Geometrical sets with forbidden configurations

Davi Castro-Silva

University of Cologne

Preprint: arXiv:2102.10018

Introduction

Motivation

General problem: How large can a set be if it does not contain a given geometrical configuration?

- What is the maximum density $\mathbf{m}_{\mathbb{R}^d}(1)$ that a subset of $\mathbb{R}^d$ can have if it does not contain pairs of points at distance 1?

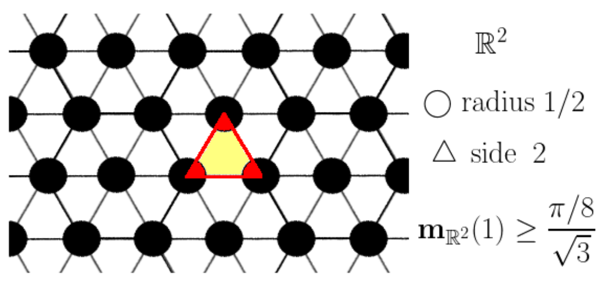

Motivation

General problem: How large can a set be if it does not contain a given geometrical configuration?

- What is the maximum density $\mathbf{m}_{\mathbb{R}^d}(1)$ that a subset of $\mathbb{R}^d$ can have if it does not contain pairs of points at distance 1?
- **Conjecture** (Erdős 82): On the plane, $\mathbf{m}_{\mathbb{R}^2}(1) < 1/4$.

  This conjecture is still open; known: $0.2293 < \mathbf{m}_{\mathbb{R}^2}(1) < 0.2545$ (lower bound by Croft 67, upper bound by Ambrus, Matolcsi 20).
- **Theorem** (Frankl, Wilson 81): The extremal density decays exponentially with the dimension: $\mathbf{m}_{\mathbb{R}^d}(1) \leq (1.2 + o(1))^{-d}$.
- What if we forbid several distances $r_1, r_2, \ldots, r_n$?

  Denote by $\mathbf{m}_{\mathbb{R}^d}(r_1, r_2, \ldots, r_n)$ the maximum density of a set in $\mathbb{R}^d$ avoiding all of these distances.

Motivation

This parameter was first studied by Székely (83), who conjectured:

If $(r_j)_{j\geq 1}$ is not bounded from above, then $\mathsf{m}_{\mathbb{R}^2}\big((r_j)_{j\geq 1}\big) = 0$.

Theorem (Furstenberg, Katznelson, Weiss 90):

If $A \subset \mathbb{R}^2$ has positive density, then there is l_0 such that for any $l \geq l_0$ one can find a pair of points $x, y \in A$ with $\|x - y\| = l$.

Theorem (Bourgain 86):

Let $P \subset \mathbb{R}^d$ be a set of d points spanning a $(d-1)$-dimensional hyperplane. If $A \subset \mathbb{R}^d$ has positive density, then there is l_0 such that A contains a congruent copy of $l \cdot P$ for all $l \geq l_0$.

Motivation

Independence density:

Given finite configurations $P_1, \ldots, P_n \subset \mathbb{R}^d$, let $\mathsf{m}_{\mathbb{R}^d}(P_1, \ldots, P_n)$ be the maximum density a set $A \subset \mathbb{R}^d$ can have without containing congruent copies of any P_i.

- This generalizes $\mathsf{m}_{\mathbb{R}^d}(r_1, \ldots, r_n)$ to higher-order configurations. In particular, $\mathsf{m}_{\mathbb{R}^d}(\{0, e_1\}) \equiv \mathsf{m}_{\mathbb{R}^d}(1)$.
- Natural analogue of *independence number* of the geometrical hypergraph encoding congruent copies of $P_1, \ldots, P_n$ on $\mathbb{R}^d$.
- Bourgain's theorem:

  If $(l_j)_{j\geq 1}$ is not bounded from above, then $\mathsf{m}_{\mathbb{R}^d}\big((l_j P)_{j\geq 1}\big) = 0$.

