Bounding the Maximum Size of a Cap Set

By: Amy Qiao

April 20, 2026

Abstract

In this paper, we explore bounds on the maximum size of a cap set in $\mathbb{F}_3^n$. We place a particular focus on the results of Ellenberg and Gijswijt, which prove an upper bound of $2.756\ldots^n$. We also analyze future directions of the study of cap sets and their applications to other areas of combinatorial number theory.

1 Introduction

The study of cap sets holds a captivating place in the interplay of additive combinatorics, discrete mathematics, and theoretical computer science. For decades, the expansive gap between the upper and lower bounds for the size of cap sets remained an open problem. In this paper, we explore properties of cap sets, the breakthrough results of Ellenberg and Gijswijt in solving the Cap Set Conjecture, and applications of the study of cap sets to the Sunflower Conjecture.

The notion of a cap was first introduced in 1947 by Bose, who defined a cap to be a subset of the projective geometry $\mathbb{P}(\mathbb{F}_q^n)$ over a finite field $\mathbb{F}_q$ that contains no three collinear points. Over $\mathbb{F}_3 = \mathbb{Z}/3\mathbb{Z}$, a cap is referred to as a cap set, which is a set of points in $\mathbb{F}_3^n$ containing no non-trivial arithmetic progressions of length three.[1] The cap set problem poses the following question: What is the size of the largest cap set in terms of n?

[1] It seems that Terence Tao incorrectly introduced the term "cap set". In [13], Tao writes, "While on the topic of past inaccuracies, it seems I may have inadvertently propagated a slightly incorrect terminology... sets in $\mathbb{F}_3^n$ free of collinear triples are known as affine caps or simply caps in the design theory literature, rather than cap sets. The latter terminology seems to have become rather entrenched, though, at least in additive combinatorics circles..."

The paper is organized in the following manner: the remainder of the introduction summarizes a tangible example of the cap set problem, connecting the problem to the popular card game SET. In section 2, we begin with a historical overview of obtaining an upper bound for the size of the cap set, starting with the results of Brian and Buhler, proved in 1982, to the Ellenberg-Gijswijt method, which closely followed the techniques of Croot, Lev, and Pach, both proved in 2016. In this section, we will especially see that the proof provided by Ellenberg and Gijswijt takes a unique turn from the typical methods of Fourier analysis that had largely influenced the problem for much of its history. We take a brief detour in section 3 to provide an overview of lower bounds on the maximum size of a cap set. In section 4, we introduce the Cap Set Conjecture along with the polynomial method, which was key to its solution. We then provide a detailed look at the proof of the Ellenberg-Gijswijt method in section 5, drawing connections to the results of Croot, Lev, and Pach. Finally, in section 6, we review applications of the Cap Set Conjecture in other problems, specifically in proving the Sunflower Conjecture.

Acknowledgments

First and foremost, I would like to thank my thesis advisor, Joseph Silverman, and concentration advisor, Isabel Vogt, for their generous guidance and for teaching me to see the full beauty of mathematics. There are no words to truly capture their unparalleled kindness. I am forever indebted to Robert Lewis for his extraordinary guidance throughout these past four years and for his class, which made me pursue mathematics. I would also like to thank Amalia Culiuc, Ellis Hershkowitz, Oanh Nguyen, and Susana Serna. To all of my advisors, thank you for believing in me, even when I did not.

Outside of academics, I am forever grateful to my parents. Their courage and perseverance guide me throughout everything. I would like to thank Sylvester, the best cat a human could ask for. Finally, thank you to Alan Mach for being my light.

1.1 SET

Example 1 (SET)**.** We begin by introducing a tangible example of the cap set problem by drawing connections to the popular card game, SET®. SET was invented by Marsha Jean Falco, a geneticist who used shape, color, fill, and quantity to visualize patterns in genetic qualities relating to epilepsy in German Shepherds. Falco realized that finding these patterns could be a game and debuted SET in 1974. SET took the world by storm, winning the Best

New Mind Game of 1991 award and finding a special place in the mathematics community.

