AMSI VACATIONRESEARCH SCHOLARSHIPS 2020–21

Get a Thirst for Research this Summer

Polynomial Methods in Additive Combinatorics

Dibyendu Roy

Supervised by Assoc. Prof Alexander Fish

University of Sydney

Vacation Research Scholarships are funded jointly by the Department of Education, Skills and Employment and the Australian Mathematical Sciences Institute.

AMSI AUSTRALIAN MATHEMATICAL SCIENCES INSTITUTE

Australian Government Department of Education, Skills and Employment

Contents

1 Prelude

1.1 Abstract

We outline the fundamental concepts in Additive Combinatorics and then look at the Cap-*SET* problem which deals with arithmetic progression free sets in $\mathbb{F}_q^n$ where n is large. The case for 3-term arithmetic progression is well understood and we present the elegant proof by Tao.

Then we present new results building off Tao's work that extend the Cap-*SET* problem. We consider sets that contain some proportion of Arithmetic progressions as well as not allowing n points that satisfy some linear equation.

We also prepared a paper [4], and submitted it to a peer-reviewed journal.

1.2 Statement of Authorship

The results, Theorem 5 and Theorem 6 are our (myself and my supervisor) own work and the rest of the results presented are from various sources which are listed in the references.

1.3 Introduction to Additive Combinatorics

Additive Combinatorics is a relatively young field of mathematics (named by Tao and Vu in their 2006 book of the same name) and has many links to other areas such as ergodic theory, graph theory, group theory and, what is of interest to us, polynomial methods.

One of the core concepts in Additive Combinatorics is the notion of the "additive structure" of a set. Let us have an additive group G which is abelian with the group action $+$ (an example of such groups are $(\mathbb{Z}_n, +)$ for $n \in \mathbb{N}$) and let $X, Y \subseteq G$.

An example of a set X that has weak additive structure would be a completely random subset of G with no predefined internal mathematical structure.

However an example of a set with strong additive structure would be an arithmetic progression in G e.g. for $a, r \in G$ and $N \in \mathbb{N}$,

$$Y = \{a,\, a + r,\, a + 2r,\, \ldots,\, a + Nr\}.$$

Furthermore we can define operations on these sets that allow us to start asking interesting questions. We can define the sum set and difference set operations quite naturally as,

Sum set:

$$A + B := \{a + b : a \in A,\, b \in B\}.$$

Difference set:

$$A - B := \{a - b : a \in A,\, b \in B\}.$$

The Cauchy-Davenport Theorem provides us with a lower bound on the size of such sum sets,

Theorem 1 (Cauchy-Davenport)**.** *If $G = (\mathbb{Z}/p\mathbb{Z}, +)$ for some prime p and $A, B \subseteq G$,*

$$|A + B| \geq \min\{|A| + |B| - 1, p\}.$$

This is a famous theorem in Additive Combinatorics and represents one of the core stepping stones to understanding the nature of “additive structure”.

In this report we look at the Cap-*SET* problem which deals with sets without arithmetic progressions, which is another kind of additive structure.

2 The Cap-*SET* problem

2.1 Origins

The Cap-*SET* problem has its roots in *SET* which is a real-time pattern matching card game. The deck has $3^4 = 81$ cards where each card has 4 different attributes, each with three variations as listed in the table.

Shape	Colour	Shading	Number
Rhombus	Yellow	Empty	1
Pill	Red	Striped	2
Squiggle	Blue	Filled	3

Table 1: Table of attributes for cards

We show the possible representations for each attribute in Figure 1 (note that these aren't all the possible cards, just the possible representations).

There is a single dealer and many players and the rules of *SET* are as follows,

- Single dealer, many players
- Dealer places 12 cards on the table
- Players call out "SET!" when they see a *SET* and take the cards
  - A *SET* is 3 cards where each of the attributes are *all same* or *all different*
  - If there is no set the game moves on

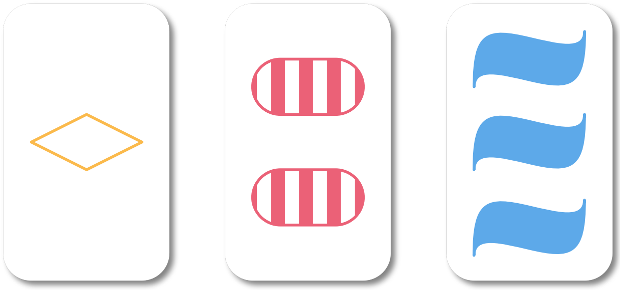

Figure 1: *SET* of 3 cards showing all possible attributes individually

- Dealer adds another 3 cards till all 81 have been dealt
- Winner is the player with the most cards

In Figure 2 we have examples of collections of three cards, one of which is a set and another is not a set. Also note that Figure 1 is also a set.

