2021

The Slice Rank Polynomial Method

Thomas C. Martinez
Harvey Mudd College

The Slice Rank Polynomial Method

Thomas Martinez

Mohamed Omar, Advisor

Arthur Benjamin, Reader

HARVEY MUDD COLLEGE

Department of Mathematics

May, 2021

Abstract

Suppose you wanted to bound the maximum size of a set in which every k-tuple of elements satisfied a specific condition. How would you go about this? Introduced in 2016 by Terence Tao, the slice rank polynomial method is a recently developed approach to solving problems in extremal combinatorics using linear algebraic tools. We provide the necessary background to understand this method, as well as some applications. Finally, we investigate a generalization of the slice rank, the partition rank introduced by Eric Naslund in 2020, along with various discussions on the intuition behind the slice rank polynomial method and other possible avenues for generalization.

Contents

List of Figures

Acknowledgments

I would like to thank my advisor, Mohamed Omar, for his guidance and flexibility throughout this thesis process. I always left our meetings with a sense of direction and ready to take on the week. I would also like to thank Arthur Benjamin for his helpful feedback throughout my many drafts. I can safely say that my introduction to combinatorics in his section of Math 55 is what solidified my desire to continue academia in mathematics. Finally, I would like to thank my peers and family for their continual support, encouragement, and feedback throughout the process of writing this thesis.

Chapter 1

Introduction

The field of extremal combinatorics obtains its name from the types of questions it studies, generally of the form: if a collection of finite objects satisfies certain restrictions, how small or how large can it be? A classic example of this is an introductory problem in Ramsey theory, which any HMC student who just took Discrete Mathematics could answer: what is the maximum number of people that can go to a party such that there is not a collection of 3 people who either all know each other or all do not know each other?[1] There are countless examples of problems in extremal combinatorics, some more familiar than others, and various ways to solve them – the principle of inclusion-exclusion, the pigeonhole principle, or induction to name a few. This thesis concerns itself with one of the more modern methods: the slice rank polynomial method.

In May of 2016, Croot et al. (2017) introduced a powerful technique named the polynomial method, which Ellenberg and Gijswijt (2017)[2] used to settle the long-standing cap-set problem (more about this in subsection 3.1.2). In August of the same year, Tao, in one of his blog posts, presented a symmetrized reformulation of the polynomial method, now known as the slice rank[3] method. As you can see by the timeline, this was an exciting period in the field of extremal combinatorics. Already in its short history, the slice rank polynomial method has been used to produce exciting results like the proof of a special case of a 50 year-old conjecture posed by Erdös

[1]The answer is 5.

[2]Discovered independently but wrote the ensuing paper together

[3]The name "slice rank" is due to Blasiak et al. (2017).

and Szemerédi by Naslund and Sawin (2017) (more about this in subsection 3.1.1).

The slice rank polynomial method is a natural generalization of a manifestation of the linear algebra method[4] in extremal combinatorics that can solve problems of the form: "how large (or how small) can a family of sets be such that there is a restriction on *single* set and/or *pairs* of sets." For example, we can restrict the size of sets and/or the size of their pairwise intersections.

The slice rank polynomial method provides a manner in which linear algebra can be used to solve problems in extremal combinatorics involving restrictions on *more* than two sets from the family. In general, the slice rank polynomial method works as follows, assuming we wish to solve a problem with a restriction on a k-tuple of sets:

1. Suppose we find the largest such set family $\mathcal{S}$ that satisfies such a restriction.
2. Create a diagonal k-tensor $T : \mathcal{S}^k \to \mathbb{F}$, where $\mathbb{F}$ is a field, with non-zero diagonal entries that encapsulates the restrictions of the k-tuple of sets.
3. By the slice rank lemma, we have that the slice rank of T is equal to $|\mathcal{S}|$.
4. We can now upper bound the slice rank of this k-tensor to achieve an upper bound[5] on $|\mathcal{S}|$.

The focus of this expository thesis is to cultivate a deeper understanding on the slice rank polynomial method. Where might we be able to apply it? Will it provide stronger bounds that those that already exist due to classical methods? Can we generalize this method even further? In order to answer these questions, our thesis begins with background on k-tensors and presents a proof of the key slice rank lemma, proved by Tao (2016b) in Chapter 2. Chapter 3 then provides some of the existing applications of the slice rank polynomial method, like that of Ellenberg and Gijswijt (2017) and Naslund and Sawin (2017), as well as some examples of what may "go wrong" when attempting to apply the slice rank polynomial method. We end Chapter 3

[4]As you can imagine, there are numerous "linear algebra methods" in solving extremal combinatorics. We leave the textbook by Babai and Frankl (1988) which focuses solely on linear algebra methods in combinatorics and, for a shorter introduction, this set of notes taken from an eight lecture series for graduate students in combinatorics at UCSD during the Fall 2019 Quarter.