The SET card deck consists of cards with four features, each with three values, leading to a total of $3^4 = 81$ cards in the deck. A Set is a collection of three cards in which for each feature, either all three cards have identical values, or all three cards have distinct values.

The following are examples and non-examples of Sets:

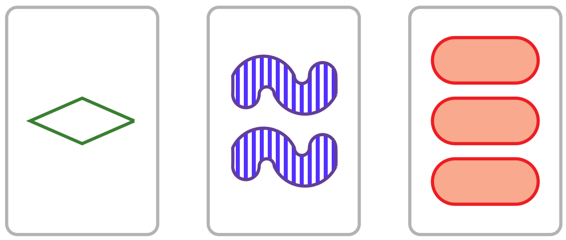

Figure 1: A collection of three cards forming a Set.

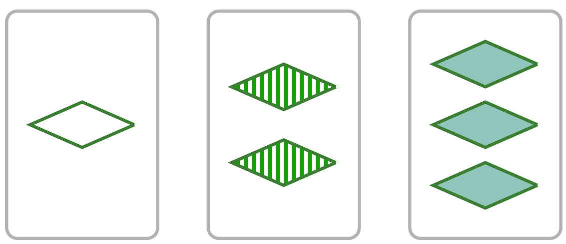

Figure 2: A collection of three cards forming a Set.

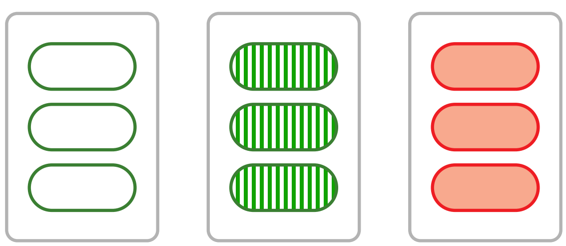

Figure 3: A collection of three cards *not* forming a Set.

Figures 1 and 2 are collections of three cards forming a Set. In Figure 1, each card has a distinct color, fill, quantity, and shape; and in Figure 2, each card has a distinct fill and quantity but the same color and shape. Figure 3 is a collection of three cards which do *not* form a Set. There are two green cards and one red card.

At the start of the game, there is approximately a $\frac{1}{34}$ chance of drawing twelve cards which contain no Set, and approximately a $\frac{1}{2500}$ chance of drawing fifteen cards which contain no Set. The likelihood of not drawing a Set only increases as the game continues [12]. This naturally leads to the question, how many cards must we draw to guarantee a Set? The answer is 21, and what is particularly incredible is that this result was proved four years before the release of SET, through the study of cap sets.

To see the connection between SET, the card game, and the study of cap sets, we redefine each card in SET as a four-dimensional vector over $\mathbb{F}_3$. Since each card has four features, each with three values, we assign each value an element in $\mathbb{F}_3$.

	-1	0	1
Quantity	1	2	3
Color	Red	Green	Purple
Fill	Solid	Striped	Open
Shape	Diamond	Oval	Squiggle

For example, a card with three purple striped squiggles becomes $\langle 1, 1, 0, 1 \rangle$. Therefore, we can represent the SET card deck as the vector space $\mathbb{F}_3^4$. Then, we can define a Set to be a collection of three distinct vectors in $\mathbb{F}_3^4$ which sum to 0, that is, an affine line in $\mathbb{F}_3^4$.

Turning back to our definition of a cap set, we see that over $\mathbb{F}_3^4$, we can think of a cap set as a collection of cards in SET in which no three cards form a Set. Then, asking the question, "how many cards must we draw to guarantee there is a Set?" is identical to asking, "what is the size of the largest cap set in $\mathbb{F}_3^4$?" The answer to the latter is 20, meaning we must draw 21 cards in a standard game of SET to guarantee at least one Set. [2]

We can also extend our SET analogy to an arbitrary n-dimensional vector space over $\mathbb{F}_3$ by modifying our version of the game to have n features. For example, if we additionally give each card one of three types of textures, we can represent the SET card deck as $\mathbb{F}_3^5$.

[2] This was proved by Giuseppe Pellegrino in 1970, predating the initial release of SET by four years.

