Slice rank of tensors and its applications

Wenjie Fang, LIP, ENS de Lyon
Work of Terence Tao and William Sawin

One Day Meeting in Discrete Structures
April 14 2017, ENS de Lyon

Two problems in extremal combinatorics

Sunflower-free set problem

Let U be a finite set with $n = |U|$. Three subsets A, B, C of U form a **sunflower** if $A \cap B = B \cap C = C \cap A$. What is the size of the largest subset family of U that has no sunflower?

Cap set problem

Three vectors $a, b, c \in \mathbb{F}_3^n$ form a progression of length 3 if $a + b + c = 0$. What is the cardinal of the largest **cap set** (set of vectors avoiding such progressions) in $\mathbb{F}_3^n$?

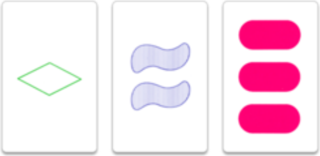

Naslund–Sawin bound on sunflower-free set

Theorem (Naslund–Sawin 2016)

Let $\mathcal{F}$ be a sunflower-free family of $\{1,2,\ldots,n\}$. Then

$$|\mathcal{F}| \le 3(n+1)\sum_{k\le n/3}\binom{n}{k} = \left(3\cdot 2^{-2/3}\right)^n e^{o(n)}.$$

Idea: A notion called **slice rank**, first used implicitly by Croot–Lev–Pach (2016) on progression-free sets in $\mathbb{Z}_4^n$.

First result that breaks $2^n e^{o(n)}$!

A polynomial model for the sunflower-free set

Let $U = \{1, 2, \ldots, n\}$, and $v_1, \ldots, v_n$ be the canonical base of $\mathbb{F}_3^n$. For $A \subseteq U$, we define $v_A = \sum_{i \in A} v_i$.

Given a polynomial $P(X_1, \ldots, X_n)$ and a vector $u = \sum_{i=1}^n x_i v_i \in \mathbb{F}_3^n$, we define $P(u) = P(x_1, \ldots, x_n)$.

Proposition

Let A, B, C be three sets without one set being the proper subset of another. The sets A, B, C form a sunflower or $A = B = C$ iff $P(v_A, v_B, v_C) = 1$, with

$$P(X_1, \ldots, X_n, Y_1, \ldots, Y_n, Z_1, \ldots, Z_n) = \prod_{i=1}^{n} (2 - (X_i + Y_i + Z_i)).$$

Proof: Since no set is a proper subset of the other, w.l.o.g., we only need to avoid $i \in (A \cap B) \setminus C$, which means $x_i = y_i = 1, z_i = 0$, which implies $x_i + y_i + z_i - 2 = 0$.

Polynomial as tensor

A polynomial $P(X_1, \dots, X_n, Y_1, \dots, Y_n, Z_1, \dots, Z_n)$ in $\mathbb{F}_3$

$$\Leftrightarrow$$

A tensor T in $\mathbb{F}_3^n \otimes \mathbb{F}_3^n \otimes \mathbb{F}_3^n$ with $T(u, v, w) = P(u, v, w)$

Let $\mathcal{F}$ be a sunflower-free family in U, and $T_{\mathcal{F}}$ the sub-tensor of T with coordinates restricted to all v_A with $A \in \mathcal{F}$.

Proposition

$T_{\mathcal{F}}$ is a diagonal tensor, that is, $T_{\mathcal{F}}(u, v, w) = 1$ iff $u = v = w$.

Idea: Upper bound on "big diagonals" $\Rightarrow$ upper bound on sunflower-free set.

We want some notion of rank to capture the size of "big diagonals".

Slice rank of a function

Let A be a finite set. A function $S : A \otimes A \otimes A \to \mathbb{F}$ is a **slice** if it has one of the following forms:

$$S(u,v,w) = f(u)g(v,w) \text{ or } f(v)g(u,w) \text{ or } f(w)g(u,v).$$

The **slice rank** of a function $F : A \otimes A \otimes A \to \mathbb{F}$, denoted by $\mathrm{sr}(F)$, is the minimum number of slices needed to sum to F.

Property: Let $T_A : A \otimes A \otimes A \to \mathbb{F}$, and T_B its restriction on $B \otimes B \otimes B$ with $B \subseteq A$. Then $\mathrm{sr}(T_B) \leq \mathrm{sr}(T_A)$.