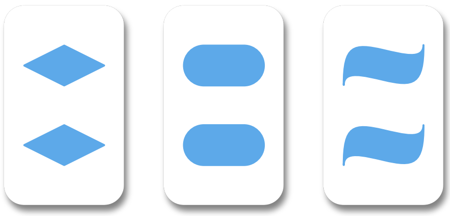

(a) This is a *SET*

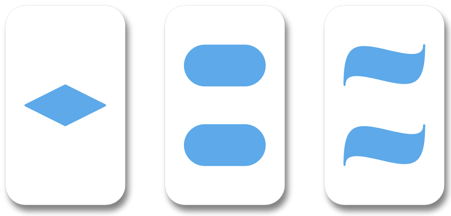

(b) This is NOT a *SET*

Figure 2: Collections of cards

In order to represent the cards in a more mathematical manner we can assign a number to each of the presentations of the attributes thereby embedding each card in $\mathbb{F}_3^4$ (4-dimensional finite field with 3 elements).

So we can represent all 81 cards as the lattice below.

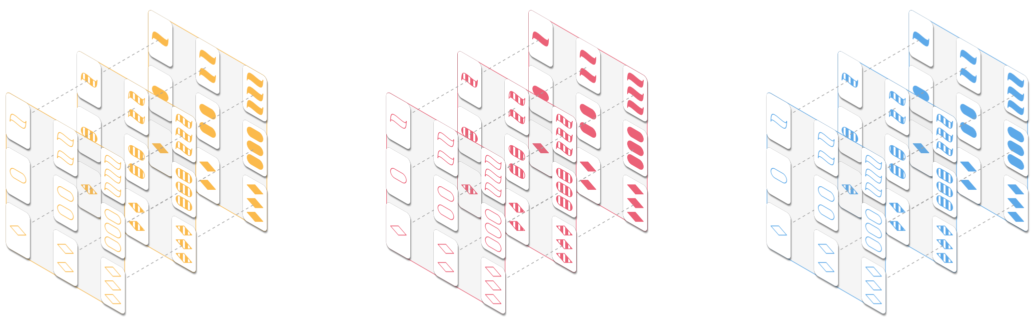

Figure 3: All the cards in *SET* represented in $\mathbb{F}_3^4$

We note that *SET*'s actually correspond to lines in $\mathbb{F}_3^4$! As before we show examples of two *SET*'s and one not

SET in Figure 4

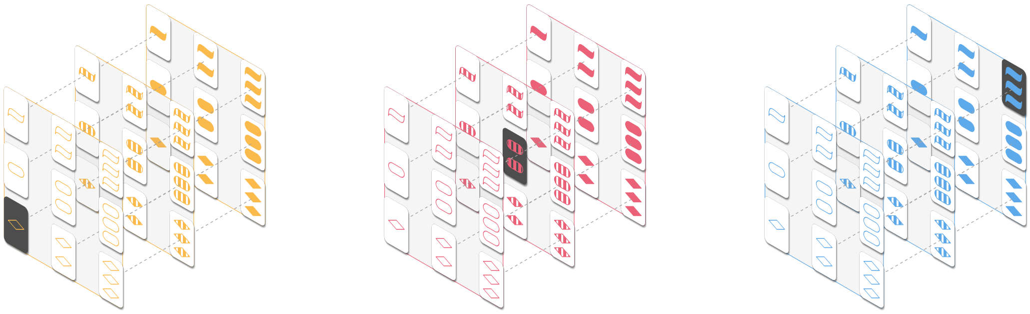

Figure 4: This is a *SET*

The question that we now ask is,

What is the maximum number of cards the dealer can put down without there being a SET in it?

A set of cards that doesn't contain a *SET* is called a **Cap-*SET*** . For this case it turns out that the largest possible Cap-*SET*s has cardinality 20 and an example is shown in Figure 5.