[5]We can suppose we find the smallest set family $\mathcal{T}$ that satisfies the restrictions, and then find a lower bound the slice rank of the k-tensor to achieve a lower bound on $|\mathcal{T}|$.

with a discussion of what we may expect the slice rank to be, given the form the k-tensor we define takes. Lastly, Chapter 4 discusses the partition rank, a generalization of the slice rank introduced by Naslund (2020b), as well as a short discussion on other possible avenues for generalizations of the slice rank polynomial method.

This thesis presumes basic knowledge in linear algebra, so the reader has sufficient background if they are an undergraduate math major who just took an introductory course or two in linear algebra.

Chapter 2

Tensors

Before exploring the slice rank polynomial method, we will discuss the linear algebra method which is well-equipped for handling extremal problems with restrictions on individual sets and/or pairs of sets.

2.1 Oddtown

Let us imagine a scenario: the case of Oddtown.[1] Oddtown has N residents and they love forming clubs and maximizing the number of them. However, they follow an *odd* tradition when forming these clubs:

(i) No two clubs may have exactly the same members;

(ii) Each club must have an **odd** number of people;

(iii) Any two clubs must share an **even** number of people.

The mayor of Oddtown would like to know how many clubs he may create, given N residents. Before we can provide the mayor with an answer, we require some notation.

[1]The problem with Oddtown is usually presented and contrasted with Eventown, where the set of rules are slightly different, in that we have each club with an even number of people but leads to a much larger maximal number of sets (an upper bound of $2^{n/2}$). The case of Eventown does not interest us as the manifestation of the linear algebra method is not what we call the matrix method. However, for the motivated reader, we leave this set of notes by Tibor Szabó and this set of notes by Calum Buchanan.

Definition 2.1.1. *Let* $[n]$ *denote the set* $\{1, 2, \ldots, n\}$ *for every* $n \in \mathbb{N}_+$. *Also, let* $2^{[n]}$ *denote the collection of subsets of* $[n]$.

We can now provide the mayor with their long-awaited answer.

Theorem 2.1.2 (Berlekamp (1969))**.** *Let* $\mathcal{F} \subseteq 2^{[n]}$ *be a set family satisfying Oddtown's rules. Then,* $|\mathcal{F}| \leq n$.

Before moving on to the proof, it is important to understand the connection between linear algebra and combinatorics. What's stopping us from solving this problem purely combinatorially?[2] As it turns out, in a typical extremal combinatorics problem, the greater the number of extremal families, the less likely a purely combinatorial proof leads to fruition. This is because a combinatorial proof has to consider all extremal families, and if these families are very combinatorially different, it may lead to an unmanageable number of case distinctions (this explanation is due to Szabó (2019)). The Oddtown theorem is one of these situations. In fact, one can prove that the number of extremal families is super-exponential[3] (see Babai and Frankl (1988)[Exercise 1.1.14]).

As it turns out, the Oddtown theorem is extremely suitable to linear algebraic methods. The connection — for this problem and, as we will see, for many others — between linear algebra and combinatorics is provided through the *characteristic vector*.

Definition 2.1.3. *The* ***characteristic vector*** *of a subset* T *of a set* S *is the vector*

$$x_T := (x_s)_{s \in S},$$

such that $x_s = 1$ *if* $s \in T$ *and* $x_s = 0$ *if* $s \notin T$. *For example, taking* $S = \{1, 2, 3, 4\}$ *and* $T = \{2, 3\}$, *our characteristic vector is* $x_T = (0, 1, 1, 0)$. *We also use the notation* $(x_T)_i$ *to denote the* i*th value of* x_T. *For example,* $(x_T)_2 = 1$ *and* $(x_T)_4 = 0$.

The reader may realize that the square of the length of the characteristic vector x_T gives us the size of our set T. So, we can easily grab hold of the condition that $|T|$ is an odd number. However, notice that the rules of Oddtown also has a condition on the sizes of pairs of sets. We must then take advantage of another algebraic object, the dot product.

[2]As it turns out, there exists a combinatorial proof, due to Petrov (2016), presented in Buchanan's notes. However, we aren't always so lucky!

[3]A function is *super-exponential* if it grows faster than any exponential function. That is, for every constant c, $\lim_{n \to \infty} f(n)/c^n = \infty$.

Definition 2.1.4. *The **dot product** of two vectors* $a = [a_1, a_2, \ldots, a_n]$ *and* $b = [b_1, b_2, \ldots, b_n]$ *of the same dimension is*

$$a \cdot b = \sum_{i=1}^{n} a_i b_i = a_1 b_1 + a_2 b_2 + \cdots + a_n b_n.$$

For example, let $a = [1, 2, 3]$ *and* $b = [2, 5, 8]$*. Then* $a \cdot b = (1 \cdot 2) + (2 \cdot 5) + (3 \cdot 8) = 2 + 10 + 24 = 36$*.*

The key insight for the Oddtown problem is to realize that for two sets $A, B \in 2^{[n]}$,

$$|A \cap B| = \sum_{i=1}^{n} (v_A)_i (v_B)_i = v_A \cdot v_B.$$

Indeed, this is due to how $(v_A)_i (v_B)_i = 1$ if and only if $(v_A)_i = (v_B)_i = 1$. That is, $(v_A)_i (v_B)_i = 1$ if and only if $i \in A \cap B$. In the context of the Oddtown problem, we have the condition that, for any two sets A, B in a set family that satisfies Oddtown's rules,

$$v_A \cdot v_B = \begin{cases} \text{odd}, & \text{if } i = j; \\ \text{even}, & \text{if } i \neq j. \end{cases}$$

We can simplify this further if we work in $\mathbb{F}_2$, the finite field of 2 elements, where our elements are just 0 and 1 and the arithmetic operations are performed modulo 2. With this, we obtain

$$v_A \cdot v_B = \begin{cases} 1, & \text{if } i = j; \\ 0, & \text{if } i \neq j. \end{cases}$$

We can now prove the Oddtown theorem.