2 History

Before we introduce the Cap Set Conjecture, we give a historical overview of major advancements in bounding the size of the maximum-sized cap set of $\mathbb{F}_3^n$.

In 1982, Tom Brown and Joe Buhler were the first to show a cap set in $\mathbb{F}_3^n$ must have size $o(3^n)$ [3], a significant improvement from the trivial upper bound of 3^n. Unlike the Fourier analytic approaches that would later dominate the cap set problem, Brown and Buhler proved for odd characteristic finite fields, the density of a subset A that is maximal with respect to the property of not containing three collinear points must be asymptotically zero, that is $|A|/3^n \to 0$ as $n \to \infty$.[3]

In 1995, Roy Meshulam established an upper bound of $\mathcal{O}(3^n/n)$ [11], extending Klaus Roth's 1953 argument for arithmetic progressions in integers to the setting of any finite Abelian group of odd order. Meshulam used Fourier analytic techniques to show that if a set $A \subseteq \mathbb{F}_3^n$ is free of three-term arithmetic progressions, then its density $\delta = |A|/3^n$ must have a decay rate of at least $1/\log_3 n$. This gives rise to the upper bound $|A| \leq C\frac{3^n}{n}$ for some constant C.

This bound stood until Michael Bateman and Nets Katz improved the upper bound to $\mathcal{O}(3^n/n^{1+\epsilon})$ in 2012, for a small positive constant $\epsilon > 0$ [2]. This was the best known upper bound until the results of Jordan Ellenberg and Dion Gijswijt in 2016, who solved the Cap Set Conjecture 1 by proving that for a set $A \subseteq \mathbb{F}_3^n$ without any three-term arithmetic progressions, $|A| \leq C^n$ for a constant $C < 3$ [7].[4] Ellenberg and Gijswijt adapted a breakthrough approach (over $\mathbb{F}_4^n$) by Ernie Croot, Vsevolod Lev, and Péter Pál Pach [4], which uses the polynomial method to demonstrate an exponential decay in the density of a cap set.

In 2021, Zhi Jiang further improved the bound set by Ellenberg and Gijswijt, proving we can gain a factor of $\sqrt{n}$ by examining the multinomial coefficients used in the polynomial method [9].

3 Lower Bounds on the Size of a Maximum Cap Set

To complement the efforts of finding an upper bound on the size of the maximum cap set, identifying a lower bound on its size is essential in narrowing

[3]A subset that is maximal with respect to the property of not containing three collinear points is referred to as an ovaloid in the projective space $\mathbb{P}^n(F)$, where F is a finite field.

[4]Specifically, they proved this holds for $C \approx 2.756$.

the window of its exact size. Lower bounds are typically proved by finding a cap set configuration in a small, fixed dimension n and lifting it to higher dimensions via a product construction. More specifically, we have the following proposition.

Proposition (Cartesian Power of Cap Sets is a Cap Set)**.** *Let A be a cap set in $\mathbb{F}_3^k$. Then, A^n is a cap set in $\mathbb{F}_3^{kn}$.*

Proof. Let A be a cap set in $\mathbb{F}_3^k$. Suppose there exist three points in arithmetic progression in A^n, that is, there exist $x, y, z \in A^n$ such that $x + z = 2y$. Then, equality must hold for each coordinate, meaning for each $i = 1, 2, \ldots, n$, we have that $x_i + z_i = 2y_i$ in A. Given that A is a cap set, A has no non-trivial three-term arithmetic progression. Thus, for each $i = 1, 2, \ldots, n$, we have that $x_i = y_i = z_i$, meaning $x = y = z$. Therefore, A^n also satisfies the property of having no non-trivial three-term arithmetic progression, and we can conclude that A^n is a cap set of $\mathbb{F}_3^{kn}$. □

From the previous proposition, we see that if we can construct a cap set A in $\mathbb{F}_3^k$, then we have a cap set in $\mathbb{F}_3^{kn}$ of size $|A^n| = |A|^n$. Thus, we get an asymptotic lower bound of the form c^n, where $c = |A|^{\frac{1}{k}}$.

