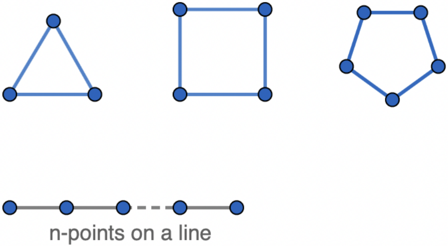

The Erdős Distinct Distances Problem

Let $\mathcal{P}$ be a set of n points in $\mathbb{R}^2$.
Define $D(\mathcal{P})$ to be the set of distances spanned by pairs of points in $\mathbb{P}$.
What is the minimum possible cardinality of $D(\mathcal{P})$?
In particular, we care about finding an asymptotic formula relying on n.
Erdős conjectured that this should be

$$D(n) := \min_{|\mathcal{P}|=n} |D(\mathcal{P})| = \Theta\left(\frac{n}{\sqrt{\log(n)}}\right),$$

and he did show that it is $O(n/\sqrt{\log(n)})$.
The conjecture was very nearly proven by Larry Guth and Nets Katz in 2010 - they achieved the asymptotic

$$D(n) = \Omega\left(\frac{n}{\log(n)}\right).$$

The Erdős Distinct Distances Problem(s)

There are lots of variants of this problem, most of which remain wide open.

1. Erdős also considered the problem in $\mathbb{R}^d$, $d > 2$, and the problem of characterizing the point sets that achieve an optimal number of distances, both of which remain wide open.
2. The problem has also been considered in complex spaces, spaces over finite fields, and for bipartite point sets, only considering distances from a point in one part to a point in the other.
3. Julia Garibaldi, in her doctoral dissertation under the supervision of Terence Tao, studied distinct distances with other metrics.

Distinct Distances with ℓ_p Metrics

Let $p > 2$ be an integer.
The ℓ_p distance between points $a = (a_x, a_y)$ and $b = (b_x, b_y)$ is

$$d_p(a,b) = (|a_x - b_x|^p + |a_y - b_y|^p)^{\frac{1}{p}}.$$

We can then talk about $D_p(n)$, the minimum number of distinct distances for an n point set in the plane under the ℓ_p metric.
In the 1990s, Székely showed for the original ℓ_2 Erdős distinct distances problem

$$D(n) = D_2(n) = \Omega(n^{4/5}).$$

Garibaldi adapted this proof to yield that

$$D_p(n) = \Omega(n^{4/5}).$$

Distinct Distances with ℓ_p Metrics

In our work, we adapted Székely's proof to yield that

$$D_p(n) = \Omega(n^{6/7-\varepsilon})$$

for any $\varepsilon > 0$.

Furthermore, we take steps towards another open problem: characterizing sets that span few distinct distances.

We characterize the sets that span an asymptotically minimal number of distinct distances under the ℓ_1 and ℓ_∞ metrics in the plane.

Székely's Proof

First, we will discuss Székely's original proof for ℓ_2 distances.

Theorem (Székely, 1993)

Let $\mathcal{P}$ be a set of n points in the plane.
Then

$$D(\mathcal{P}) = \Omega(n^{4/5}).$$

Proof.

Bla bla bla...

Székely's Proof

First, we will discuss Székely's original proof for ℓ_2 distances.

Theorem (Székely, 1993)

Let $\mathcal{P}$ be a set of n points in the plane.
Then

$$D(\mathcal{P}) = \Omega(n^{4/5}).$$

Proof.

More bla bla bla...

Székely's Proof

First, we will discuss Székely's original proof for ℓ_2 distances.

Theorem (Székely, 1993)

Let $\mathcal{P}$ be a set of n points in the plane.
Then

$$D(\mathcal{P}) = \Omega(n^{4/5}).$$

Proof.

And more bla bla bla... End of proof. □

Theorem

Let $\mathcal{P}$ be a set of n points in the $\mathbb{R}^2$ and let $p > 2$ be an integer. Under the $\updownarrow_p$ metric, for and $\varepsilon > 0$, there exists a point $u \in \mathcal{P}$ such that Then

$$D_p(u, \mathcal{P}) = \Omega(n^{6/7-\varepsilon}).$$

Sketch of proof: We count cr(G) in two different ways...

First, we have two useful results:

Theorem (Sharir & Zahl, 2017)

Let $\mathcal{P}$ be a set of points and let Γ be a set of n curves that belong to an s-dimensional family of curves of degree at most D, no two sharing an irreducible component. Then for every $\varepsilon > 0$,

$$I(\mathcal{P},\Gamma) = O_{s,D,\varepsilon}\left(m^{\frac{2s}{5s-4}} n^{\frac{5s-6}{5s-4}+\varepsilon} + m^{2/3}n^{2/3} + m + n\right)$$

Lemma (Crossing lemma for multigraphs)

Let $G = (V,E)$ be a multigraph, with maximum edge multiplicity m. Then

$$cr(G) = \Omega\left(\frac{|E|^3}{m|V|^2}\right)$$

Fix some integer $p \geq 2$, and set $t = \max_{u \in \mathcal{P}} D_p(u, \mathcal{P})$. We may assume $t = O(n/\log n)$. There are at most t circles defined by the $p - norm$ which are centered at a given point and incident to another point in $\mathcal{P}$. There are $\mathcal{I}(\mathcal{P}, \mathcal{C}) = n(n-1)$ incidences among all such circles. We can remove the circles incident to at most two points of $\mathcal{P}$, and we still have $I(\mathcal{P}, \mathcal{C}) = \Theta(n^2)$

Consider the graph $G = (V, E)$, where V contains a vertex for each point in $\mathcal{P}$ and E contains an edge between every two pairs of points that are consecutive along a circle in $\mathcal{C}$.