Proof of Theorem 2.1.2. Let $\mathcal{F} = \{C_i \mid 1 \leq i \leq m\}$ be a set family satisfying Oddtown's rules. For each set C_i, we will associate v_{C_i} as its characteristic vector. Working over $\mathbb{F}_2$, we note that we have $v_{C_i} \cdot v_{C_i} = 1$ for every $1 \leq i \leq m$, and $v_{C_i} \cdot v_{C_j} = 0$ for all $i \neq j$.

Suppose we have a linear combination $\sum_{i=1}^{m} \alpha_i v_{C_i} = 0$, with $\alpha_i \in \mathbb{F}_2$. For every $1 \leq j \leq m$, we take the dot product of this linear combination and obtain

$$0 = 0 \cdot v_{C_j} = \left(\sum_{i=1}^{m} \alpha_i v_{C_i} \right) \cdot v_{C_j} = \sum_{i=1}^{m} \alpha_i (v_{C_i} \cdot v_{C_j}) = \alpha_j.$$

This implies that $\alpha_j = 0$ for all $1 \leq j \leq m$, which implies that the characteristic vectors form a linearly independent set. So, their number cannot exceed the dimension of the space, we have $m \leq \dim \mathbb{F}_2^n = n$. □

As you can see from this proof and discussion, we are well-equipped for handling problems which handle restrictions on single sets and/or pairs of sets, thanks to the connection to objects like the characteristic vector and the dot product. There are numerous examples of problems in extremal combinatorics being solved in this manner, like Fisher's Inequality, and the neighboring Eventown problem. An issue arises, however, when we wish to go beyond that, and handle restrictions on, say, triples of sets. Thankfully, this is precisely what the slice rank polynomial method is equipped for. Before we introduce the slice rank polynomial method, we require some introduction to k-tensors and what slices actually are.

2.2 Basics

As we have seen before, there are methods, using linear algebra, to bound the size of a family of sets $\mathcal{F}$, where there were restrictions on each set, or each *pair* of sets. However, how can we generalize this idea, so we can place restrictions on each triple of sets, or quadruple of sets, and so on. To do so, we provide a generalization of a matrix, a k-tensor.[4]

Definition 2.2.1. *Let X be a finite set and $\mathbb{F}$ a field. A* **k-tensor** *is a function $T : X^k \to \mathbb{F}$. A k-tensor is* **diagonal** *if $T(x_1, \dots, x_k) \neq 0$ implies $x_1 = \cdots = x_k$.*

It may come as a surprise, but there is no restrictions on our function T for it to be a k-tensor. That is, it does not have to be linear or injective or anything of the sort. To aid with this definition, let us look at some simple examples.

Example 1 (The Zero Tensor)**.** *The function $Z : X^k \to \mathbb{F}$ where $Z(x_1, \dots, x_k) = 0$ for all $x_1, \dots, x_k \in X$ is a k-tensor for all $k \geq 1$.*

Example 2. *The function $M : \{1,2,3\}^3 \to \mathbb{R}$ where $M(x, y, z) = xy + yz$ for all $x, y, z \in \{1,2,3\}$ is a 3-tensor. We note that M is not diagonal as $M(1,2,1) = 1 * 2 + 2 * 1 = 2 + 2 = 4 \neq 0$.*

[4]This chapter will focus on building an intuition for the reader. For a more rigorous introduction, using the notion of tensor products and other concepts in exterior and multilinear algebra, see Tao's original blog post and this follow-up blog post dedicated to the slice rank.

Definition 2.2.2. *Let X be a finite set and $\mathbb{F}$ a field. Then, for any k-tensors, f, g, and any scalar $c \in \mathbb{F}$, we have*

i. $(f+g)(x_1, \ldots, x_n) = f(x_1, \ldots, x_n) + g(x_1, \ldots, x_n)$, *and*

ii. $(c \cdot f)(x_1, \ldots, x_n) = c \cdot f(x_1, \ldots, x_n)$.

That is, the sum of two k-tensors and the scalar multiple of a k-tensor is still a k-tensor.