In 2004, Yves Edel achieved a significant breakthrough in lower bounds, establishing that a cap set in $\mathbb{F}_3^n$ must have size at least 2.2174^n. In 2022, Fred Tyrell built on the construction of Edel to improve the lower bound to 2.2180^n. We can see that even the best lower bounds are significantly smaller than the currently best proven upper bounds. Although several breakthroughs have been made with respect to identifying the size of the maximum cap set, there remains significant progress to be gained.[5]

4 The Cap Set Conjecture

Definition 1. Let G be an Abelian group. We set

$$r_3(G) = \begin{pmatrix} \text{the size of the largest subset of } G \text{ with} \\ \text{no three-term arithmetic progression} \end{pmatrix}.$$

Conjecture 1 (Cap Set Conjecture)**.**

$$r_3((\mathbb{Z}/3\mathbb{Z})^n) \leq 3^{\lambda n}$$

[5]This is good news for us, as it means we too can contribute to this exciting problem.

for some $\lambda < 1$.

Before 2016, virtually all upper bounds on $r_3(\mathbb{F}_3^n)$ were based on Fourier analytic techniques. Ellenberg and Gijswijt made a sharp departure from previous methods by adapting Croot, Lev, and Pach's argument for $\mathbb{F}_4^n$ to $\mathbb{F}_3^n$, which used the polynomial method to prove that $r_3(\mathbb{F}_4^n) \leq 4^{0.926\ldots \cdot n}$.[6] We formally introduce the polynomial method in the next section.

4.1 Polynomial Method

The polynomial method lies in the intersection of algebraic geometry and combinatorics, using polynomial equations to describe combinatorial structures and the systems of polynomial equations to reason about said structures. Although the framework of the polynomial method has been used since the 1990s, only in the 2010s was a formal framework for the polynomial method developed. Recently, the polynomial method has also been instrumental in proofs of several long-standing open problems, namely the finite field Kakeya problem [6] and the Cap Set Conjecture.[7]

In particular, the polynomial method is based on various extensions of the following facts about univariate polynomials to multivariate polynomials:

1. Every nonzero polynomial of degree d has at most d roots.
2. For every set S of points, there exists a nonzero polynomial f of degree at most $|S|$ which vanishes on S.

Therefore, to obtain an upper bound on the size of a given set S, it suffices to identify a nonzero polynomial f of low degree vanishing on S [10]. In the context of the Cap Set Conjecture, we will use the polynomial method to bound the size of an arbitrary cap set.

[6] In case readers are curious as to why Croot, Lev, and Pach worked over $\mathbb{Z}/4\mathbb{Z}$ rather than $\mathbb{F}_3$ or $\mathbb{F}_p$ for some p prime, which seem much easier to work with, according to personal communication between Joshua Grochow and Lev [8], their initial motivation was to improve Sanders's results in "Roth's Theorem in $\mathbb{F}_4^n$."

[7] With all of the applications of the polynomial method in recent years, one cannot help but ask, what exciting applications of the polynomial method are waiting to be uncovered?

5 Proof of the Cap Set Conjecture

5.1 Preliminaries

We start by introducing some preliminary notation. Let q be an odd prime and $\mathbb{F}_q$ a finite field of order q.

We say a polynomial is reduced in $\mathbb{F}_q$ if each variable has degree at most $q-1$. Let $\mathscr{M}_n$ be the set of reduced monomials in n variables and $\mathscr{S}_n$ be the $\mathbb{F}_q$-vector space spanned by $\mathscr{M}_n$, that is, the space of reduced polynomials in n variables over $\mathbb{F}_q$. Then, let $\mathscr{M}_n^d$ be the set of monomials in $\mathscr{M}_n$ of degree at most d and $\mathscr{S}_n^d$ be the subspace of $\mathscr{S}_n$ spanned by $\mathscr{M}_n^d$, that is, the subspace of reduced polynomials in n variables over $\mathbb{F}_q$ of total degree at most d. Let m_d denote the dimension of $\mathscr{S}_n^d$.