Motivation

Our goal: Study the independence density and related geometrical parameters.

Among others, we consider the questions:

(Q1) What is the rate of decay of $\mathbf{m}_{\mathbb{R}^d}(t_1 P, t_2 P, \ldots, t_n P)$ with n as the ratios t_{j+1}/t_j between consecutive scales get large?

(Q2) What possible values can be taken by the independence density $\mathbf{m}_{\mathbb{R}^d}(t_1 P, t_2 P, \ldots, t_n P)$ of n distinct dilates of P?

(Q3) Are there analogous results valid for other (non-Euclidean) spaces?

The Euclidean space

Notation

For $x \in \mathbb{R}^d$, $R > 0$ and a measurable set $A \subseteq \mathbb{R}^d$:

- $Q(x, R)$ is the axis-parallel cube of side R and center x:

$$Q(x, R) = x + [-R/2,\ R/2]^d$$

- $d_{Q(x,R)}(A) := \mathrm{vol}(A \cap Q(x, R))/R^d$ is the density of A inside the cube $Q(x, R)$.
- $\overline{d}(A) := \limsup_{R\to\infty} d_{Q(0,R)}(A)$ is the upper density of A.
  If the limit exists, denote it by $d(A)$.
- A **configuration** P is a finite subset of $\mathbb{R}^d$, and is **admissible** if $|P| \leq d$ and it spans a $(|P| - 1)$-dimensional affine hyperplane.

Independence density

Two configurations $P, Q \subset \mathbb{R}^d$ are **congruent** ($P \simeq Q$) if they can be made equal using rigid transformations.

A set $A \subseteq \mathbb{R}^d$ **avoids** P if no subset of A is congruent to P.

Given $n \geq 1$ configurations $P_1, \ldots, P_n \subset \mathbb{R}^d$, we define

$$\mathrm{m}_{\mathbb{R}^d}(P_1, \ldots, P_n) := \sup\left\{\overline{d}(A) : A \subset \mathbb{R}^d \text{ avoids } P_1, \ldots, P_n\right\} \quad \text{and}$$

$$\mathrm{m}_{Q(0,R)}(P_1, \ldots, P_n) := \sup\left\{d_{Q(0,R)}(A) : A \subset Q(0,R) \text{ avoids } P_1, \ldots, P_n\right\}$$

Independence density

Easy bounds: For all finite configurations $P \subset \mathbb{R}^d$:

$$\frac{1}{4^d} \leq \mathsf{m}_{\mathbb{R}^d}(P) \leq 1 - \frac{1}{|P|}$$

$$\frac{\mathsf{m}_{Q(0,R)}(P)}{\left(1 + \frac{\operatorname{diam} P}{R}\right)^d} \leq \mathsf{m}_{\mathbb{R}^d}(P) \leq \mathsf{m}_{Q(0,R)}(P) \qquad \forall R > 0$$

- We analyze $\mathsf{m}_{\mathbb{R}^d}(P)$ by studying $\mathsf{m}_{Q(0,R)}(P)$.

The counting function I_P

I_P counts how many congruent copies of P are contained in A:

$$I_P(A) \cong \int_{Q \text{ congruent to } P} \mathbb{1}\{Q \subset A\}\, dQ$$

Given a configuration $P = \{v_1, v_2, \ldots, v_k\} \subset \mathbb{R}^d$ and a set $A \subseteq \mathbb{R}^d$

$$I_P(A) := \int_{\mathbb{R}^d} \int_{\mathrm{O}(d)} A(x + Tv_1)\, A(x + Tv_2) \cdots A(x + Tv_k)\, d\mu(T)\, dx,$$

where: A is the indicator function of set A,
μ is the uniform (Haar) probability measure on $\mathrm{O}(d)$.

- Note: If $R \gg \operatorname{diam} P$, then $I_P(Q(x, R)) \approx R^d$.
- We similarly define its weighted version $I_P(f)$, for $f : \mathbb{R}^d \to \mathbb{R}$.

