Three-wise restrictions and the slice rank

Tibor Szabó

Extremal Combinatorics, FU Berlin, WiSe 2017–18

In applications of the linear algebra method we have seen so far the restrictions on the set family included *single* sets and/or *pairs* of sets. For example, the size of sets and/or the size of pairwise intersection of sets from the family was restricted in some way. This restriction could then be encoded into information about the dot product of charateristic vectors, which then later could be utilzed to bound the dimension of an appropriate space of vectors or functions associated with the members of the family.

These applications all fundamentally depended on the dot product, translating the combinatorial restrictions to linear algebra, being a *pairwise* operation. Until very recently there was no known method that could effectively use linear algebra to problems involving restrictions involving *more than two* sets from the family. In 2016 there was a breakthrough in this direction. In a matter of just a couple of weeks a technique was developed to study exactly the problems of this sort. The technique was used to shatter the best known bounds in several well-studied problems of extremal combinatorics. Even more excitingly, everything that happened can be fully explained in a Master course. In this section we introduce two of the problems and apply the method of slice rank of tensors to solve them.

1 Sunflower-free families

In the section on the classics we discussed the Erdős-Rado Sunflower Conjecture about the maximum number of sets a k-uniform family can contain without containing a sunflower with ℓ petals (an ℓ-sunflower). This conjecture is *very* open, even in the case $\ell = 3$, i.e. when we are looking for three sets in the family such that all three pairwise intersections are the same set.

An important feature of this problem and also the connected theorem of Erdős and Rado is the independence of the conjectured bound on the number of vertices. This presents a substantial difficulty for any approach. In order to get grip on this notoriously difficult problem, Erdős and Szemerédi [4] suggested to obtain a non-trivial upper bound involving the number n of vertices instead of the uniformity. The best they could come up with was that the maximum size of a 3-sunflower-free family on n vertices was at most $2^n/e^{c\sqrt{n}}$. This is hardly better than the trivial bound of 2^n, while they believed that an exponential factor improvement should also be true.
Conjecture 1.1. *For every $\ell \in \mathbb{N}_+$ there exists a constant $c_\ell < 2$ such that every family $\mathcal{F} \subseteq 2^{[n]}$ of size at least c_ℓ^n contains an ℓ-sunflower.*

Erdős and Szemerédi proved that the Erdős-Rado Conjecture implies their conjecture.

The k-uniform construction we gave for the Erdős-Rado Conjecture on 3-sunflowers had $n = 2k$ vertices and contained 2^k sets. This gives a lower bound of $\sqrt{2}$ for the constant c_3 in the Erdős-Szemerédi Conjecture.

Similarly to the Erdős-Rado Sunflower Conjecture, the Erdős-Szemerédi Sunflower Conjecture was already wide open for $\ell = 3$. Recently the conjecture was proved for $\ell = 3$ and next we show the proof. The conjecture is still wide open for $\ell > 3$.

Theorem 1.2 ([6]). *For the constant $c = \frac{3}{\sqrt[3]{4}} < 1.89$ we have that every family $\mathcal{F} \subseteq 2^{[n]}$ of size at least c^n contains a 3-sunflower.*

Before going into the actual proof, we motivate the general proof strategy we use in this section. In the intersection theorems of the previous section we encountered restrictions on every single set and/or every *pair* of sets from the family. In some of these proofs, e.g. for Fischer's Inequality and for the Mod-p-Town Theorem, we converted this information into a $|\mathcal{F}| \times |\mathcal{F}|$ diagonal *matrix* with non-zero diagonal entries. Then, from the way this matrix was constructed, we have shown that its rank, that is equal to its order $|\mathcal{F}|$, can be bounded by n.

In this argument it was crucial that the restrictions concerned *pairs* of members of the family, so we could build a matrix. In the 3-sunflower problem at hand, we have a restriction about every *triple* of sets of the family: they cannot form a sunflower. We will translate this information into a $|\mathcal{F}| \times |\mathcal{F}| \times |\mathcal{F}|$ matrix, aka a 3-*tensor*, which only has non-zero entries in its main diagonal. We appropriately generalize the concept of rank of a matrix to 3-tensors. This new rank concept, called *slice rank*, is always at most the usual tensor rank[1], and in the special case of a diagonal tensor with non-zero diagonal entries it is equal to the order $|\mathcal{F}|$ of the tensor. Furthermore we will be able to effectively bound the slice rank of the particular 3-tensors and hence obtain an upper bound on $|\mathcal{F}|$.