The core argument of Ellenberg and Gijswijt is Proposition 2 of [7], which generalizes Lemma 1 of Croot, Lev, and Pach [4]. We state Lemma 1 of Croot, Lev, and Pach, and then state and prove Proposition 2 of Ellenberg and Gijswijt.

5.2 Lemma 1 of Croot-Lev-Pach

Lemma (Lemma 1 of Croot-Lev-Pach). *Suppose that $n \geq 1$ and $d \geq 0$ are integers, P is a multilinear polynomial in n variables of total degree at most d over a field $\mathbb{F}$, and $A \subseteq \mathbb{F}^n$ is a set with*

$$|A| > 2 \sum_{0 \leq i \leq d/2} \binom{n}{i}.$$

If $P(a-b) = 0$ for all $a, b \in A$ with $a \neq b$, then also $P(0) = 0$.

The proof of Proposition 2 of [7] generalizes the proof of Lemma 1 of [4].[8] For the sake of brevity, we omit the proof of Lemma 1 of [4] and provide the proof for the generalized case in the following proposition.

5.3 Proposition 2 of Ellenberg-Gijswijt

Proposition (Proposition 2 of Ellenberg-Gijswijt). *Suppose $\mathbb{F}_q$ is a finite field and $A \subseteq \mathbb{F}_q^n$. Let α, β, γ be three elements of $\mathbb{F}_q$ which sum to 0. Suppose*

[8]Taking $\alpha = 1, \beta = -1$, and $\gamma = 0$, we see that Lemma 1 of [4] is a special case of Proposition 2 of [7].

$P \in \mathscr{S}_n^d$ satisfies $P(\alpha a + \beta b) = 0$ for every pair a, b of distinct elements in A. Then, the number of $a \in A$ for which $P(-\gamma a) \neq 0$ is at most $2m_{d/2}$.

Proof. Suppose $P \in \mathscr{S}_n^d$, and $x, y \in \mathbb{F}_q^n$. We look at the expansion of $P(\alpha x + \beta y)$. By the definition of $\mathscr{S}_n^d$, P is a linear combination of reduced monomials in n variables of total degree at most d. Therefore, we write

$$P(\alpha x + \beta y) = \sum_{m, m' \in \mathscr{M}_n^d : \deg(mm') \leq d} c_{m,m'} m(x) m'(y).$$

Taking an arbitrary term $c_{m,m'} m(x) m'(y)$, the total degree of $m(x)m'(y)$ is at most d. Since the degree of a product of monomials is the sum of their individual degrees, we have that $\deg(m) + \deg(m') \leq d$. Then, at least one of $\deg(m) \leq d/2$ or $\deg(m') \leq d/2$ must hold.

Thus, we can rearrange the terms in $P(\alpha x + \beta y)$ to write (not necessarily uniquely)

$$P(\alpha x + \beta y) = \sum_{m \in \mathscr{M}_n^{d/2}} F_m(y) m(x) + \sum_{m \in \mathscr{M}_n^{d/2}} G_m(x) m(y)$$

for some collection of polynomials F_m and G_m in $\mathscr{S}_n$ indexed by $m \in \mathscr{M}_n^{d/2}$.

Now, let B be the $|A| \times |A|$ matrix indexed by $a, b \in A$ whose (a, b)-th entry is $P(\alpha a + \beta b)$, that is,

$$B_{ab} = \sum_{m \in \mathscr{M}_n^{d/2}} F_m(b) m(a) + \sum_{m \in \mathscr{M}_n^{d/2}} G_m(a) m(b).$$

Since each term in each summation is a matrix of rank at most 1 and there are $m_{d/2}$ terms in each summation, the rank of B is bounded by $2m_{d/2}$.

By our hypothesis, for all $a, b \in A$ distinct, $P(\alpha a + \beta b) = 0$. This means B is a diagonal matrix. Given that $\alpha + \beta + \gamma = 0$ in $\mathbb{F}_q$, then for all $a \in A$, we have that $\alpha a + \beta a = -(\gamma a)$. Therefore, the diagonal entries of B are given by

$$B_{a,a} = P(\alpha a + \beta a) = P(-\gamma a).$$

Since B is a diagonal matrix, the rank of B is exactly the number of nonzero diagonal entries. Therefore, the number of $a \in A$ for which $P(-\gamma a) \neq 0$ is at most $2m_{d/2}$. □

Now, we move on to the main theorem of Ellenberg and Gijswijt, that any cap set $A \subseteq \mathbb{F}_q^n$ is bounded by $3m_{(q-1)n/3}$.