Lemma (Special case of Tao (2016))

The slice rank of the function $F(u,v,w) = \sum_{a \in A} c_a \delta_a(u)\delta_a(v)\delta_a(w)$ is the number of non-zero coefficients $c_a \in \mathbb{F}$.

Proof: delayed.

Slice rank of the sunflower polynomial

$$P(\underline{X}, \underline{Y}, \underline{Z}) = \prod_{i=1}^{n} (2 - (X_i + Y_i + Z_i)).$$

For a monomial $X_1^{a_1} \cdots X_n^{a_n} Y_1^{a_1} \cdots Y_n^{a_n} Z_1^{a_1} \cdots Z_n^{a_n}$ in $P(\underline{X}, \underline{Y}, \underline{Z})$, we have $\sum_{i=1}^{n} a_i + \sum_{i=1}^{n} b_i + \sum_{i=1}^{n} c_i \leq n$. One of the total powers of X, Y and Z must be $\leq n/3$.

$$P(\underline{X}, \underline{Y}, \underline{Z}) = \sum_{a_1 + \cdots + a_n \leq n/3} X_1^{a_1} \cdots X_n^{a_n} P_{a_1, \ldots, a_n}(\underline{Y}, \underline{Z}) + \cdots .$$

Thus we have (since all $a_i \leq 1$)

$$\mathrm{sr}(P) \leq 3 \sum_{k \leq n/3} \binom{n}{k}.$$

Proof of upper bound

Let $\mathcal{F}$ be a sunflower-free family, with $\mathcal{F} = \bigcup_{\ell \geq 0} \mathcal{F}_\ell$ the partition by number of elements. Sets in $\mathcal{F}_\ell$ are never proper subset of each other.

Let $A_\ell = \{v_A \mid A \in \mathcal{F}_\ell\}$. The function P is diagonal on A_ℓ, thus $|\mathcal{F}_\ell| = \mathrm{sr}_{A_\ell}(P) \leq \mathrm{sr}(P)$.

We thus have

$$|\mathcal{F}| \leq 3(n+1) \sum_{k \leq n/3} \binom{n}{k} = \left(3 \cdot 2^{-2/3}\right)^n e^{o(n)}.$$

New bound on cap set problem

Polynomial:

$$P(\underline{X}, \underline{Y}, \underline{Z}) = \prod_{i=1}^{n} (1 - (X_i + Y_i + Z_i)^2).$$

Theorem (Ellenberg–Gijswijt (2016))

The size of a cap set in $\mathbb{F}_3^n$ *is* $o(2.756^n)$.

General result for any finite field. Kleinberg–Sawin–Speyer gave a concrete construction on a lower bound that matches within a subexponential factor.

A general strategy

Given a problem concerning avoiding some structure.

1. Construct a polynomial P whose zeros are exactly on "everything equal" or "things forming the structure", which is a product of the same polynomial on different sets of variables in many cases;
2. The function P restricted to an avoiding family $\mathcal{F}$ will then be diagonal;
3. Compute the slice rank of P, which is an upper bound of the size of $\mathcal{F}$;
4. Hopefully this bound will be a breakthrough, or not.

Can we know the power of the method?

Slice rank for tensors

We consider tensors in $V_1 \otimes V_2 \otimes \cdots \otimes V_k$. We define in the natural way the j^{th} tensor product

$$\otimes_j : V_j \otimes \bigotimes_{1 \leq i \leq k, i \neq j} V_i \to \bigotimes_{1 \leq i \leq k} V_i.$$

A **slice** is any element of the form $v_j \otimes_j v_{\neq j}$ for any j. The **slice rank** of a tensor T is the minimum number of slices that sum to T.

Example: For V_1 (resp. V_2, V_3) the space of polynomials of X_i (resp. Y_i, Z_i) in $\mathbb{F}$, the slice rank of tensors in $V_1 \otimes V_2 \otimes V_3$ is the slice rank of polynomials.

Property: Let T be a tensor in $V_1 \otimes V_2 \otimes \cdots \otimes V_k$ and T' a sub-tensor of T. Then $\mathrm{sr}(T') \leq \mathrm{sr}(T)$.

Slice rank of a polynomial and its value tensor

Let P be a polynomial in a finite field $\mathbb{F}$ with k sets of n variables. P is a tensor in $V_1 \otimes \cdots \otimes V_k$, where V_i is spanned by monomials in the i^{th} set of variable.