Robustness of the counting function

For $\delta > 0$, let $\mathcal{Q}_\delta(x) = \delta^{-d} Q(0,\delta)(x)$. Then

$$A * \mathcal{Q}_\delta(x) = \delta^{-d} \int_{Q(x,\delta)} A(y)\, dy = d_{Q(x,\delta)}(A)$$

is the 'local density' of A at $x \in \mathbb{R}^d$.

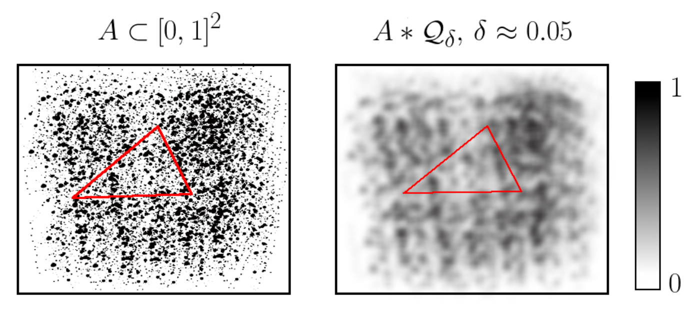

Counting lemma

Motto: "Blurring a set does not significantly change the count of admissible configurations."

Counting lemma (C.S.):

For every admissible $P \subset \mathbb{R}^d$ there exists $C_P > 0$ such that:
For every $R > 0$ and measurable set $A \subseteq Q(0, R)$, we have

$$|I_P(A) - I_P(A * \mathcal{Q}_\delta)| \leq C_P \delta^{1/4} R^d \qquad \forall \delta \in (0, 1].$$

- Important: The bound is *uniform* over all sets $A \subseteq Q(0, R)$.
- Proven using Bourgain's Fourier analytic methods.
- Similar in spirit to counting lemmas in (hyper)graph theory.

Supersaturation theorem

Motto: "If a set is *slightly denser* than the independence density of P, then it must contain *a positive proportion* of all copies of P."

Supersaturation theorem (C.S.):

For every admissible $P \subset \mathbb{R}^d$ and $\varepsilon > 0$ there are $R_0,\, c > 0$ such that:

For all $R \geq R_0$, if $A \subseteq Q(0,R)$ satisfies

$$d_{Q(0,R)}(A) \geq \mathsf{m}_{Q(0,R)}(P) + \varepsilon,$$

then $I_P(A) \geq cR^d$.

- Proven using weak* continuity of the counting function I_P and the Lebesgue Density Theorem.
- Very similar to supersaturation results in (hyper)graph theory.

Theorem (Bukh 08): For any configurations $P_1, P_2, \ldots, P_n$ we have

$$\mathsf{m}_{\mathbb{R}^d}(P_1, P_2, \ldots, P_n) \geq \prod_{i=1}^{n} \mathsf{m}_{\mathbb{R}^d}(P_i).$$

Proof: Take $\varepsilon > 0$ and $R \gg 1$ sufficiently large.

- For $1 \leq i \leq n$, take a P_i-avoiding set A_i with $\overline{d}(A_i) \geq \mathsf{m}_{\mathbb{R}^d}(P_i) - \varepsilon$. We can assume each A_i is periodic with same period R.

$$d_{Q(x, R-\operatorname{diam} P_i)}(A_i) \geq \overline{d}(A_i) - \varepsilon$$

$$A_i' = A_i \cap Q(x, R - \operatorname{diam} P_i)$$

$A_i' + R\,\mathbb{Z}^d$:

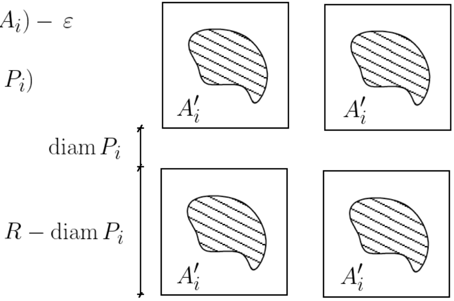

Results on the independence density

Theorem (Bukh 08): For any configurations $P_1, P_2, \ldots, P_n$ we have

$$\mathsf{m}_{\mathbb{R}^d}(P_1, P_2, \ldots, P_n) \geq \prod_{i=1}^{n} \mathsf{m}_{\mathbb{R}^d}(P_i).$$

Proof: Take $\varepsilon > 0$ and $R \gg 1$ sufficiently large.