These definitions are relatively natural to make. More abstractly, the sum of two functions with the same domain and co-domain will *still be* a function with the same domain and co-domain, and the scalar multiplication property is obtained due to $\mathbb{F}$ being a field. A natural question may be how multiplication works, and this is a little more complicated. We provide another definition as to how multiplication works, justifying these definitions in another section. First, some notation. Given variables $x_1, \ldots, x_n$ and a set $S \subseteq \{1, \ldots, n\}$ with $S = \{s_1, \ldots, s_k\}$, we use the notation x_S to denote the subset of variables

$$x_{s_1}, \ldots, x_{s_k}.$$

So, for a function f of k variables, we have

$$f(x_S) = f(x_{s_1}, \ldots, x_{s_k}).$$

We can now define multiplication of k-tensors.

Definition 2.2.3. *Let X be a finite set and $\mathbb{F}$ a field. $I, J \subseteq \{1, \ldots, n\}$, $|I| = k_1$, $|J| = k_2$, and $I \cap J = \emptyset$. Then, for a k_1-tensor $f(x_I)$ where and a k_2-tensor $g(x_J)$, we have,*

$$(f \cdot g)(x_1, \ldots, x_n) = f(x_I) \cdot g(x_J).$$

Let us further illuminate this definition with a couple of examples.

Example 3. *Let $f : \mathbb{F}_3^2 \longrightarrow \mathbb{R}$ be the 2-tensor $f(x_1, x_3) = x_1 + x_3$, and let $g : \mathbb{F}_3 \longrightarrow \mathbb{R}$ be the 1-tensor defined by $g(x_2) = x_2$. Then, we see that*

$$(f \cdot g)(x_1, x_2, x_3) = f(x_1, x_3) \cdot g(x_2) = (x_1 + x_3)(x_2) = x_1x_2 + x_2x_3.$$

We see that $f \cdot g$ is indeed a 3-tensor. It also has a special property, we will investigate shortly, so stay tuned.

Example 4. *Let* $f : \mathbb{F}_3^2 \to \mathbb{R}$ *be the* 2*-tensor* $f(x_1, x_3) = x_1 + x_3$*, and let* $g : \mathbb{F}_3^2 \to \mathbb{R}$ *be the* 2*-tensor defined by* $g(x_2, x_4) = x_2 - x_4$*. Then, we see that*

$$\begin{aligned}(f \cdot g)(x_1, x_2, x_3, x_4) = f(x_1, x_3) \cdot g(x_2, x_4) &= (x_1 + x_3)(x_2 - x_4) \\ &= x_1x_2 - x_1x_4 + x_2x_3 - x_3x_4.\end{aligned}$$

We see that $f \cdot g$ *is indeed a* 4*-tensor.*

Now that we have defined a tensor, we want a way to understand how "complicated" a k-tensor is. This leads to a nominal notion of the rank of a tensor.

Definition 2.2.4. *Let* T *be a* k*-tensor. The* **tensor rank** *of* T *is the smallest non-negative integer* r *such that we can write*

$$T(x_1, \ldots, x_k) = \sum_{i=1}^{r} \prod_{j=1}^{k} f_{i,j}(x_j),$$

where each $f_{i,j}$ *is a* 1*-tensor. For notational purposes, we denote* $r = \operatorname{trk}(T)$*.*

In general, calculating the tensor rank is very difficult. Håstad (1989) showed that, even for 3-tensors, calculating the tensor rank is an NP-Complete[5] problem, and while we do not go into detail as to what exactly that means for this thesis, one may imagine it as a *difficult* problem.

However, there is more than one way to define the rank of a tensor. To do so, we introduce another object.

Definition 2.2.5. *A* k*-tensor* S *is a* **slice** *if there exists a* 1*-tensor* f *and a* $(k-1)$*-tensor* g *such that for some* $1 \le i \le k$*,*

$$S(x_1, \ldots, x_k) = f(x_i)\, g(x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_k),$$

for all $x_1, \ldots, x_k \in X$*.*

What is so special about a slice? Intuitively, if a k-tensor f is a slice, we can imagine is as a simpler tensor, a $(k-1)$-tensor g, where we are modifying g very slightly. Let's give a more concrete example as to how to think about it. Imagine $\mathbb{R}^3$ for a moment, and a function $f : \mathbb{R}^3 \to \mathbb{R}$. If we can describe f

[5] For a short explanation for what it means to be an NP-Complete problem, we refer the reader to this Stack Overflow post.

as the product of two functions $g_1 : \mathbb{R}^2 \to \mathbb{R}$ and $g_2 : \mathbb{R} \to \mathbb{R}$, it is a slice. That is, if f can instead be defined on a plane, which is one dimension lower than a three-dimensional manifold, multiplied some other function that keeps into account our final coordinate, it is a slice. Essentially, we can think of f being encoded in a lower dimensional space, which makes it a much simpler kind of tensor.