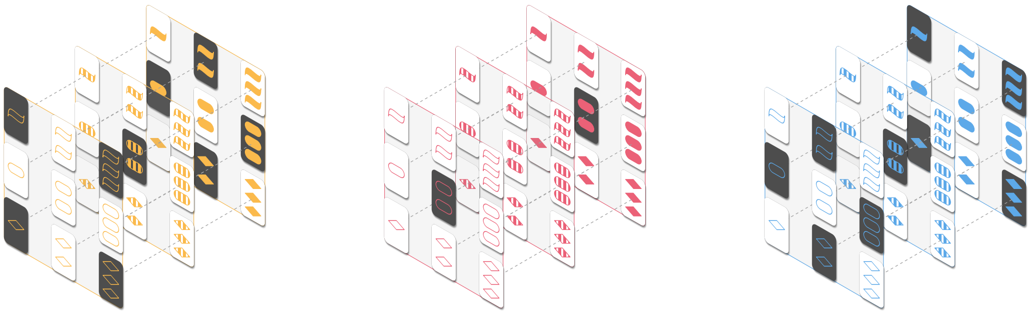

Figure 5: Example of a maximal Cap-Set with 20 cards

2.2 Mathematical Formulations

Formally the Cap-Set problem deals with the case where there are n attributes, each with 3 presentations (as opposed to 4 attributes with 3 presentations in the card game). Furthermore, we note lines in $\mathbb{F}_3^4$ have three distinct points, x, y and z which must satisfy $x + y + z = 0$.

Therefore the classic Cap-Set problem can be stated as below,

Classic Cap-Set Problem

Let $A \subseteq \mathbb{F}_3^n$ such that A contains no lines, ie.

$$x + y + z \neq 0 \quad \forall x, y, z \in A \text{ (distinct)}.$$

How does the maximum size of A grow with respect to n?

The bound of $\mathcal{O}(2.76^n)$ was achieved by Croot-Lev-Pach [1] and Ellenberg-Gijswijt [3] using "conventional" polynomial methods.

In [6] Tao achieved the same bound with a symmetric version of the Classic Cap-Set Problem and we build on the techniques that he laid out in this report.

In fact Ellenberg-Gijswijt proved the result for a stronger problem which we will refer to as the Cap-Set Problem,

Cap-Set Problem

Let $A \subseteq \mathbb{F}_q^n$ such that A doesn't contain any three points $(x, y, z) \in A^3$ that satisfy the following equation,

$$ax + by + cz \neq 0 \quad \forall x, y, z \in A \text{ (distinct)}.$$

Where are $(a, b, c) \in (\mathbb{F}_q \setminus \{0\})^n$ are fixed coefficients that satisfy $a + b + c = 0$. How does the maximum size of A grow with respect to n?

In fact, Tao's approach can be trivially extended to this version of the problem ([6] Remark 3).

2.3 Slice-Rank and Tensors

The concept of slice-rank plays a key role in the work of Tao and is heavily used by us in our results.

Definition 1 (d-Tensor [5])**.** *Let V be an n dimensional linear space over $\mathbb{F}_q$ with $d \geq 1$. Then any multilinear map $T : V^d \to \mathbb{F}_q$ is a d-Tensor and can be represented as,*

$$T\left(x^1, \ldots, x^d\right) = \sum_{i_1, \ldots i_d \in [n]} T_{i_1, \ldots i_d} x^1_{i_1}, \ldots x^d_{i_d}.$$

Where $[n] = \{1, \ldots, n\}$ and $x^i = \left(x^i_1 \ldots, x^i_n\right) \in \mathbb{F}_q^n$ for $i \in [d]$. We can think of this tensor as a d-dimensional array of the coefficients, $T_{i_1, \ldots i_d}$

Definition 2 (Tensor Slice-Rank [5])**.** *Let V denote a finite dimensional vector space over $\mathbb{F}_q$. The slice rank of a d-Tensor $T : V^d \to \mathbb{F}_q$, is denoted as* $\operatorname{srank}(T)$*.*

$\text{srank}(T) = 1$ ***iff*** *there exists a 1-Tensor* $G : V \to \mathbb{F}_q$, *a* $(d-1)$*-Tensor* $H : V \to \mathbb{F}_q$ *and some* $i \in \{1, 2, \ldots d\} = [d]$ *such that*

$$T(v_1, \ldots, v_d) = G(v_i)H\left(v_j | j \in [d] \setminus \{i\}\right).$$

$\text{srank}(T) \leq k$ ***iff*** *there is a sequence of* k *rank one functions on* A^d, $(T_i)_{i\in[k]}$ *such that,*

$$T = \sum_{i=1}^{k} T_i.$$

Tao defines a nearly identical notion of slice-rank for functions $A^d \to \mathbb{F}_q$

Definition 3 (Tao Slice-Rank [6])**.** *The Tao slice-rank of a function* $f : A^d \to \mathbb{F}_q$, *is denoted as* $\text{T-srank}(f)$.