5.4 Theorem 4 of Ellenberg-Gijswijt

Theorem (Theorem 4 of Ellenberg-Gijswijt)**.** *Let α, β, γ be elements of $\mathbb{F}_q$ such that $\alpha + \beta + \gamma = 0$ and $\gamma \neq 0$. Let $A \subseteq \mathbb{F}_q^n$ such that*

$$\alpha a_1 + \beta a_2 + \gamma a_3 = 0$$

has no solutions (a_1, a_2, a_3) in A^3 apart from trivial solutions of the form $a_1 = a_2 = a_3$. As previously stated, let m_d be the dimension of $\mathscr{S}_n^d$, that is, the number of monomials in $x_1, \ldots, x_n$ with total degree at most d in which each variable has degree at most $q-1$. Then,

$$|A| \leq 3m_{(q-1)n/3}.$$

Proof. Suppose $\alpha, \beta, \gamma \in \mathbb{F}_q$ such that $\alpha+\beta+\gamma = 0$ and $\gamma \neq 0$. Suppose $A \subseteq \mathbb{F}_q^n$ such that the only solutions $(a_1, a_2, a_3) \in A^3$ to the equation $\alpha a_1+\beta a_2+\gamma a_3 = 0$ are trivial solutions of the form $a_1 = a_2 = a_3$. Then,

$$\{\alpha a + \beta b : a, b \in A, a \neq b\} \cap \{-\gamma a : a \in A\} = \emptyset.$$

We denote $-\gamma A = \{-\gamma a : a \in A\}$.

Suppose $d \in \mathbb{Z}$, where $0 \leq d \leq (q-1)n$. Then, suppose $V \subseteq \mathscr{S}_n^d$ is the subspace of polynomials vanishing on $(-\gamma A)^c$, that is,

$$V = \{P \in \mathscr{S}_n^d : P(x) = 0 \text{ for all } x \notin -\gamma A\}.$$

Since any $P \in V$ must vanish on $(-\gamma A)^c = \mathbb{F}_q^n \setminus (-\gamma A)$, which has size $q^n - |A|$, we have that

$$\dim(V) \geq \dim(\mathscr{S}_n^d) - (q^n - |A|) = m_d - q^n + |A|.$$

Given that

$$\{\alpha a + \beta b : a, b \in A, a \neq b\} \cap (-\gamma A) = \emptyset,$$

then

$$\{\alpha a + \beta b : a, b \in A, a \neq b\}$$

is contained in $(-\gamma A)^c$. This means for all $a, b \in A$ distinct, $P(\alpha a + \beta b) = 0$. By Proposition 2, the number of elements $a \in A$ for which $P(-\gamma a) \neq 0$ is at most $2m_{d/2}$.

Let

$$\Sigma_P = \{x \in \mathbb{F}_q^n : P(x) \neq 0\}$$

be the support of P. We have that $|\Sigma_P| \geq \dim(V)$, otherwise, there would exist a nonzero $Q \in V$ vanishing on Σ_P. Since $P \in V$ vanishes on $(-\gamma A)^c$, then its support must be contained in $-\gamma A$, that is, $\Sigma_P \subseteq -\gamma A$. Therefore,

$$|\Sigma_P| \leq 2m_{d/2}$$

and hence,

$$\dim(V) \leq 2m_{d/2}.$$

It follows that

$$m_d - q^n + |A| \leq 2m_{d/2}$$

and thus,

$$|A| \leq 2m_{d/2} + (q^n - m_d).$$

The expression $q^n - m_d$ represents the number of reduced monomials with total degree strictly greater than d. By the mapping of exponents

$$e_i \mapsto (q-1) - e_i,$$

we see that these monomials are in bijection with the monomials of total degree at most $n(q-1) - d - 1$. Therefore,