Definition 1.3. *Let X be a finite set and $\mathbb{F}$ be a field. A function $M : X^k \to \mathbb{F}$ is called a* k-tensor*. We call a k-tensor* diagonal *if $M(x_1, \dots, x_k) \neq 0$ implies $x_1 = \dots = x_k$.*

A 1-tensor is just a vector, and a 2-tensor is a matrix.

Definition 1.4. *A k-tensor $S : X^k \to \mathbb{F}$ is called a* slice *if there exist a 1-tensor $f : X \to \mathbb{F}$ and a $(k-1)$-tensor $g : X^{k-1} \to \mathbb{F}$, such that for some $i \in |X|$ we have $S(x_1, \dots, x_k) = f(x_i)g(x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_k)$ for every $x_1, \dots, x_k \in X$.*

The slice rank *of a k-tensor $M : X^k \to \mathbb{F}$ is the smallest integer r such that M can be written as a sum of r slices. The slice rank of M is denoted by $srk(M)$.*

When $k = 2$, a slice S is just a matrix of rank at most 1. In a matrix S of rank at most 1 each row is a constant multiple of some vector v^T, so S has the form uv^T, where u is the column vector of the multipliers. That is, S is the product of u and v, when they are viewed as 1-tensors on the first and second coordinate of S, respectively. As the rank of a matrix M is the smallest number r such that M is the sum of r matrices of rank 1, the slice rank of a 2-tensor is just the usual matrix rank.

Observe that every k-tensor can be written as $M(x_1, \dots, x_k) = \sum_{x \in X} \delta_{x_1 x} M(x, x_2, x_3, \dots, x_k)$, where δ_{ab} is the Kronecker delta function. Here the terms are slices, so the slice rank is never more than $|X|$.

In the definition of the usual tensor rank the k-tensor has to be written as the sum of terms of the form $f_1(x_1) \cdot \dots \cdot f_k(x_k)$, i.e. product of k 1-tensors. Consequently, the slice rank is always less than or equal to the usual tensor rank.

The following lemma is the generalization of the trivial fact that a diagonal matrix with non-zero diagonal entries has rank $|X|$.

Lemma 1.5. *Let $M : X^k \to \mathbb{F}$ be a diagonal k-tensor with non-zero diagonal entries. Then*

$$srk(M) = |X|.$$

We give the proof after our two application.

Proof. Let $\mathcal{S} \subseteq 2^{[n]}$ be a 3-sunflower-free family. To create our 3-tensor we observe that three distinct sets form a sunflower if and only if NO element of $[n]$ occurs in exactly two of them.

[1]The tensor rank of a function $M : X^k \to \mathbb{F}$, where X is a finite set and $\mathbb{F}$ is a field, is equal to the minimum non-negative integer r for which we can write $M(x_1, \dots, x_k) = \sum_{i=1}^{r} f_1(x_1) f_2(x_2) \cdots f_k(x_k)$ where each f_i is a function from X to $\mathbb{F}$.

Formulated in terms of characteristic vectors this says that for any three distinct elements A, B, C of $\mathcal{S}$ there exists an $i \in [n]$ such that $(v_A + v_B + v_C)_i = 2$. This motivates us to define the function

$$M(x, y, z) := \prod_{i=1}^{n} \left(2 - (x + y + z)_i\right),$$

where $x, y, z \in \mathbb{F}^n$. Then $M(v_A, v_B, v_C) = 0$ for every three distinct members A, B, C of $\mathcal{S}$. To make M a diagonal 3-tensor on $\mathcal{S}$ we also need to take care of the substitutions where two of the three sets are equal. For any $A \neq B = C$ we have that $v_A + v_B + v_C$ has a coordinate 2 if and only if B has an element that is not in A, i.e., $B \not\subseteq A$. For this reason we partition $\mathcal{S}$ to antichains, say classifying the members according to their size: let $\mathcal{S}_j = \{S \in \mathcal{S} : |S| = j\}$. Then M is a diagonal tensor on each $\mathcal{S}_j$. Since there are only $n + 1$ different values of j and we are seeking an exponential upper bound, it won't make much of a difference if we handle each $\mathcal{S}_j$ separately.