$\text{T-srank}(f) = 1$ ***iff*** *there exists a function* $f : A \to \mathbb{F}_q$, *a function* $h : A^{d-1} \to \mathbb{F}_q$ *and some* $i \in [d]$ *such that*

$$f(x_1, x_2 \ldots x_d) = g(x_i)h\left(x_j | j \in [d] \setminus \{i\}\right).$$

$\text{T-srank}(f) \leq k$ ***iff*** *there are a sequence of* k *rank one d-Tensors,* $(f_i)_{i\in[k]}$ *such that,*

$$f = \sum_{i=1}^{k} f_i.$$

Now with these notions of rank we can relate tensors over V^3 to functions over $\mathbb{A}^3$.

Lemma 1. *For all functions* $F : A^3 \to \mathbb{F}_q$ *of the form,*

$$F(x, y, z) = \sum_{\substack{\alpha,\beta,\gamma \in I \\ I \subseteq A^3}} \delta_\alpha(x)\, \delta_\beta(y)\, \delta_\gamma(z).$$

there exists a 3-Tensor $T_F : V^3 \to \mathbb{F}_q$ *(on* $V = \{A - \mathbb{F}_q\}$*) such that for all* $\alpha, \beta, \gamma \in A^3$,

$$F(\alpha, \beta, \gamma) = T_F(\delta_\alpha, \delta_\beta, \delta_\gamma).$$

Proof. We define the tensor T_F as,

$$T_F(f, g, h) = \sum_{a,b,c \in A} F(a, b, c) f(a) g(b) h(c).$$

Clearly this is a 3-Tensor and for delta functions, we have,

$$\begin{aligned} T_F(\delta_\alpha, \delta_\beta, \delta_\gamma) &= \sum_{a,b,c\in A} F(a, b, c)\delta_\alpha(a)\delta_\beta(b)\delta_\gamma(c) \\ &= F(\alpha, \beta, \gamma) \end{aligned}$$

as required. ∎

Theorem 2 (Tensor slice-rank bounds Tao slice-rank from below)**.** *For a function $F : A^3 \to \mathbb{F}_q$ and its corresponding tensor $T_F : V^3 \to \mathbb{F}_q$ constructed from Lemma 1*

$$\text{T-srank}(F) \geq \text{srank}(T_F).$$

In order to prove this bound we establish some basic properties of the slice-rank in both Tensor and Tao's form.

Lemma 2.

$$\text{T-srank}(F) = 1 \implies \text{srank}(T_F) = 1.$$

Proof. T-srank$(F) = 1$ means that we can express (WLOG),

$$F(\alpha, \beta, \gamma) = G(\alpha)H(\beta, \gamma).$$

From an elementary extension of Lemma 1 we can define tensors $T_G : V \to \mathbb{F}$ and $T_H : (V)^2 \to \mathbb{F}$. Obviously we know, for all $\alpha, \beta, \gamma \in A^3$,

$$T_F(\delta_\alpha, \delta_\beta, \delta_\gamma) = T_G(\delta_\alpha)T_H(\delta_\beta, \delta_\gamma).$$

We must extend this equality to all functions $A \to \mathbb{F}$. We note that the set of delta functions $\{\delta_a | a \in A\}$ form basis vectors for the vector space of functions $A \to \mathbb{F}$. Therefore we can represent any function $f : A \to \mathbb{F}$ as,

$$f(x) = \sum_{\Delta \in A} \mathcal{C}^f_\Delta \delta_\Delta(x).$$

where $\{\mathcal{C}^f_\Delta\}_{\Delta \in A}$ is a sequence of coefficients in $\mathbb{F}_q$ for each delta function. Ie $f(a) = \mathcal{C}^f_a$. Therefore we see,

$$\begin{aligned}
T_F(f, g, h) &= \sum_{a,b,c \in A} F(a, b, c)f(a)g(b)h(c) \\
&= \sum_{a,b,c \in A} F(a, b, c)\mathcal{C}^f_a \mathcal{C}^g_b \mathcal{C}^h_c \\
&= \sum_{a,b,c \in A} G(a)H(b, c)\mathcal{C}^f_a \mathcal{C}^g_b \mathcal{C}^h_c \\
&= \left(\sum_{a \in A} G(a)\mathcal{C}^f_a\right)\left(\sum_{b,c \in A} H(b, c)\mathcal{C}^g_b \mathcal{C}^h_c\right) \\
&= \left(\sum_{a \in A} G(a)f(a)\right)\left(\sum_{b,c \in A} H(b, c)g(b)h(c)\right) \\
&= T_G(f)T_H(g, h)
\end{aligned}$$

Therefore srank$(T_F) = 1$ as required. □

Proof (Tensor slice-rank bounds Tao slice-rank from below). This follows directly from Lemma 2, as a slice-rank one decomposition of F is also a slice-rank one representation of T_F which is an upper bound for the slice rank of T_F. $\square$

2.4 Caro-Wei Theorem and independent sets

Although Slice-rank is a very useful notion of rank it is quite hard do deal with as we see in Tao [6], finding the slice rank of a diagonal matrix requires a complex inductive argument. Therefore, invoke the paper of Lovett [5] to introduce the concept of an **independent set**.