Let T_P be the **value tensor** of P in $(\mathbb{F}^n)^{\otimes k}$ defined by

$$T_P = \sum_{\underline{v}_1, \ldots, \underline{v}_k \in \mathbb{F}^n} P(\underline{v}_1, \ldots, \underline{v}_k) \underline{v}_1 \otimes \cdots \otimes \underline{v}_k.$$

Proposition

We have $\mathrm{sr}(P) = \mathrm{sr}(T_P)$.

Proof: Equivalence on slices.

Slice rank and diagonal

We now consider tensors of the form $V^{\otimes k}$. Let S be a basis of V.

Lemma (Special case of Tao (2016))

The slice rank of the tensor $T = \sum_{a \in S} c_a a^{\otimes k}$, denoted by $\mathrm{sr}(F)$, is the number of non-zero coefficients $c_a \in \mathbb{F}$.

Proof: Again delayed.

For $\mathcal{S} \subseteq V^k$ structures to avoid (*e.g.* sunflowers), suppose we have a polynomial P in $\mathbb{F}$ with non-zero values only on $u_1 = \cdots = u_k$ or S.

An avoiding family $\mathcal{F} \subseteq V$ gives a sub-tensor $T_P|_{\mathcal{F}^{\otimes k}}$ that is a diagonal.

We thus have $|\mathcal{F}| = \mathrm{sr}(T_P|_{\mathcal{F}^{\otimes k}}) \leq \mathrm{sr}(T_P) = \mathrm{sr}(P)$.

Upper bound on $\mathrm{sr}(P) \Rightarrow$ upper bound on $\mathcal{F}$.

Slice rank (dual version)

Let T be a tensor in $V = V_1 \otimes V_2 \otimes \cdots \otimes V_k$. Let W_i be the dual space of V_i, with the canonical pairing $\langle \cdot, \cdot \rangle_i$. Let $W = W_1 \otimes \cdots \otimes W_k$, and we define the pairing

$$\langle w_1 \otimes \cdots \otimes w_k, v_1 \otimes \cdots \otimes v_k \rangle = \prod_{i=1}^{k} \langle w_i, v_i \rangle_i.$$

Proposition

We have $\mathrm{sr}(T) \leq r$ iff there are sub-spaces W_i^T for all i such that the co-dimensions of W_i^T for all i sum to r, and that $\langle \cdot, v \rangle$ is zero on $\bigotimes_{i=1}^{k} W_i^T$.

Proof: There must be a component that annihilates the pairing.

Projections and upper bound

We fix a basis S_i for each V_i. We define $\pi_i(s_1 \otimes \cdots \otimes s_k) = s_i$ for all v_i in S_i.

Proposition

Let T be a tensor in $V_1 \otimes \cdots \otimes V_k$, and Γ its support w.r.t. $(S_i)_{1 \leq i \leq k}$. We have

$$\mathrm{sr}(T) \leq \min_{\Gamma = \Gamma_1 \cup \cdots \cup \Gamma_k} \sum_{i=1}^{k} |\pi_i(\Gamma_i)|.$$

Proof: Decompose by the vector obtained after projection:

$$\begin{aligned} T &= \sum_{i=1}^{k} \sum_{(s_1 \otimes \cdots \otimes s_k) \in \Gamma_i} c_* s_1 \otimes \cdots \otimes s_k \\ &= \sum_{i=1}^{k} \sum_{s_i \in \pi_i(\Gamma_i)} c_* s_i \otimes_i v_{s_i, \neq i}. \end{aligned}$$

Each summand is a slice.

Lower bound

We suppose that, for each S_i, we have a total order $\leq_i$. They induce a partial order on vectors $s_1 \otimes \cdots \otimes s_k$ for $s_i \in S_i$.

Proposition

Let T be a tensor in $V_1 \otimes \cdots \otimes V_k$, Γ its support w.r.t. $(S_i)_{1\leq i\leq k}$, and Γ' the set of **maximal** *elements in Γ. We have*

$$\operatorname{sr}(T) \geq \min_{\Gamma'=\Gamma'_1\cup\cdots\cup\Gamma'_k} \sum_{i=1}^{k} |\pi_i(\Gamma'_i)|.$$

Remark: $\operatorname{sr}(T)$ does not depend on basis.