Finally note that the diagonal entries $M(v_A, v_A, v_A) = (-1)^{|A|} 2^{n-|A|}$ are all non-zero if we choose to work over a field whose characteristic is not 2, say $\mathbb{F} = \mathbb{R}$. Applying Lemma 1.5 to the 3-tensor M on $\mathcal{S}_j$, we obtain that $|\mathcal{S}_j| = srk(M)$.

All that is left is to bound the slice rank of M, which we do by finding a particular decomposition of M into slices. M is a polynomial of total degree n in the $3n$ variables $x_1, \ldots, x_n, y_1, \ldots, y_n, z_1, \ldots, z_n$. Expanding the parenthesis we obtain

$$M(x, y, z) = \sum_{I \sqcup J \sqcup K \sqcup L = [n]} 2^{|L|} x_I y_J z_K,$$

where we use the notation $w_I = \prod_{i \in I} w_i$ for the product of some coordinates of vector w, and $\sqcup$ denotes disjoint union. We classify the terms according to the smallest of the sizes $|I|, |J|, |K|$, breaking ties arbitrarily. Since I, J and K are disjoint, this minimum size is definitely at most $\frac{n}{3}$. Hence we can write

$$M(x, y, z) = \sum_{\substack{I \subseteq [n] \\ |I| \leq n/3}} x_I f_I(y, z) + \sum_{\substack{J \subseteq [n] \\ |J| \leq n/3}} y_J g_J(x, z) + \sum_{\substack{K \subseteq [n] \\ |K| \leq n/3}} z_K h_K(x, y)$$

for some 2-tensors f_I, g_J, h_K, by collecting all the corresponding terms. Each of the terms in this sum are slices, so the slice rank of M is at most $3 \sum_{i=0}^{n/3} \binom{n}{i}$. Therefore,

$$|\mathcal{S}| = \sum_{j=0}^{n} |\mathcal{S}_j| \leq 3(n + 1) \sum_{i=0}^{n/3} \binom{n}{i} = 2^{(H(1/3)+o(1))n} < 1.89^n.$$

□

2 The Capset Problem

In the visual perception game SET [2] one plays with cards depicting objects having four features: shape, color, number, and shading. Each feature has three possible values: the shape can be oval, squiggles, or diamonds, the color can be red, purple or green, the number can be one, two or three, and the shading can be solid, striped or outlined. Each of the combinations of these occurs on exactly one card, which makes up the deck of $3^4 = 81$ cards. The game starts by the players putting out 12 cards face up and then staring at it, looking for a configuration of three cards called SET. A *SET* consists of three cards in which each of the cards' features, looked at one-by-one, are the same on each card, or, are different on each card.