- For $1 \leq i \leq n$, take a P_i-avoiding set A_i with $\overline{d}(A_i) \geq \mathsf{m}_{\mathbb{R}^d}(P_i) - \varepsilon$. We can assume each A_i is periodic with same period R.
- For any fixed $x_1, \ldots, x_n \in \mathbb{R}^d$ the intersection $\bigcap_{i=1}^n (x_i + A_i)$ avoids each of the P_i.
- If the parameters x_i are chosen independently over $Q(0, R)$, by averaging the expected density of $\bigcap_{i=1}^n (x_i + A_i)$ is $\prod_{i=1}^n d(A_i)$.
- There exist $x_1, \ldots, x_n \in Q(0, R)$ such that

$$d\left(\bigcap_{i=1}^{n} (x_i + A_i)\right) \geq \prod_{i=1}^{n} d(A_i) \geq \prod_{i=1}^{n} \left(\mathsf{m}_{\mathbb{R}^d}(P_i) - \varepsilon\right). \qquad \square$$

Results on the independence density

Intuitively, if $\mathsf{m}_{\mathbb{R}^d}(P_1,\, P_2,\, \ldots,\, P_n) \approx \prod_{i=1}^{n} \mathsf{m}_{\mathbb{R}^d}(P_i)$ then the constraints of avoiding each configuration P_i are *uncorrelated*, or *independent*.

This might be expected to happen if the *natural sizes* of each P_i are very different from each other.

Theorem (C.S.): If $P_1, P_2, \ldots, P_n \subset \mathbb{R}^d$ are admissible, then

$$\mathsf{m}_{\mathbb{R}^d}(t_1 P_1,\, t_2 P_2,\, \ldots,\, t_n P_n) \to \prod_{i=1}^{n} \mathsf{m}_{\mathbb{R}^d}(P_i)$$

as $t_2/t_1,\, t_3/t_2,\, \ldots,\, t_n/t_{n-1} \to \infty$.

- Proven using the counting lemma and the supersaturation theorem, based on an argument of Bukh.

Results on the independence density

Corollary (Q1): If $P \subset \mathbb{R}^d$ is admissible, then

$$\mathsf{m}_{\mathbb{R}^d}(t_1P,\, t_2P,\, \ldots,\, t_nP) \to \mathsf{m}_{\mathbb{R}^d}(P)^n \quad \text{as} \quad t_2/t_1,\, \ldots,\, t_n/t_{n-1} \to \infty.$$

From this we can easily deduce Bourgain's theorem:

$$\mathsf{m}_{\mathbb{R}^d}\left((l_jP)_{j\geq 1}\right) = 0 \text{ for all unbounded positive sequences } (l_j)_{j\geq 1}.$$

Remark: This result is *not* true for non-admissible configurations.

(Graham 94): If $P \subset \mathbb{R}^d$ is *nonspherical*, then

$$\mathsf{m}_{\mathbb{R}^d}\big(P,\, \sqrt{3}P,\, \sqrt{5}P,\, \sqrt{7}P,\, \ldots\big) > 0.$$

Theorem (C.S.): The function $(P_1, P_2, \ldots, P_n) \mapsto \mathsf{m}_{\mathbb{R}^d}(P_1, P_2, \ldots, P_n)$ is continuous on the set of n admissible configurations.