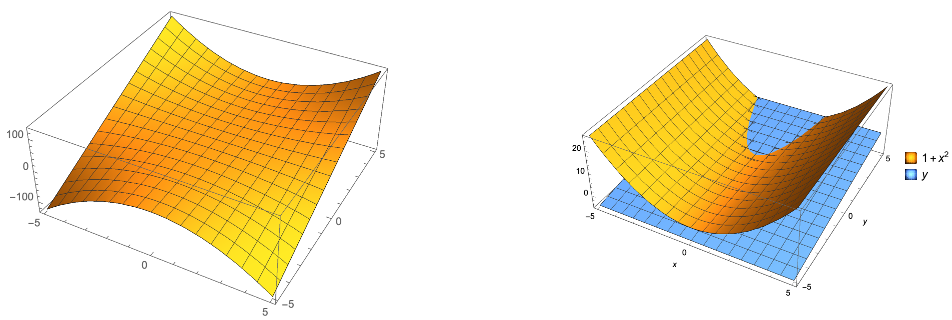

a. The slice $f(x, y) = x^2 y + y$.

b. It's components $g_1(x) = x^2 + 1$ and $g_2(y) = y$

Figure 2.1 A visualization of a slice into its components.

We can now define the slice rank.

Definition 2.2.6. *Let T be a k-tensor. The* **slice rank** *of T is the smallest non-negative integer r such that T can be written as the sum of r slices. For notational purposes, we denote $r = \operatorname{srk}(T)$.*

Let us compute the slice rank of a simple tensor.

Example 5. *Let us consider the 3-tensor $T : X^3 \to \mathbb{F}$ where $T(x, y, z) = xy + yz$. We see that*

$$T(x, y, z) = F(y)\, G(x, z),$$

where $F : X \to \mathbb{F}$ is a 1-tensor with $F(y) = y$ and $G : X^2 \to \mathbb{F}$ is a 2-tensor with $G(x, z) = x + z$. So, $\operatorname{srk}(T) \leq 1$. Also, clearly $\operatorname{srk}(T) > 0$ as T is not the zero tensor. So, T is a slice and $\operatorname{srk}(T) = 1$.

Example 6. *Let us consider the 3-tensor $M : X^3 \to \mathbb{F}$ where $M(x, y, z) = xy + yz + xz$. We note that M is not a slice because we cannot factor out a variable* *from this sum. So,* $\operatorname{srk}(M) \geq 2$. *However, we can write*

$$M(x, y, z) = x(y + z) + yz.$$

So, $\operatorname{srk}(M) = 2$.

The slice rank may seem to be a completely new idea, but in fact, it is more familiar than it may seem. Let us recall some linear algebra. In a matrix M of rank 1, each row is a constant multiple of some vector v^T, so we can write $M = uv^T$ for some vector u. The rank of a matrix M is just the smallest number r such that M is the sum of r rank 1 matrices. Rewriting in the language of 2-tensors, we note that a slice S is just a matrix of rank at most 1. So, a slice S can be written as the product of the two vectors u and v, if we write them as 1-tensors. This argument also works for the tensor rank, as when $k = 2$, the tensor rank is the minimal sum of the product of two 1-tensors. From this, we see the slice rank *and* the tensor rank of a 2-tensor is just the usual matrix rank. From this, we see that these two notions of rank are just different ways of generalizing the matrix rank to larger k.

We see that the tensor rank and the slice rank agree when $k = 2$. Let us now compare the slice rank and the tensor rank in a more general sense.

Proposition 2.2.7. *Let T be a k-tensor. Then* $\operatorname{srk}(T) \leq \operatorname{trk}(T)$. *If $k = 1$ or $k = 2$, then* $\operatorname{srk}(T) = \operatorname{trk}(T)$.

Proof. Let $\operatorname{trk}(T) = m$. Then, we see that

$$T = \sum_{i=1}^{m} f_i^1(x_1) f_i^2(x_2) \dots f_i^k(x_k) = \sum_{i=1}^{m} f_i^1(x_1) g_i(x_2, \dots, x_k),$$

for some $(k-1)$-tensor g_i, which is the product of the $k-1$ 1-tensors $f_i^2, \dots, f_i^k$. We see that each term on the right is a slice. From this, we have $\operatorname{srk}(T) \leq \operatorname{trk}(T)$.

Note that if $k = 1$, then $\operatorname{srk}(T) = \operatorname{trk}(T) = 1$. If $k = 2$, then $T = \sum_{i=1}^{m} f_i(x_1) g_i(x_2)$ so, $\operatorname{srk}(T) = m = \operatorname{trk}(T)$. □

We now have a relation between the slice rank and the tensor rank. However, our end goal is to relate the slice rank and our set X. The following proposition captures an important truth of the slice rank.

Proposition 2.2.8. *Let $T : X^k \to \mathbb{F}$ be a k-tensor. Then,* $\operatorname{srk}(T) \leq |X|$.

Proof. Every k tensor $T : X^k \to \mathbb{F}$ can be written as

$$T(x_1, \ldots, x_k) = \sum_{x \in X} \delta_{x_1 x} T(x, x_2, \ldots, x_k),$$

where δ_{ij} denotes the Kronecker delta function. We see each term on the right hand side is indeed a slice. □

It is interesting to note that while we always have this upper bound for the slice rank of a k-tensor T, the tensor rank can be much larger. In fact, Lickteig (1985) proved that the tensor rank of a typical[6] k-tensor T is $\approx |X|^2/3$ if $|X| \neq 3$ and our field $\mathbb{F}$ is algebraically closed.