[2] `https://www.setgame.com/set/puzzle`

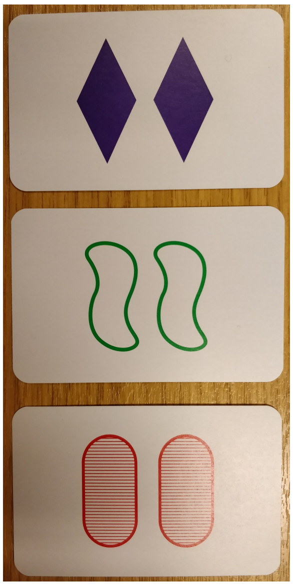

Figure 1: A SET in which exactly one feature is the same, the number

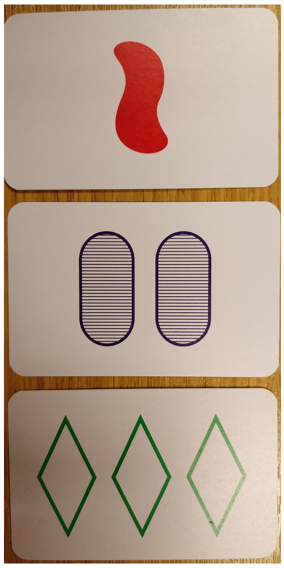

Figure 2: A SET in which all features are different

Whoever spots a SET first and lets this known to the opponents with a emphatic shout “SET!”, can take the three cards. Three new cards are put out and the staring for a new SET continues. Whoever has the most cards at the end wins. Sometimes it happens that after some staring the players realize that the particular 12 cards in front of them does not contain any SET. In that case they add three more cards to the table and try to find a SET among the 15 cards. Very rarely, but it could still happen that no SET is found and three more cards need to be added. Apparently even then it could happen that the 18 cards does not contain any SET (though I have never encountered such a configurations during my playing career). What is then the minimum number of cards where no matter what there is a SET? This number is 21 and it was proved much before the invention of the game by the Italian mathematician Giuseppe Pellegrino [5]. The motivation for the study of these objects for mathematicians came from finite geometry and coding theory. Let us see now finite geometry comes into the picture.

We can encode the cards of SET as vectors of length four, where each coordinate corresponds to a feature and can take three different values. This is exactly the vector space $\mathbb{F}_3^4$. What does a SET corresponds to in this vector space? The elements $a, b, c \in \mathbb{F}_3$ sum up to 0 if they are pairwise distinct and so do they if they are all the same. If $a \neq b = c$ however, then $a + b + c = a + 2b = a - b \neq 0$ in $\mathbb{F}_3$. Hence three distinct cards corresponding to the vectors $x, y, z \in \mathbb{F}_3^4$ form a SET if and only $x + y + z = 0$

Rewriting this further $x + y + z = x + y - 2z$ over $\mathbb{F}_3$, so a SET-free set of vectors is just a 3-AP-free set of vectors. We already studied this problem extensively within the positive integers, but the concept of a 3-AP makes sense in any abelian group . Here we will be looking for the largest number $r_3\left(\mathbb{F}_3^d\right)$ such that there is a 3-AP-free set of that size in $\mathbb{F}_3^d$. In terms of the card game SET, if cards had d features instead of 4, then $r_3\left(\mathbb{F}_3^d\right) + 1$ would be the smallest number of cards you would need to put out face up on the table to guarantee to find three of them forming a SET. The determination of $r_3\left(\mathbb{F}_3^d\right)$ is known as the *capset problem*, where the “cap” in the name originates in finite geometry.[3]

[3] The name “cap” in the “capset problem” comes from the fact that 3-APs are also equivalent to lines in the corresponding affine space over $\mathbb{F}_3$. Analogously to the real d-space, a *line* in $\mathbb{F}^d$, over an arbitrary field $\mathbb{F}$, is an affine subspace of dimension 1. Over $\mathbb{F}_3$ this means the three-element subsets of the form $\{\alpha m + t : \alpha \in \mathbb{F}_3\}$, where $m \in \mathbb{F}_3^d \setminus \{0\}$ could be considered the “slope” and $t \in \mathbb{F}_3^d$ is the translation vector. Clearly, the sum

The greedy construction of Erdős and Turán for $r_3(n)$, containing the integers without any digit 2 in their ternary expansion, can also be interpreted in $\mathbb{F}_3^d$: the set $\{0,1\}^d$ is a 3-AP-free set, so $2^d \le r_3(\mathbb{F}_3^d)$. From above, obviously, $r_3(\mathbb{F}_3^d) \le 3^d$. The small values were also determined up to some point: for example, we have $r_3(\mathbb{F}_3^2) = 4$, $r_3(\mathbb{F}_3^3) = 9$ and $r_3(\mathbb{F}_3^4) = 20$.

Besides its relevance in finite geometry, where it is mostly studied in fixed dimension over the field $\mathbb{F}_q$ with a growing q, the problem received significant attention from combinatorial number theorists due to its connection to the Density Hales-Jewett Theorem and to 3-AP-free sets in the integers in particular. From the various different proofs of $r_3(n) = o(n)$, by far the best upper bound was delivered by the analytic number theory approach of Roth from the 50's, using Fourier-techniques. Since then this was improved, transformed and developed several times to reach the current record of $n/\log^{1-o(1)} n$ by Sanders (2011). Going beyond the exponent 1 of the logarithm in the denominator is considered a conceptually important benchmark for the $r_3(n)$ problem, with further implications.