Drawing each each and vertex, edges intersect only at intersection points of the circles of $\mathcal{C}$. Any two circles have at most two intersection points. This implies

$$cr(G) \leq 2\binom{|\mathcal{C}|}{2} \leq 2\binom{nt}{2} = \Theta(n^2 t^2)$$

.

By considering different types of bisectors, we are able to remove edges with high multiplicity, and apply the bound of Sharir's and Zahl's to algebraic curves which comprise the bisectors. We obtain a lower bound on $\mathrm{cr}(G)$...

Structure of Optimal Sets in $\mathbb{R}^2$

Q: What do sets achieving minimum number of distinct distances look like?

- Maybe a lattice? *Too hard!*
- Maybe there is a line with many (i.e. $\Theta(\sqrt{n})$ points on it?
- Maybe there is a line with $\Theta(n^{\varepsilon})$ points on it for some $\varepsilon > 0$?

Theorem

(Citation) There exists a line with $\Theta(\log n)$ *points on it.*

In Other Metrics

The structural problem is much more approachable for ℓ_1 and ℓ_∞ metrics.
Here unit "circles" look like squares. (talk more, less?)
ℓ_1 and ℓ_∞ differ from each other only by rotation and scaling, so WLOG let's focus on ℓ_∞.
First guess: optimal sets look like a lattice.

First guess: lattice

When does a set $\mathcal{P}$ have a "lattice structure"?
Answer: $\mathcal{P} \subset A \times B$ for arithmetic progressions A (think on x-axis) and B (think on y-axis), of size $\Theta(\sqrt{n})$.
This is far from true! Optimal sets can be very different from Cartesian products of arithmetic progressions. Consider the following example:

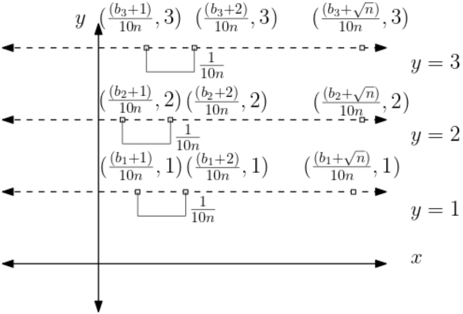

Generalized Arithmetic Progressions

A *generalized arithmetic progression* of *dimension d* is defined as

$$\left\{ a + \sum_{j=1}^{d} k_j b_j : \text{ for integers } 0 \le k_j \le n_j - 1 \text{ for every } 1 \le j \le d \right\},$$

for fixed $a, b_1, b_2, \ldots, b_d \in \mathbb{R}$ and $n_1, \ldots, n_d \in \mathbb{N}$. The size of such a progression is $\prod_{j=1}^{d} n_j$.

Theorem (Polymath 2020)

Let $\mathcal{P}$ be a set of n points such that $D_\infty(\mathcal{P}) = \Theta(\sqrt{n})$. Then:

(a) There exists a set $\mathcal{L}$ of $\Theta(\sqrt{n})$ lines such that $\mathcal{P} \subset \bigcup_{\ell \in} \ell$. Either all lines of $\mathcal{L}$ are horizontal, or all are vertical.

(b) After removing $o(n)$ points from $\mathcal{P}$, we also have:

- Every line $\ell \in \mathcal{L}$ satisfies $|\ell \cap \mathcal{P}| = \Theta(\sqrt{n})$. When thinking of ℓ as $\mathbb{R}$, the points of $\ell \cap \mathcal{P}$ are contained in a generalized arithmetic progression of dimension $\Theta(1)$ and size $\Theta(\sqrt{n})$.
- There exist generalized arithmetic progressions $A_1, \ldots, A_s$ with $s = \Theta(1)$, each of constant dimension and size $O(\sqrt{n})$, that satisfy the following. For every $\ell \in \mathcal{L}$, when thinking of ℓ as $\mathbb{R}$, there exists $r \in \mathbb{R}$ and $1 \leq j \leq s$ with $|(\ell \cap \mathcal{P}) \cap (r + A_j)| = \Theta(\sqrt{n})$.
- The lines of $\mathcal{L}$ can be partitioned into $\Theta(1)$ disjoint subsets, each of size $\Theta(\sqrt{n})$. In each subset, the intercepts of the lines (with the orthogonal axis) are contained in a generalized arithmetic progression of dimension $\Theta(1)$ and size $\Theta(\sqrt{n})$.

Corollary

Let $\mathcal{P}$ be a set of n points such that $D_\infty(\mathcal{P}) = \Theta(\sqrt{n})$. Then there exist a subset $\mathcal{P}' \subseteq \mathcal{P}$ of $\Theta(n)$ points with the following properties:

- There exists a set $\mathcal{L}$ of $\Theta(\sqrt{n})$ lines, either all vertical or all horizontal. Every line $\ell \in \mathcal{L}$ satisfies $|\ell \cap \mathcal{P}'| = \Theta(\sqrt{n})$.
- There exists a generalized arithmetic progression A with of dimension $\Theta(1)$ and size $\Theta(\sqrt{n})$. For every $\ell \in \mathcal{L}$, when thinking of ℓ as $\mathbb{R}$, there exists $r \in \mathbb{R}$ such that $(\ell \cap \mathcal{P}') \subseteq (r + A)$.
- The intercepts of the lines (with the orthogonal axis) are contained in another generalized arithmetic progression of dimension $\Theta(1)$ and size $\Theta(\sqrt{n})$.