Assume that T is a d-tensor on the space $V = \mathbb{F}_q^n$. By multilinearity of T we have:

$$T(x^1, \ldots, x^d) = \sum_{\alpha \in [n]^d} c_\alpha x^1_{\alpha_1} \ldots x^d_{\alpha_d}, \text{ for } (x^1, \ldots, x^d) \in V^d,$$

where any vector $v \in V$ is represented in the coordinates as $v = (v_1, \ldots, v_n)$, and $\alpha \in [n]^d = \{1, \ldots, n\}^d$ has coordinates $\alpha_1, \ldots, \alpha_d$.

Definition 4 (Independent set)**.** *We define that a set $\mathcal{I} \subset \{1, \ldots, n\}$ is an independent set for T if for any $\alpha = (\alpha_1, \ldots, \alpha_d) \in \mathcal{I}^d$ such that $c_\alpha \neq 0$ implies $\alpha_1 = \ldots = \alpha_d$.*

Theorem 3 (Lovett [5], Theorem 1.7)**.** *For any d-tensor T we have*

$$\text{srank}(T) \gtrsim |\mathcal{I}|,$$

for any independent set $\mathcal{I} \subset \{1, \ldots, n\}$.

In our case we are dealing with $\{0, 1\}$ valued 3-Tensors which can be interpreted as the adjacency matrix for a 3-Uniform Hypergraph. Hypergraphs generalise the concept of edges, as an "edge" can contain any number of vertices (as opposed to 2). A 3-Uniform Hypergraph has exactly 3 vertices for every edge and we can represent it as a $\{0, 1\}$ valued 3-Tensor.

In fact the notion of an independent set also applies in the same way to 3-Hypergraphs, essentially corresponding to the largest subset of vertices such that no two of them are adjacent. The Caro-Wei [2] Theorem provides us with a lower bound on the size of the largest independent set in a graph and furthermore this was extended by Dutta et. al. [2] for k-uniform Hypergraphs.

Theorem 4 (Caro-Wei Generalised)**.** *Let the maximal cardinality of an independent set for a k-uniform Hypergraph $G = (V, E)$ be denoted as $\mathcal{I}(G)$ and let d_v be the degree of a vertex $v \in V$. Then,*

$$\mathcal{I}(G) \gtrsim \sum_{v \in V} \frac{1}{(d_v + 1)^{1/k}}.$$

From the concavity of the function we can easily get a weaker bound with respect to the average degree, d_{ave} as,

Lemma 3 (Caro-Wei Weak)**.**

$$\mathcal{I}(G) \gtrsim \frac{|V|}{(d_{ave}+1)^{1/k}}.$$

2.5 Our Contribution

First we extend this result in the case where we allow some proportion of solutions to the equation $ax+by+cz=0$ with $a+b+c=0$. As an application, we generalise the Ellenberg-Gijswijt Theorem to the multivatiable case.

We define an **Almost Cap-Set** to be a set that allows *some* solutions to the equation $ax+by+cz=0$ for some $a, b, c \in (\mathbb{F}_q \setminus \{0\})^3$ that satisfy $a+b+c=0$. We can formalise this notion by defining an (ϵ, δ)-Cap-Set for some $\epsilon, \delta > 0$ such that there exists some $A' \subset A$ with $|A'| > \delta |A|$ where, for every element $x \in A'$, the number of pairs (y, z) that are solutions to $ax+by+cz=0$ is less than $|A|^\epsilon$.