The $r_3\left(\mathbb{F}_3^d\right)$-problem was always considered a plausible terrain to develop and test ideas and methods for the $r_3(n)$-problem. Still, the main question for a long time was just to decide whether $r_3\left(\mathbb{F}_3^d\right)$ behaves similarly to $r_3(n)$ in the sense that it is larger than any power of the size 3^d of the base set. We know that $r_3(n)$ is larger than $n^{1-\epsilon}$ for any $\epsilon > 0$, is $r_3\left(\mathbb{F}_3^d\right) > 3^{(1-\epsilon)d}$ for any $\epsilon > 0$? Progress on this problem was really slow. In 1982 Brown and Buhler showed that $r_3\left(\mathbb{F}_3^d\right) = o(3^d)$. In 1995 Meshulam used Fourier analysis on $\mathbb{F}_3^d$ to get an explicit upper bound of $\frac{3^d}{d}$. Note that this is comparable to Sanders' upper bound on $r_3(n)$, where the density is the logarithm of the base set. In 2012 Bateman and Katz [2] managed to beyond Meshulam's bound by a tiny ϵ power of the denominator. Their argument was long and technical and was born in hope of inspiring similar progress for $r_3(n)$.

While such small improvement is still outstanding for $r_3(n)$, a much greater one was achieved in 2016 for $r_3(\mathbb{F}_3^d)$. To much surprise the argument used linear algebra instead of Fourier analysis and was simple enough to be presented in a Masters course. Ellenberg and Gijswijt improved the upper bound on $r_3(\mathbb{F}_3^d)$ by an exponential factor [3], hence establishing that the behaviour of $r_3(\mathbb{F}_3^d)$ and $r_3(n)$ are significantly different. They were building on a breakthrough, achieved by Croot, Lev, and Pach just a few days earlier, for the analogous 3-AP-free set problem in $\left(\mathbb{Z}/4\mathbb{Z})^d\right)$.

Theorem 2.1 (Ellenberg-Gijswijt, 2016). *For large d*

$$r_3\left(\mathbb{F}_3^d\right) < 2.76^d.$$

Proof. Let $S \subseteq \mathbb{F}_3^d$ be a 3-AP-free set. We have seen above that this equivalent to that for every distinct $a, b, c \in S$, the sum $a + b + c \neq 0$. This motivates us to define the following function

$$M(x, y, z) = \prod_{i=1}^{d} (x_i + y_i + z_i - 1)(x_i + y_i + z_i - 2).$$

By the above $M(a, b, c) = 0$ for any distinct $a, b, c \in S$. The same holds for any $a \neq b = c$ as well, since $\alpha + 2\beta \neq 0$ in $\mathbb{F}_3$ for any $\alpha \neq \beta$. So M is a diagonal 3-tensor on S. Since $M(a, a, a) = 2^d \neq 0$ for any $a \in \mathbb{F}_3^d$, Lemma 1.5 applies, and we have $|S| = srk(M)$.

To estimate the slice rank of M we expand M as a polynomial in the $3d$ variables $x_1, \ldots, x_d, y_1, \ldots, y_d, z_1, \ldots, z_d$.

$t + (m + t) + (2m + t) = 3m + 3t$ of the members of a line is always 0 over $\mathbb{F}_3$. In the other direction if $x + y + z = 0$, then the set $\{x, y, z\}$ is a line with slope $m = y - x$ and translation vector $t = x$.

A set of points is referred to a "cap" in finite geometry if it does not contain three collinear points. A typical example of a cap in in $\mathbb{F}_q^3$ comes from the solution set of the equation of an elliptic quadric $z = f(x, y)$ (where f is an irreducible degree 2 polynomial over $\mathbb{F}_q$), which has size q^2. In $\mathbb{R}^3$ some of these solution sets sort of look like a cap, hence the name capset. In dimension 3 over $\mathbb{F}_q$ the elliptic quadric in fact gives the largest possible cap. When $q = 3$, then 3-AP consist of exactly three collinear points, so the name capset problem for $r_3(\mathbb{F}_3^d)$. Note that while a cap is always a 3-AP free set, for $q > 3$ the converse is not true.