$$q^n - m_d = m_{n(q-1)-d-1}.$$

Returning to our previous inequality, we have that

$$|A| \leq 2m_{d/2} + m_{n(q-1)-d-1}.$$

Given that $d \in \mathbb{Z}$ with $0 \leq d \leq (q-1)n$ was arbitrary, then, taking $d = \frac{2(q-1)n}{3}$, we have that

$$|A| \leq 2m_{(q-1)n/3} + m_{(q-1)n/3-1} \leq 3m_{(q-1)n/3}. \qquad \square$$

5.5 Corollary 5 of Ellenberg-Gijswijt

Corollary (Corollary 5 of Ellenberg-Gijswijt). *Suppose $A \subseteq (\mathbb{Z}/3\mathbb{Z})^n$ is a cap set, that is, A contains no three-term arithmetic progression. Then,*

$$|A| = o(2.756^n).$$

Proof. Let X be a random variable taking values $0, 1, \ldots, q-1$ with probability $1/q$ each. The expression m_d counts the number of ways to choose n exponents $e_1, \ldots, e_n$ such that each e_i satisfies $0 \leq e_i \leq q-1$ and $\sum e_i \leq d$. Thus,

$$\frac{m_{(q-1)n/3}}{q^n}$$

represents the probability that n independent copies of X have mean at most $\frac{q-1}{3}$. Applying Cramér's Theorem from large deviation theory, we have that

$$\lim_{n\to\infty} \log\left(\frac{m_{(q-1)n/3}}{q^n}\right) = -I\left(\frac{q-1}{3}\right)$$

where I is the rate function of X calculated as

$$I(x) = \sup_{\theta\in\mathbb{R}} \left(\theta x - \log\left(\frac{1 + e^\theta + \ldots + e^{(q-1)\theta}}{q}\right)\right).$$

When $q = 3$ and $x = 2/3$, the supremum is attained when $e^\theta = \frac{\sqrt{33}-1}{8}$, and we obtain the bound $3e^{-I(2/3)} < 2.756$. Applying Theorem 4 with $\alpha = \beta = \gamma = 1$, we can conclude that $|A| = o(2.756^n)$. □

6 Applications and Generalizations

In this section, we explore applications of the Croot-Lev-Pach polynomial method and generalizations of the Cap Set Conjecture.

6.1 Sunflower Conjecture

Definition 2 (Sunflower). A sunflower is a collection of sets $A_1, \ldots, A_k$ such that all pairwise intersections are the same as their k-wise intersection, that is, if $i \neq j$, then

$$A_i \cap A_j = A_1 \cap \ldots \cap A_k.$$

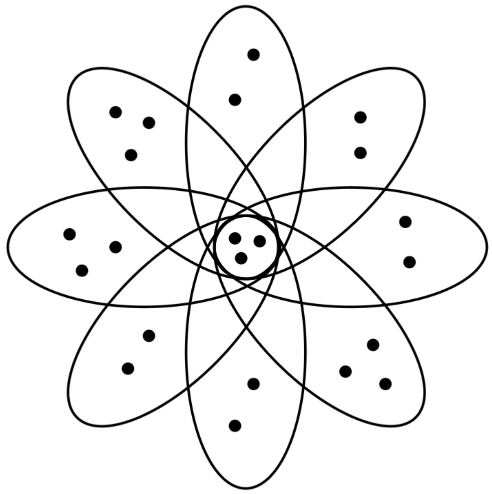

An example of an 8-sunflower.

Conjecture 2 (Sunflower Conjecture)**.** *Suppose $\mathcal{F}$ is a collection of sets, each of size s. For all $k > 0$, there exists a constant c_k such that if $|\mathcal{F}| \geq (c_k)^s$, then $\mathcal{F}$ contains a k-sunflower.*

When studying the Sunflower Conjecture in connection to the complexity of matrix multiplication, Alon, Shpilka, and Umans developed a new notion of sunflowers, which more clearly displays their connection to cap sets.