Results on the independence density

(Q2) Characterize all possible independence densities when forbidding n distinct dilates of some configuration P.

$$\mathcal{M}_n(P) := \{\mathbf{m}_{\mathbb{R}^d}(t_1P, t_2P, \ldots, t_nP) : 0 < t_1 < t_2 < \cdots < t_n < \infty\}$$

Theorem (C.S.): If $P \subset \mathbb{R}^d$ is admissible, then

$$\overline{\mathcal{M}_n(P)} = \left[\mathbf{m}_{\mathbb{R}^d}(P)^n, \mathbf{m}_{\mathbb{R}^d}(P)\right].$$

Proof: Clearly $\mathbf{m}_{\mathbb{R}^d}(t_1P, t_2P, \ldots, t_nP) \leq \mathbf{m}_{\mathbb{R}^d}(t_1P) = \mathbf{m}_{\mathbb{R}^d}(P)$, and

$$\mathbf{m}_{\mathbb{R}^d}(t_1P, t_2P, \ldots, t_nP) \geq \prod_{i=1}^{n} \mathbf{m}_{\mathbb{R}^d}(t_iP) = \mathbf{m}_{\mathbb{R}^d}(P)^n.$$

But $\mathbf{m}_{\mathbb{R}^d}(P)$ and $\mathbf{m}_{\mathbb{R}^d}(P)^n$ are accumulation points, and the result follows from continuity. $\square$

The sphere

Notation

Now we work on the d-dimensional unit sphere $S^d \subset \mathbb{R}^{d+1}$:

- σ denotes the uniform probability measure on S^d.
- A **spherical configuration** is a finite subset of $\mathbb{R}^{d+1}$ which *is congruent to* a set on S^d.
- A spherical configuration P is **admissible** if:
  - it has at most d points;
  - it is congruent to a collection $P' \subset S^d$ which is *linearly independent*.
- For $n \geq 1$ spherical configurations $P_1, \ldots, P_n$ we define

$$\mathsf{m}_{S^d}(P_1, \ldots, P_n) := \sup \left\{ \sigma(A) : \ A \subset S^d \text{ avoids } P_1, \ldots, P_n \right\}$$

Contractible configurations

Issue: The unit sphere is not compatible with dilations.

Contractible configurations: Finite sets $P \subset \mathbb{R}^{d+1}$ for which $t\,P$ is a spherical configuration for all $0 < t \leq 1$.

Example: Any set of at most $d+1$ points on S^d.

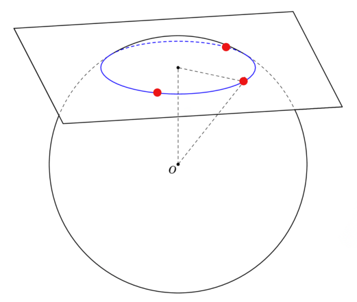

Contractible configurations

Issue: The unit sphere is not compatible with dilations.

Contractible configurations: Finite sets $P \subset \mathbb{R}^{d+1}$ for which tP is a spherical configuration for all $0 < t \leq 1$.

Example: Any set of at most $d+1$ points on S^d.

- Even for contractible configurations, there is no easy relationship between $\mathsf{m}_{S^d}(tP)$ for different $0 < t \leq 1$.

Lemma: For any fixed contractible configuration $P \subset S^d$ we have

$$\inf_{0<t\leq 1} \mathsf{m}_{S^d}(tP) > 0 \quad \text{and} \quad \sup_{0<t\leq 1} \mathsf{m}_{S^d}(tP) < 1.$$

The long and winding road

We can define the counting function

$$I_P(A) = \int_{\mathrm{O}(d+1)} A(Tv_1)\, A(Tv_2) \cdots A(Tv_k)\, d\mu(T)$$

for a set $A \subseteq S^d$ and a configuration $P = \{v_1, \ldots, v_k\} \subset S^d$.