The total degree is $2d$ and the degree in each variable is 2, so we can write

$$M(x,y,z) = \sum_{\alpha,\beta,\gamma\in\{0,1,2\}^d} c_{\alpha,\beta,\gamma}\left(\prod_{i=1}^{d} x_i^{\alpha_i}\right)\left(\prod_{i=1}^{d} y_i^{\beta_i}\right)\left(\prod_{i=1}^{d} z_i^{\gamma_i}\right).$$

We classify the terms according to which of x, y, or z has the smallest total degree. Since the overall total degree is $2d$, at least one of $\sum\alpha_i$, $\sum\beta_i$, or $\sum\gamma_i$ is at most $2d/3$, and then

$$M(x,y,z) = \sum_{\substack{\alpha\in\{0,1,2\}^d\\ \sum\alpha_i\le 2d/3}}\left(\prod_{i=1}^{d} x_i^{\alpha_i}\right) f_\alpha(y,z) + \sum_{\substack{\beta\in\{0,1,2\}^d\\ \sum\beta_i\le 2d/3}}\left(\prod_{i=1}^{d} y_i^{\beta_i}\right) g_\beta(x,z) + \sum_{\substack{\gamma\in\{0,1,2\}^d\\ \sum\gamma_i\le 2d/3}}\left(\prod_{i=1}^{d} z_i^{\gamma_i}\right) h_\gamma(x,y),$$

for some 2-tensors $f_\alpha, g_\beta, h_\gamma$. All these terms are slices, hence the slice rank is bounded by 3 times the number of ways to select a vector $\alpha\in\{0,1,2\}^d$ such that the sum of its coordinates is at most $\frac{2}{3}d$. This number is

$$3\cdot\sum_{\substack{a+b+c=d\\ b+2c\le 2d/3}}\frac{d!}{a!b!c!},$$

where a, b, and c represent the number of $0, 1$, and 2 coordinates of α, respectively.

To estimate this multinomial coefficient consider the expression

$$(1+x+x^2)^d = \sum_{\substack{a,d,c\in\mathbb{N}_0\\ a+b+c=d}}\frac{d!}{a!b!c!}x^{b+2c},$$

which is true for every real x by the Multinomial Theorem. To make the terms of our interest, i.e. $b+2c\le\frac{2d}{3}$, the dominating ones, we first divide through by $x^{2d/3}$ and then estimate when $0<x<1$:

$$(x^{-2/3}+x^{1/3}+x^{4/3})^d > \sum_{\substack{a+b+c=d\\ b+2c\le 2d/3}}\frac{d!}{a!b!c!}x^{b+2c-\frac{2}{3}d} > \sum_{\substack{a+b+c=d\\ b+2c\le 2d/3}}\frac{d!}{a!b!c!}.$$

Here in the first estimate we used that $x>0$ and in the second one that $x<1$. To obtain an upper as strong as possible, we use the estimate for the particular x_0, $0<x_0<1$, that minimizes the function $f(x)=x^{-2/3}+x^{1/3}+x^{4/3}$ on this interval. Basic calculus shows that the minimum is taken at $x=\frac{\sqrt{33}-1}{8}$ and its value is roughly 2.7551 which is less than 2.76. This gives us the bound of 2.76^d. □

3 Proof of the Slice Rank Lemma

In this section we prove that the slice rank of a diagonal tensor with non-zero diagonal entries is equal to its order.

Proof of Lemma 1.5. To simplify notation we only present the proof for $k=3$, which is the case we used in both of our applications anyway. The proof of the general case is similar.

Let $M: X^3\to\mathbb{F}$ be an arbitrary diagonal 3-tensor with non-zero diagonal entries. The upper bound of $|X|$ holds for any 3-tensor and was proved in a remark right after the definition.

To show the lower bound, we assume that there is a decomposition of M into the sum of $\gamma = srk(M) < |X|$ slices and will arrive at a contradiction. So suppose

$$M(x,y,z) = \sum_{i=1}^{\alpha} f_i(x)G_i(y,z) + \sum_{i=\alpha+1}^{\beta} f_i(y)G_i(x,z) + \sum_{i=\beta+1}^{\gamma} f_i(z)G_i(x,y),$$

where $0 \leq \alpha \leq \beta \leq \gamma < |X|$ are integers, the f_is are 1-tensors, and the G_is are 2-tensors on X. Consider the subspace V orthogonal to the vectors $f_1, \dots f_\alpha$, that is

$$V := \langle f_1, \dots, f_\alpha \rangle^\perp = \left\{ v : X \to \mathbb{F} : \sum_{x \in X} v(x) f_i(x) = 0 \text{ for all } i = 1, \dots, \alpha \right\}.$$