Proposition 2.2.8 gives us an avenue to lower bound $|X|$, by providing a lower bound for $\text{srk}(T)$ where $T : X^k \to \mathbb{F}$ is a k-tensor. We have no restrictions on our k-tensor here, any T will work. However, we cannot provide an upper bound with this method. Although we have $\text{srk}(T) \leq \text{trk}(T)$ and $\text{srk}(T) \leq |X|$ for any k-tensor $T : X^k \to \mathbb{F}$, we do not have always have a relation between $\text{trk}(T)$ and $|X|$. Naturally, we can ask when is equality for either of these values achieved, and we provide precisely that. This result, due to Tao, is the crux of our thesis.

Lemma 2.2.9 (Slice Rank Lemma, Tao (2016b))**.** *Let $T : X^k \to \mathbb{F}$ be a diagonal k-tensor with non-zero diagonal entries. Then* $\text{srk}(T) = |X|$.

Proof. We prove this lemma by induction. When $k = 2$, we note that the slice rank is equivalent to the matrix rank. So, we can recall from linear algebra that the rank of an $|X| \times |X|$ diagonal matrix with non-zero entries is $|X|$.

In order to simplify notation, we prove this lemma for $k = 3$, as the proof of the general case is similar.

Let $T : X^3 \to \mathbb{F}$ be a diagonal 3-tensor with non-zero diagonal entries. By Proposition 2.2.8, we see that $\text{srk}(T) \leq |X|$. Suppose toward a contradiction that $\gamma = \text{srk}(T) < |X|$. That is

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

[6] The word typical here is somewhat of a misnomer. This is to say the set of k-tensors with this tensor rank is non-empty and Zariski-open. We do not go into much detail as to what a Zariski-open set is, but the important property is that this set is a dense set in the set of all k-tensors.

where $0 \le \alpha \le \beta \le \gamma < |X|$ are integers, each f_i is a 1-tensor, and each G_i is a 2-tensor on X.

Consider the subspace V orthogonal to the vectors $f_1, \dots, f_\alpha$. That is,

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

Let $v \in V$ be a vector with the largest support $S_v = \{x \in X \mid v(x) \neq 0\}$. We claim that $|S_v| \ge \dim V$.

Suppose that $|S_v| < \dim V$, and consider the subspace $W = \{w : X \to \mathbb{F} \mid w(x) = 0 \text{ for all } x \in S_v\}$. From this, we see that if we were to represent each w as a vector, it can only be non-zero for all $x \notin S_v$. This gives us

$$\dim W = |X| - |S_v| > |X| - \dim V.$$

Therefore, there exists a non-zero vector $w \in V$ such that $w(x) = 0$ for all $x \in S_v$. So, we see that $|S_{v+w}| \ge |S_v|$, as for every $x \in S_v$, $(v + w)(x) = v(x) + w(x) = v(x) \neq 0$. However, as w is a non-zero vector, we see there exists some $x \notin S_v$ such that $(v + w)(x) = v(x) + w(x) = w(x) \neq 0$. Thus, $|S_{v+w}| > |S_v|$, contradicting the maximality of S_v.

So, we see that

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

Now, consider the 2-tensor $Q : S_v^2 \to \mathbb{F}$ defined by

$$Q(y, z) = \sum_{x \in X} v(x) T(x, y, z).$$

By substituting T as the sum of γ slices, we see that

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

However, note that, by definition, for $1 \leq i \leq \alpha$, $\sum_{x \in X} v(x) f_i(x) = 0$, and thus,

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

for some 1-tensors g_i. So, we see that $\mathrm{srk}(Q) \leq \gamma - \alpha$.

Let us now try examining $\mathrm{srk}(Q)$ via direct substitution. Since T is diagonal, we have that $Q(y,z) = \sum_{x \in X} v(x) T(x,y,z) = 0$ whenever $y \neq z$. So, for $y \in S_v$, we have $Q(y,y) = \sum_{x \in X} v(x) T(x,y,y)$. Again, the diagonality of T implies that for all $x \neq y$, we see that $v(x)T(x,y,y) = 0$, so, $Q(y,y) = v(y)T(y,y,y) \neq 0$. Thus, Q is a diagonal matrix with non-zero diagonal entries, so $\mathrm{srk}(Q) = |S_v| > \gamma - \alpha$, which is a contradiction. □

This result is what allows us to solve problems in extremal combinatorics. Recalling from the introduction, a problem, in which we consider restrictions on k-tuple of sets, we may solve it in the following manner:

1. Suppose we find the largest such set family $\mathcal{S}$ that satisfies such a restriction.
2. Create a diagonal k-tensor $T : \mathcal{S}^k \to \mathbb{F}$, where $\mathbb{F}$ is a field, with non-zero diagonal entries that encapsulates the restrictions of the k-tuple of sets.
3. By the slice rank lemma, we have that $\mathrm{srk}(T) = |\mathcal{S}|$.
4. We can now upper bound the slice rank of this k-tensor to achieve an upper bound[7] on $|\mathcal{S}|$.