Both the counting lemma and the supersaturation theorem continue to hold on the spherical setting. However:

- The proof of the counting lemma is much more technical.
- The proof of the supersaturation theorem needs first a change of spaces to $\mathrm{O}(d+1)$.
- We need to relate ‘local’ notions of independence density on spherical caps to the ‘global’ notion on S^d.

Results on the spherical independence density

Theorem (C.S.): For all configurations $P_1, P_2, \ldots, P_n \subset S^d$ we have

$$\mathrm{m}_{S^d}(P_1,\, P_2,\, \ldots,\, P_n) \geq \prod_{i=1}^{n} \mathrm{m}_{S^d}(P_i).$$

Theorem (C.S.): If $P_1, P_2, \ldots, P_n \subset S^d$ are admissible and $\varepsilon > 0$, then

$$\left| \mathrm{m}_{S^d}(t_1 P_1,\, t_2 P_2,\, \ldots,\, t_n P_n) - \prod_{i=1}^{n} \mathrm{m}_{S^d}(t_i P_i) \right| < \varepsilon$$

whenever $t_1 \gg t_2 \gg \cdots \gg t_n$.

Corollary (Bourgain's theorem on the sphere):

Suppose $P \subset S^d$ is admissible, and let $A \subset S^d$ be a set with $\sigma(A) > 0$. There exists a number $t_0 > 0$ such that A contains a congruent copy of tP for all $t \leq t_0$.

Results on the spherical independence density

Theorem (C.S.): The independence density $P \mapsto \mathsf{m}_{S^d}(P)$ is continuous on the set of admissible spherical configurations.

Theorem (C.S.): If P is admissible, then there exists a P-avoiding measurable set $A \subset S^d$ attaining $\sigma(A) = \mathsf{m}_{S^d}(P)$.

Remark: Both of these results are *false* for forbidden angles on S^1.

> (DeCorte, Pikhurko 15): The function $\theta \mapsto \mathsf{m}_{S^1}(\theta)$ is *discontinuous* at every rational multiple of 2π with odd denominator.
>
> If $\theta/2\pi$ is irrational, then $\mathsf{m}_{S^1}(\theta) = 1/2$ but there is not θ-avoiding set $A \subset S^1$ with measure 1/2.

Open problems

Open problems

- Do our theorems hold if the configuration has $d+1$ points?

  Conjecture: Let $u, v, w \in \mathbb{R}^2$ be non-collinear points.

  If $A \subset \mathbb{R}^2$ has positive upper density, then there is l_0 such that A contains a congruent copy of $\{lu, lv, lw\}$ for any $l > l_0$.

- *Compatibility condition:* If P is contractible, is it true that $\mathbf{m}_{S^d}(tP) \to \mathbf{m}_{\mathbb{R}^d}(P)$ as $t \to 0$?

- Suppose $A \subset S^d$ contains few copies of a configuration P: $I_P(A) \ll 1$. Can we erase all copies of P by removing from A a set of measure $\ll 1$?

Thank you for your attention!

-Ambrus and Matolcsi, *Density estimates of 1-avoiding sets via higher order correlations*, Discrete Comput. Geom. (2020).

-Bourgain, *A Szemerédi type theorem for sets of positive density in* $\mathbb{R}^k$, Israel J. Math. (1986).

-Bukh, *Measurable sets with excluded distances*, Geom. Funct. Anal. (2008).

-Croft, *Incidence incidents*, Eureka (1967).

-DeCorte and Pikhurko, *Spherical sets avoiding a prescribed set of angles*, Int. Math. Res. Not. (2016).

-Frankl and Wilson, *Intersection theorems with geometric consequences*, Combinatorica (1981).

-Furstenberg, Katznelson and Weiss, *Ergodic theory and configurations in sets of positive density*, in Mathematics of Ramsey theory, Springer (1990).

-Graham, *Recent trends in Euclidean Ramsey theory*, Discrete Math. (1994).

-Székely, *Remarks on the chromatic number of geometric graphs*, in Graphs and other combinatorial topics, Teubner-Texte zur Mathematik (1983).