Let $v \in V$ be a vector with support $S_v = \{x \in X : v(x) \neq 0\}$ that is as large as possible. Then

$$|S_v| \geq \dim V = |X| - \dim \langle f_1, \dots, f_\alpha \rangle \geq |X| - \alpha.$$

Indeed, if $|S_v| < \dim V$ then the dimension $|X| - |S_v|$ of the subspace $\{w \in \mathbb{F}^X : w(x) = 0 \text{ for all } x \in S_v\}$ is strictly larger than $|X| - \dim V$, so it intersects V non-trivially. Hence there is a non-zero vector $w \in V$ such that $w(x) = 0$ for every $x \in S_v$ and then the support of $v + w \in V$ is larger than $|S_v|$, contradicting the maximality of S_v.

We will arrive at our final contradiction through examining the rank of the 2-tensor Q, defined by $Q(y, z) := \sum_{x \in X} v(x) M(x, y, z)$, on the support S_v of v. Substituting the definition of M and exchanging the sums, we can write Q as a sum of $\gamma - \alpha$ slices:

$$\begin{aligned} Q(y, z) &= \sum_{i=1}^{\alpha} G_i(y, z) \sum_{x \in X} f_i(x) v(x) + \sum_{i=\alpha+1}^{\beta} f_i(y) \sum_{x \in X} v(x) G_i(x, z) + \sum_{i=\beta+1}^{\gamma} f_i(z) \sum_{x \in X} v(x) G_i(x, y) \\ &= \sum_{i=\alpha+1}^{\beta} f_i(y) h_i(z) + \sum_{i=\beta+1}^{\gamma} f_i(z) h_i(y), \end{aligned}$$

for some 1-tensors h_i. In the second equality we used that v is orthogonal to every f_i, $1 \leq i \leq \alpha$. We have thus shown that the (slice) rank of Q is at most $\gamma - \alpha$.

On the other hand direct substitution contradicts this. We have $Q(y, z) = \sum_{x \in X} v(x) M(x, y, z) = 0$ whenever $y \neq z$ since M is diagonal. For $y \in S_v$ we have $Q(y, y) = \sum_{x \in X} v(x) M(x, y, y) = v(y) M(y, y, y) \neq 0$. Consequently Q is a diagonal matrix with non-zero entries in the diagonal so its rank is its order $|S_v| \geq |X| - \alpha > \gamma - \alpha$, a contradiction.

□

Remark. Alon, Shpilka and Umans [1] showed in 2011 that the Erdős-Szemerédi Conjecture follows from the solution of the Capset Problem. It is this implication through which the Sumflower problem was resolved when Ellenberg and Giiswijt solved the Capset Problem in 2016. The concept of slice rank was distilled from the paper of Croot, Lev, and Pach by Tao in his blog [7]. The proof of we presented here is based on the paper of Naslund and Sawin [6] who applied the slice rank lemma straight to the Sunflower Problem and obtained a better numerical value for the base of the exponent than one that follows through the Alon, Shpilka and Uman simplication. Alon, Spilka, and Umans also established the connection of these problems to the computational complexity of matrix multiplication.

References

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

[2] M. Bateman and N. H. Katz. New bounds on cap sets. JAMS 25 (2012), 585–613.

[3] J. Ellenberg and D. Gijswijt. Ann. Math. 185 (2017), 339–343.

[4] P. Erdős and E. Szemerédi. Combinatorial properties of systems of sets. J. Combin. Theory Ser. A 24 (1978), 308–313.

[5] G. Pellegrino. Sul massimo ordine delle calotte in $S_{4,3}$, Matematiche (Catania) 25 (1970), 1–9.

[6] E. Naslund and W. F. Sawin. Upper bounds for sunflower-free sets. Forum Math. Sigma 5 (2017).

[7] Terence Tao. A symmetric formulation of the Croot-Lev-Pach-Ellenberg-Gijswijt capset bound, 2016. URL: `https://terrytao.wordpress.com/2016/05/18/a-symmetric-formulation-of-the-croot-lev-pach-ellenberg-gijswijt-capset-bound/`