To aid in our proof, let us define the $A^\epsilon_{\mathbf{a}}$ with respect to some $\epsilon > 0$ and coefficients $\mathbf{a} = (a, b, c) \in (\mathbb{F}_q \setminus 0)$ that satisfy $a+b+c=0$ as the set,

$$A^\epsilon_{\mathbf{a}} = \{x \in A \big| \left|\{(y,z) \in A^2 | ax+by+cz=0\}\right| \leq |A|^\epsilon\}.$$

Theorem 5 (Almost Cap-Sets)**.** *There exist $\epsilon > 0$ and $c_q < q$ such that for any $\delta > 0$, $A \subset \mathbb{F}_q^n$ with $|A| > c_q^n$ and $\mathbf{a} = (a, b, c) \in (\mathbb{F}_q \setminus \{0\})^3$ for sufficiently large n we have*

$$|A^\epsilon_{\mathbf{a}}| \geq (1-\delta)|A|.$$

In other words, there exists $\epsilon > 0$ such that for any $\delta > 0$, the $(\epsilon, \delta)-$cap sets $A \subset \mathbb{F}_q^n$ satisfy that $|A| \leq c_q^n$ for sufficiently large n.

This result shows that there are power-law bounds that also hold for these almost cap-sets. The work of Lovett was crucial in this proof as it allowed us to bound the slice-rank of certain functions by the independent set. Regardless, this allows us to obtain a similar power upper law for the Ellenberg-Gijswijt Cap-Set problem, now extended to the multivariable case.

Theorem 6 (Multivariable Cap-Sets)**.** *Let us have $a_1, a_2, \ldots, a_d \in \mathbb{F}_q \setminus \{0\}$ ($d > 3$) such that $\sum_{i=1}^d a_i = 0$. Then if $|A| > c_q^n$ (for large enough n) implies that there exists distinct $x_1, \ldots x_d \in A$ such that*

$$\sum_{i=i}^{d} a_i x_i = 0.$$

3 Proof of Almost Cap-Set bound (Theorem 5)

We can define the function,

$$\begin{aligned} F : A^3 &\longrightarrow \mathbb{F}_q \\ x, y, z &\longmapsto F(x,y,z) = \delta_{0^n}\left(ax+by+cz\right). \end{aligned}$$

In this proof we will find an upper and lower bound of the slice-rank of this function which will give our desired result. Alternately we can write it as,

$$F(x,y,z) = \sum_{\alpha=(\alpha_1,\alpha_2,\alpha_3)\in A^3} c_\alpha \delta_{\alpha_1}(x)\,\delta_{\alpha_2}(y)\,\delta_{\alpha_3}(z)\,,.$$

where,

$$c_\alpha = \begin{cases} 1 & \text{if } a\alpha_1 + b\alpha_2 + c\alpha_3 = 0 \\ 0 & \text{else} \end{cases}.$$

By Lemma 1 we construct the corresponding 3-Tensor,

$$T_F(f_1,f_2,f_3) = \sum_{\alpha\in A^3} c_\alpha f_1(\alpha_1)f_2(\alpha_2)f_3(\alpha_3)..$$

Since A is an (ϵ,δ)-Cap-Set there exists a subset $A' \subset A$ with $|A'| \geq \delta|A|$ and such that for every $x \in A'$ there are at most $|A|^\epsilon$ pairs $(y,z) \in A^2$ with

$$ax + by + cz = 0.$$

Therefore we deduce,

$$\delta|\{c_\alpha \neq 0 \,|\, \alpha \in (A')^3\}| \leq \delta\,|A|^\epsilon\,|A'| \leq |A'|^{1+\epsilon}.$$

This follows from the fact that for each $x \in A'$ there are at most $|A|^\epsilon$ additional solutions.

Now we can invoke the Caro-Wei lower bound on the independent set from Lemma 3. In this case the Tensor represents a 3-Uniform Hypergraph with $|A'|$ vertices. The degree of each vertex is, $d_x = |\{c_\alpha \neq 0 \,|\, \alpha_1 = x,\, \alpha \in (A')^3\}|$, therefore the average degree is $d_{ave} \leq \delta^{-1}\,|A'|^\epsilon$

Therefore,

$$|\mathcal{I}| \gtrsim \frac{|A'|}{(\delta^{-1}\,|A'|^\epsilon + 1)^{\frac{1}{3}}} \approx |A'|^{1-\epsilon/3}.$$

Therefore by Theorem 3 we get the final bound,

$$\operatorname{srank}(F) \gtrsim |A|^{1-\epsilon/3}.$$

We can shift our focus to obtaining an upper bound on the rank of $F(x,y,z)$ by considering it as a polynomial in the components of x, y, z.