We only need to show that there is a covering $\Gamma'_1, \ldots, \Gamma'_k$ of Γ' such that $\operatorname{sr}(T) \geq \sum_{i=1}^{k} |\pi_i(\Gamma'_i)|$.

Proof using the dual definition

Suppose that $S_i = \{s_{i,1} \leq \cdots \leq s_{i,d_i}\}$, with $d_i = \dim(V_i)$. Let $s^*_{i,j}$ be the dual of $s_{i,j}$ in W_i.

Consider $W^T = W_1^T \otimes \cdots \otimes W_k^T$ that annihilates T on the pairing $\langle \cdot, \cdot \rangle$. There is a basis $(w_{i,j})_{1 \leq j \leq e_i}$ of W_i^T in a row-echelon form:

$$
\begin{aligned}
w_{i,1} &= s^*_{i,t_1} + \cdots + * s^*_{i,t_2} + \cdots + * s^*_{i,t_{e_i}} + \cdots \\
w_{i,2} &= \phantom{s^*_{i,t_1} + \cdots + *} s^*_{i,t_2} + \cdots + * s^*_{i,t_{e_i}} + \cdots \\
&\vdots \\
w_{i,e_i} &= \phantom{s^*_{i,t_1} + \cdots + * s^*_{i,t_2} + \cdots + *} s^*_{i,t_{e_i}} + \cdots .
\end{aligned}
$$

Let $S'_i = \{s_{i,t_1}, \ldots, s_{i,t_{e_i}}\}$. We claim that $v = s'_1 \otimes \cdots \otimes s'_k$ with $s'_i \in S'_i$ for all i is not in Γ'.

Suppose the contrary. By maximality of elements in Γ', all $s^\dagger_1 \otimes \cdots \otimes s^\dagger_k$ with $s^\dagger_i \geq s'_i$ for all i are not in Γ, except for v itself.

Then $\langle v, T \rangle \neq 0$ by row-echelon form.

Proof using the dual definition (cont'd)

Any $v = s'_1 \otimes \cdots \otimes s'_k$ with $s'_i \in S'_i$ for all i is not in Γ'.

We now take the covering $\Gamma'_i = \{s_1 \otimes \cdots \otimes s_k \mid s_i \notin S'_i\}$. We have $\pi_i(\Gamma'_i) = d_i - e_i$, which is also the co-dimension of W_i^T.

Therefore, for all annihilator W^T, there is a covering $\Gamma'_1, \ldots, \Gamma'_k$ of Γ' such that

$$\sum_{i=1}^{k} \mathrm{codim}(W_i^T) \leq \sum_{i=1}^{k} |\pi_i(\Gamma'_i)|.$$

We conclude by the dual definition of slice rank.

Corollary on diagonal tensor

We consider diagonal tensors in $V^{\otimes k}$ over a field $\mathbb{F}$, with S a basis of V.

Corollary

Let $T = \sum_{a \in S} c_a a^{\otimes k}$. Then $\mathrm{sr}(T)$ is the number of non-zero coefficients c_a. In particular, for $T = \sum_{a \in S} a^{\otimes k}$, we have $\mathrm{sr}(T) = |S|$.

Proof: Consider a total order $\leq_S$ on S, and we form a partial order by taking $\leq_S$ on all components except the last, which has the reversed total order. Then the diagonal is an anti-chain without overlapping elements in projections.

Slice rank of tensor powers

Recall that many problems lead to polynomials that are product of the same polynomial on different set of variables, which leads to consider the slice rank of **tensor powers**.

Given a tensor T in $V_1 \otimes \cdots V_k$, with S_i a basis of V_k, we want to compute asymptotically $\mathrm{sr}(T^{\otimes n})$ for $T^{\otimes n}$ in $(V_1 \otimes \cdots \otimes V_k)^{\otimes n} \cong V_1^{\otimes n} \otimes \cdots \otimes V_k^{\otimes n}$.

We suppose that all S_i come with a total order $\leq_i$. We denote by Γ the support of T w.r.t. all S_i, and Γ' the set of maximal elements of Γ.

Upper and lower bounds

Proposition

For $n \to \infty$, we have

$$\exp(n(H' + o(1))) \leq \mathrm{sr}(T^{\otimes n}) \leq \exp(n(H + o(1))),$$

where

$$H = \sup_X \min(h(\pi_1(X)), \ldots, h(\pi_k(X))),$$
$$H' = \sup_{X'} \min(h(\pi_1(X')), \ldots, h(\pi_k(X'))),$$

with X (resp. X') a probability distribution on Γ (resp. Γ'), and $h(\cdot)$ the entropy function.