Definition 3 (Sunflowers in $(\mathbb{Z}/m\mathbb{Z})^n$)**.** A k-sunflower in $(\mathbb{Z}/m\mathbb{Z})^n$ is a collection of k vectors $v_1, \ldots, v_k \in (\mathbb{Z}/m\mathbb{Z})^n$ such that for every coordinate $i \in [n]$, either all vectors v_j have the same value in their i-th coordinate or all i-th coordinates are distinct. Equivalently, for each i, we have that

$$\left|\left\{(v_1)_i, (v_2)_i, \ldots, (v_k)_i\right\}\right|$$

must be either 1 or k.

We then introduce a new notion of the Sunflower Conjecture.

Conjecture 3 (Sunflower Conjecture in $(\mathbb{Z}/m\mathbb{Z})^n$)**.** *For every k, there is a constant b_k such that for every m, n, every set of at least b_k^n vectors in $(\mathbb{Z}/m\mathbb{Z})^n$ contains a k-sunflower.*

While this conjecture seems distinct from the classical Sunflower Conjecture 2, Alon, Shpilka, and Umans proved that the two are indeed equivalent.[1]

We observe that a 3-sunflower in $(\mathbb{Z}/3\mathbb{Z})^n$ is the same as a cap set. Thus, the Cap Set Conjecture solves the Sunflower Conjecture when taking $k = m = 3$. More generally, the same method used to prove the Cap Set Conjecture can be used to prove a weak version of the Sunflower Conjecture in $(\mathbb{Z}/m\mathbb{Z})^n$. Blasiak et al. [5] extended the Ellenberg-Gijswijt bound from vector spaces over $\mathbb{Z}/q\mathbb{Z}$ to $\mathbb{Z}/m\mathbb{Z}$ for arbitrary m, and then for any Abelian group of bounded exponent. Given that this bound depends on m, these results only solve the Sunflower Conjecture in $(\mathbb{Z}/m\mathbb{Z})^n$ for a fixed m. While the full Sunflower Conjecture in $(\mathbb{Z}/m\mathbb{Z})^n$ along with the original Sunflower Conjecture remain open problems, recent applications of the polynomial method in combinatorial number theory have demonstrated new approaches in determining their solutions.

References

[1] N. Alon, A. Shpilka, and C. Umans, On sunflowers and matrix multiplication, computational complexity, 22 (2013), pp. 219–243.

[2] M. Bateman and N. H. Katz, New Bounds on cap sets, 2011.

[3] T. Brown and J. Buhler, A density version of a geometric ramsey theorem, Journal of Combinatorial Theory, Series A, 32 (1982), pp. 20–34.

[4] E. Croot, V. Lev, and P. Pach, Progression-free sets in Z_4^n are exponentially small, 2016.

[5] Drexel University, J. Blasiak, T. Church, Stanford University, H. Cohn, J. A. Grochow, E. Naslund, W. F. Sawin, and C. Umans, On cap sets and the group-theoretic approach to matrix multiplication, Discrete Analysis, (2017).

[6] Z. Dvir, S. Kopparty, S. Saraf, and M. Sudan, Extensions to the Method of Multiplicities, with applications to Kakeya Sets and Mergers, 2009.

[7] J. S. Ellenberg and D. Gijswijt, On large subsets of F_q^n with no three-term arithmetic progression, 2016.

[8] J. Grochow, New applications of the polynomial method: The cap set conjecture and beyond, Bulletin of the American Mathematical Society, 56 (2018), pp. 29–64.

[9] Z. Jiang, Improved explicit upper bounds for the Cap Set Problem, 2021.

[10] S. Jukna, Extremal Combinatorics: With Applications in Computer Science, Texts in Theoretical Computer Science. An EATCS Series, Springer Berlin Heidelberg, Berlin, Heidelberg, 2011.

[11] R. Meshulam, On subsets of finite abelian groups with no 3-term arithmetic progressions, Journal of Combinatorial Theory, Series A, 71 (1995), pp. 168–172.

[12] P. Norvig, The Odds of Finding a SET in The Card Game SET.

[13] T. Tao, A symmetric formulation of the Croot-Lev-Pach-Ellenberg-Gijswijt capset bound, May 2016.