We can expand the delta function as,

$$\delta_{0^n}(ax+by+cz) = \prod_{i=1}^{n}\left(1-(ax_i+by_i+cz_i)^{q-1}\right).$$

Expanding the RHS out we get a polynomial made of the monomials,

$$\prod_{i=1}^{n}(ax_i)^{j_i}(by_i)^{k_i}(cz_i)^{l_i}.$$

With $j_i, k_i, l_i \in \{0, \dots, q-1\}$ and $\sum_{i=1}^{n} j_i + k_i + l_i \leq (q-1)n$.

By the Pigeonhole principle, for each monomial one of the contributions from the x or y or z components to the total degree must be less than $\frac{(q-1)n}{3}$. Ie., $\exists m \in \{i, j, k\}$ s.t.

$$\sum_{i=1}^{n} m_i \leq \frac{(q-1)n}{3}.$$

Therefore, for each monomial we can "extract" the smallest contribution in one of the three variables giving us the following representation in srank 1 functions,

$$\delta_{0^n}(ax+by+cz) = \sum_{\substack{\beta=(\beta_1\dots\beta_n)\in\{0,\dots q-1\}^n \\ \sum\beta_i \leq \frac{(q-1)n}{3}}} f_\beta(x)g_\beta(y,z) + f_\beta(y)g_\beta(z,x) + f_\beta(z)g_\beta(x,y),$$

where $f_\beta(x)$ is the monomial with exponents being the components of β ie.,

$$f_\beta(x) = \prod_{i=1}^{n} x_i^{\beta_i},$$

and g_β represent the contributions dependent on the other two variables which also include the constants.

Let the number of possible monomials (f_β) be N. Therefore we have represented $F(x,y,z)$ in $3N$ slice-rank one functions. Using Cramers theorem from the theory of large deviation reveals that $3N \leq b_q^n$ for some $b_q < q$.

Using the upper and lower bound for the rank of F we have,

$$|A|^{1-\epsilon/3} \lesssim \mathrm{srank}(F) \lesssim b_q^n.$$

Then we choose some $\epsilon > 0$ that satisfies,

$$b_q^{\frac{1}{1-\epsilon/3}} < q,$$

and some c_q, $b_q^{\frac{1}{1-\epsilon/3}} < c_q < q$. This directly implies $|A| < c_q^n$ leading to the contradiction.

4 Proof of the Multivariable Cap-Set Bounds (Theorem 6)

WLOG we can re-arrange the sequence of co-efficients in order to ensure that the partial sums are never 0 ie., $\sum_{i=1}^{k} a_i \neq 0$ for $k = 2, 3, \ldots d-2$.

For the sake of convenience we define the co-efficients from the partial sums as,

$$b_k = \sum_{i=1}^{k} a_i.$$

Consider the following equation in three variables with coefficients b_{d-2}, a_{d-1}, a_d for $t_{d-2}, x_{d-1}, x_d \in A^3$

$$b_{d-2}t_{d-2} + a_{d-1}x_{d-1} + a_d x_d = 0.$$

By Theorem 5, for some $\delta > 0$ and sufficiently large n there exists $\epsilon > 0$ and a set $A^{\epsilon}_{b_{d-2},a_{d-1},a_d} = A_{d-2} \subset A$ with $|A_{d-2}| \geq (1-\delta)\,|A|$. Recall, for any $t_{d-2} \in A_{d-2}$ there are at least $|A|^{\epsilon}$ pairs $(x_{d-1}, x_d) \in A^2$ satisfying the above equation.

Then we can apply Theorem 5 recursively on the equations indexed by $k = \{d-3,\, d-2, \ldots 3,\, 2\}$, and

$$\begin{aligned} b_{d-3}t_{d-3} + a_{d-2}x_{d-2} &= b_{d-2}t_{d-2} & (k = d-3) \\ &\vdots & \vdots \\ b_k t_k + a_{k+1}x_{k+1} &= b_{k+1}t_{k+1} & (k) \\ &\vdots & \vdots \\ b_2 t_2 + a_3 x_3 &= b_3 t_3 & (k = 2) \end{aligned}$$

and finally,

$$a_1 x_1 + a_2 x_2 = b_2 t_2..$$

For example in the next step we apply Theorem 5 on the set A_{d-2} with the equation corresponding to $d-3$ giving us a set $A_{d-3} \subset A_{d-2}$ such that for every $t_{d-3} \in A_{d-3}$ there are at least $|A_{d-2}|^{\epsilon}$ pairs $(x_{d-2}, t_{d-2}) \in A^2_{d-2}$ that satisfy $b_{d-3}t_{d-3} + a_{d-2}x_{d-2} = b_{d-2}t_{d-2}$.