Sawin and Tao also provided some criteria for the maximizing distribution X.

A sketch of proof

We only need to show for any Γ that

$$\min_{\Gamma^{\otimes n}=\Gamma_{n,1}\cup\cdots\cup\Gamma_{n,k}} \sum_{i=1}^{k} |\pi_{n,i}(\Gamma_{n,i})| = \exp(n(H+o(1))).$$

By compactness, we can take X that reaches the sup H.

$\geq$: consider vectors in $\Gamma^{\otimes n}$ that are "ϵ-close" to X, there are roughly $\exp(n(H+o(1)))$ such vectors, and at least one partition contains $1/k$ of them.

$\leq$: we can cover Γ by $O(\exp(o(n))$ "ϵ-close" balls centered at some X.

Sunflower: bounds

We recall that the "sunflower polynomial" is $2 - X - Y - Z$ in $\mathbb{F}_3$. We now consider the polynomial space.

We have

$$\Gamma = \{(1,0,0),(0,1,0),(0,0,1),(0,0,0)\},$$
$$\Gamma' = \{(1,0,0),(0,1,0),(0,0,1)\}.$$

A maximizing distribution for both is $X = \frac{1}{3}(1,0,0) + \frac{1}{3}(0,1,0) + \frac{1}{3}(0,0,1)$, which leads to

$$H = \frac{1}{3}\log(3) + \frac{2}{3}\log(3/2) = \log(3 \cdot 2^{-2/3}).$$

This also shows that we cannot do better $\left(\mathrm{sr}(T^{\otimes n}) = \exp(nH + o(n))\right)$.

Capset: bounds

We recall that the "cap set polynomial" is $(1-(X+Y+Z)^2)$ in $\mathbb{F}_3$.

Reason: Cap set condition on a coordinate is that $X+Y+Z=0$.

We have

$$\Gamma = \{(2,0,0),(0,2,0),(0,0,2),(1,1,0),(1,0,1),(0,1,1),(0,0,0)\},$$
$$\Gamma' = \{(2,0,0),(0,2,0),(0,0,2),(1,1,0),(1,0,1),(0,1,1)\}.$$

A maximizing distribution should take the form

$$X = \alpha((2,0,0)+(0,2,0)+(0,0,2))+\beta((1,1,0),(1,0,1),(0,1,1))+\gamma(0,0,0).$$

By maximizing the corresponding H, we have the result, which has $\gamma=0$. It means that we cannot do better $(\mathrm{sr}(T^{\otimes n})=\exp(nH+o(n)))$.

Limitation of the polynomial method

Proposition

Let $k \geq 8$, and G a finite abelian group. Let $\mathbb{F}$ be any field, and $V_1 = \cdots = V_k$ the space of functions from G to $\mathbb{F}$.
Let F be any $\mathbb{F}$-valued function that is zero only on k-progressions or on the diagonal. Then $\mathrm{sr}(F) = |G|$.

Proof ideas: first reduce the problem to the cyclic group $\mathbb{Z}/n\mathbb{Z}$, then show an ordering that makes every constant progression a maximal element (thus in Γ').

Ordering: $(\leq, \leq, \leq, \geq, \leq, \geq, \geq, \geq)$.

Change of basis

We consider the polynomial $P = 1 + (1 + Z)(X + Y)$ in $\mathbb{F}_2$.

Meaning: Three sets A, B, C such that $A \Delta B \subseteq C$.

$$\Gamma = \{(1,0,0), (0,1,0), (1,0,1), (0,1,1), (0,0,0)\}$$

$X = \frac{1}{2}((1,0,0),(0,1,1))$ maximized $H = \log(2)$.

But a change of variable $Z \leftarrow 1 + Z$ gives $Q = 1 + XZ + YZ$, with entropy $H = \log(3 \cdot 2^{-2/3})$.

Discussion

Observations:

- The lower bounds are limits of the method, and does not give concrete construction on original problems.
- Not limited to sub-tensors with only zeros outside the diagonal.
- The bounds does not depend on degree, but on the monomials in the defining polynomial.

Further directions:

- More applications?
- Synergies with other methods?
- Use the fact that slice rank is basis-independent?