As we will see in the next chapter, it will require great ingenuity to create diagonal k-tensors which encapsulate the restrictions of the problem while simultaneously having "small" slice rank. It will also require some machinery to bound $\mathrm{srk}(T)$, but we will go over those steps in due time.

[7]We can suppose we find the smallest set family $\mathcal{T}$ that satisfies the restrictions and then find a lower bound the slice rank of the k-tensor to achieve a lower bound on $|\mathcal{T}|$.

Chapter 3

Slice Rank Polynomial Method

This chapter is focused on the slice rank polynomial method itself. We begin by providing some applications due to other mathematicians, then with some examples in which the slice rank polynomial method does not provide any meaningful improvement of previous results. We end with a discussion on how the slice rank grows, and what we may expect the slice rank of a given k-tensor to be.

3.1 Applications of the Slice Rank Polynomial Method

Now that we have introduced the basics of the slice rank polynomial method, let us put it to use. These next two examples offer meaningful insight as to the power of the method, as well as where the difficulties lie. In fact, the second example can be regarded as the birthplace of the slice rank polynomial method. We go into more detail on the short history in subsection 3.1.2. Without further ado, let us get our hands dirty.

3.1.1 Sunflowers

This first application features the Erdős-Rado Sunflower Conjecture, which asks about the maximum number of sets a k-uniform family can contain without containing a sunflower with ℓ petals, that is, an ℓ-sunflower. To this end, let us define what a sunflower is *mathematically*.

Definition 3.1.1. *Let a family $\mathcal{F} \subseteq 2^{[n]}$ be given.*

1. *We say that $\mathcal{F}$ is k-**uniform** if every set $S \in \mathcal{F}$ satisfies $|S| = k$.*
2. *We say that $\mathcal{F}$ contains an ℓ-**sunflower** if there exists subsets $\{S_1, S_2, \ldots, S_\ell\}$ of $[n]$ such that for any $i \neq j$, $S_i \cap S_j = S$ for some set $S \in 2^{[n]}$. That is, the pairwise intersection of ℓ sets in $\mathcal{F}$ is the same set S.*

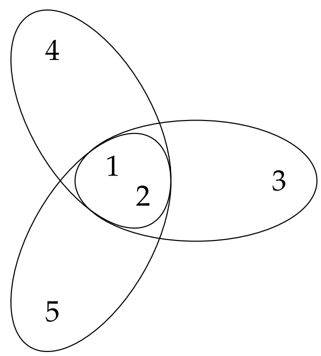

Figure 3.1 Visualization of the 3-sunflower $\{\{1,2,3\},\{1,2,4\},\{1,2,5\}\}$

While these objects are easy to define, the conjecture is very much open, and has been since its formulation in 1960.[1] In order to get a handle on this problem, Erdős and Szemerédi suggested to obtain a non-trivial upper bound involving the number n instead of the uniformity. The best they could come up with was that the maximum size of a 3-sunflower-free family on $[n]$ was at most $2^n/e^{c\sqrt{n}}$, but unfortunately, this was not much better than the trivial bound of 2^n. However, Erdős and Szemerédi did believe an exponential factor improvement should be true.

Conjecture (Erdös and Szemerédi (1978))**.** *For every $\ell \in \mathbb{N}_+$, there exists a constant $c_\ell < 2$ such that for 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 implied their conjecture. However, similarly to the Erdős-Rado Conjecture, this conjecture is hard to prove, and was only recently proved for $\ell = 3$. The problem is still unsolved for $\ell > 3$.

[1] Excitingly, there has been recently progress due to Alweiss, Lovett, Wu and Zhang (2019) and Rao (2020), however the methods used in those arguments are not of interest for this thesis.

Theorem 3.1.2 (Naslund and Sawin (2017)). *Let $c = \frac{3}{\sqrt[3]{4}} < 1.89$. Then every family $\mathcal{F} \subseteq 2^{[n]}$ of size at least c^n contains a 3-sunflower.*

Proof. Let $\mathcal{S} \subseteq 2^{[n]}$ be a 3-sunflower-free family. In order to create our 3-tensor, we wish to find a characterization of when 3 distinct sets do not form a sunflower. We see that three distinct sets, A, B, C, do not form a sunflower if and only if an element of $[n]$ is shared with exactly two of A, B, or C. Less colloquially, 3 distinct sets $A, B, C \in \mathcal{S}$ do not form a 3-sunflower if and only if there exists $i \in [n]$ such that $(v_A + v_B + v_C)_i = 2$. We are now closer to defining our diagonal 3-tensor. We define the function

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

where $x, y, z \in \mathbb{F}^{[n]}$. From this definition, we see that $T(v_A, v_B, v_C) = 0$ for any three distinct members $A, B, C \in \mathcal{S}$. However, we still need to take care of the cases where two of the three sets are equal in order to make T diagonal. To this end, we partition $\mathcal{S}$ into antichains, classifying the elements according to their size, say $\mathcal{S}_j = \{S \in \mathcal{S} \mid |S| = j\}$. If we now consider $A, B, C \in \mathcal{S}_j$, where we have $A = B \neq C$, then, we necessarily have $A, B \nsubseteq C$. This necessarily implies that there exists an i such that $(v_A + v_B + v_C)_i = 2$. From this, we see that T is a diagonal 3-tensor on $\mathcal{S}_j$.