Using such a method we can construct a chain of nested sets,

$$A \supset A_{d-2} \supset A_{d-3} \supset \ldots \supset A_2,$$

that satisfy the following properties for $k \in \{d-2,\, d-3,\, \ldots 2\}$.

- $|A_{k-1}| \geq (1-\delta)\,|A_k|$
- For any $t_k \in A_k$ there exists at least $|A_{k+1}|^{\epsilon}$ pairs $(x_{k+1}, t_{k+1}) \in A^2_{k+1}$ that satisfy $b_k t_k + a_{k+1}x_{k+1} = b_{k+1}t_{k+1}$.

Now working our way down the chain we can construct a solution to the original multivariable equation using the ($|A_{k+1}|^\epsilon$) abundance of solutions to our advantage.

Take a pair of distinct $(x_1, x_2) \in A_2^2 \subset A^2$ satisfying that $a_1x_1 + a_2x_2 = b_2t_2$ for $t_2 \in A_2$. Then there exist at least $|A_3|^\epsilon$ pairs $(x_3, t_3) \in A_3^2 \subset A^2$ satisfying $b_2t_2 + a_3x_3 = b_3t_3$ for t_2 that we already chosen. Find (x_3, t_3) among these solutions such that $x_3 \notin \{x_1, x_2\}$. Assume that we already constructed distinct $\{x_1, \ldots, x_k\} \in A^k$ satisfying that $a_1x_1 + \ldots + a_kx_k = b_kt_k$, for some $t_k \in A_k$. Since there exist at least $|A_{k+1}|^\epsilon$ pairs $(x_{k+1}, t_{k+1}) \in A_{k+1}^2$ satisfying

$$b_kt_k + a_{k+1}x_{k+1} = b_{k+1}t_{k+1},$$

we can choose one of the solutions $(x_{k+1}, t_{k+1}) \in A_{k+1}^2$ satisfying that $x_{k+1} \notin \{x_1, \ldots, x_k\}$. Notice that there exists $t_{k+1} \in A_{k+1}$ such that the sequence $(x_1, \ldots, x_{k+1}) \in A^{k+1}$ satisfies

$$a_1x_1 + \ldots + a_{k+1}x_{k+1} = b_{k+1}t_{k+1}.$$

We continue this process till we reach distinct $\{x_1, \ldots, x_{d-2}\} \in A$ satisfying

$$a_1x_1 + \ldots + a_{d-2}x_{d-2} = b_{d-2}t_{d-2}$$

for some $t_{d-2} \in A_{d-2}$. Since there are at least $|A|^\epsilon$ pairs $(x_{d-1}, x_d) \in A^2$ satisfying

$$b_{d-2}t_{d-2} + a_{d-1}x_{d-1} + a_dx_d = 0,$$

we can choose the solution $(x_{d-1}, x_d) \in A^2$ such that $x_d \neq x_{d-1}$ and $x_{d-1}, x_d \notin \{x_1, \ldots, x_{d-2}\}$. This finishes the proof of Theorem 6.

5 Acknowledgements

I would like to thank the University of Sydney and the AMSI VRS program for allowing me to take this project.

Most of all I would like to thank my supervisor Assoc. Prof Fish for his invaluable mentorship as well as being a key driving force in exploring this problem.

References

[1] Ernie Croot, Vsevolod Lev, and Péter Pach. "Progression-free sets in $\mathbb{Z}_4^n$ are exponentially small". In: *Annals of Mathematics* 185.1 (2017), pp. 331–337. DOI: 10.4007/annals.2017.185.1.7.

[2] Kunal Dutta, Dhruv Mubayi, and C. R. Subramanian. "New Lower Bounds for the Independence Number of Sparse Graphs and Hypergraphs". In: *SIAM Journal on Discrete Mathematics* 26.3 (2012), pp. 1134–1147. DOI: 10.1137/110839023.

[3] Jordan Ellenberg and Dion Gijswijt. "On large subsets of $\mathbb{F}_q^n$ with no three-term arithmetic progression". In: *Annals of Mathematics* 185.1 (2017), pp. 339–343. DOI: 10.4007/annals.2017.185.1.8.

[4] Alexander Fish and Dibyendu Roy. *On almost Cap sets in three variables and the multivariable Cap set problem*. 2021. arXiv: 2102.11496 [math.NT].

[5] Shachar Lovett. "The Analytic Rank of Tensors and Its Applications". In: *Discrete Analysis* 7 (2019).

[6] Terrence Tao. *A symmetric formulation of the Croot-Lev-Pach-Ellenberg-Gijswijt capset bound*. http://terrytao.wordpre.com/2016/05/18/a. June 2016.