We must finally ensure that our diagonal entries are non-zero. However, if we work in a field with characteristic not equal to 2, say $\mathbb{R}$, we see that $T(v_A, v_A, v_A) = (-1)^{|A|} 2^{n-|A|}$, which is non-zero. We can now apply Lemma 2.2.9 to our 3-tensor T on $\mathcal{S}_j$, we obtain that $|\mathcal{S}_j| = \text{srk}(T)$. What is left is bounding $\text{srk}(T)$.

We note that T is 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, we obtain

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

where $w_I = \prod_{i \in I} w_i$. To break this up into slices, we classify the terms according to the smallest of the sizes $|I|, |J|, |K|$, with ties broken arbitrarily.

Since I, J, and K are disjoint, the minimum size is at most $n/3$. So, we write

$$\begin{aligned} T(x,y,z) &= \sum_{\substack{I\sqcup J\sqcup K\sqcup L=[n]\\ |I|\le n/3}} 2^{|L|} x_I y_J z_K + \sum_{\substack{I\sqcup J\sqcup K\sqcup L=[n]\\ |J|\le n/3}} 2^{|L|} x_I y_J z_K \\ &\qquad + \sum_{\substack{I\sqcup J\sqcup K\sqcup L=[n]\\ |K|\le n/3}} 2^{|L|} x_I y_J z_K \\ &= \sum_{\substack{I\subseteq[n]\\ |I|\le n/3}} x_I f_I(y,z) + \sum_{\substack{J\subseteq[n]\\ |J|\le n/3}} y_J g_J(x,z) + \sum_{\substack{K\subseteq[n]\\ |K|\le n/3}} z_K h_K(x,y), \end{aligned}$$

for some 2-tensors f_I, g_J, h_K. We see that each of these terms are slices, so $\operatorname{srk}(T) \le 3\sum_{i=0}^{n/3}\binom{n}{i}$. Thus,

$$|\mathcal{S}| = \sum_{j=0}^{n} |\mathcal{S}_j| \le 3\,(n+1)\sum_{i=0}^{n/3}\binom{n}{i}.$$

We analyze this using probability theory.[2] We need only consider the sum of binomials, since the other terms grow linearly, and this term will grow exponentially.

Let $X \sim \operatorname{Bin}(n, \frac{1}{2})$ be a binomial distribution with n trials, where the probability of success is $1/2$. Then, for $0 < \lambda < 1$, we have

$$\sum_{k=0}^{\lambda n} \frac{\binom{n}{k}}{2^n} = \sum_{k=0}^{\lambda n} \underbrace{\binom{n}{k}\left(\frac{1}{2}\right)^k \left(\frac{1}{2}\right)^{n-k}}_{\mathbb{P}[X=k]} = \mathbb{P}[X \le \lambda n],$$

where, in general $\mathbb{P}[X = m]$ is the probability that X has m successes in n trials. We then see, in our sum,

$$\sum_{i=0}^{n/3}\binom{n}{i} = 2^n\,\mathbb{P}\left[X \le \frac{1}{3}\,n\right].$$

So, we just need to analyze $\mathbb{P}\left[X \le \frac{1}{3}\,n\right]$. However, note that X can be written as the sum of indicator variables Y_i, where $Y_i \sim \operatorname{Bin}(1, \frac{1}{2})$, is a Bernoulli

[2] We do not provide an introduction to probability theory in this thesis. For the curious reader, see texts like Ross (2014) for an introduction.

distribution. Then, by the theory of large deviations,

$$\mathbb{P}\left[X \leq \frac{1}{3}n\right] = \underbrace{\mathbb{P}\left[\frac{X}{n} \leq \frac{1}{3}\right]}_{\text{average of the } Y_i's} \approx e^{-n\, I\left(\frac{1}{3}\right)},$$

where, for a binomial distribution,

$$I\left(\frac{1}{3}\right) = \frac{1}{3}\ln\left(\frac{1}{3}\right) + \left(1 - \frac{1}{3}\right)\ln\left(1 - \frac{1}{3}\right) + \ln 2.$$

Therefore, we find that

$$\begin{aligned}
\mathbb{P}\left[X \leq \frac{1}{3}n\right] &\approx e^{-n\left(\frac{1}{3}\ln\left(\frac{1}{3}\right)+\left(1-\frac{1}{3}\right)\ln\left(1-\frac{1}{3}\right)+\ln 2\right)} \\
&= e^{-n\,\ln\left(\frac{2\cdot 2^{2/3}}{3}\right)} \\
&= \left(e^{\ln\left(\frac{3}{2\cdot 2^{2/3}}\right)}\right)^n \\
&= \left(\frac{3}{2\cdot 2^{2/3}}\right)^n.
\end{aligned}$$

We are now able to bound the sum of binomials, in turn, giving us a bound of $|\mathcal{S}|$:

$$\sum_{i=0}^{n/3}\binom{n}{i} = 2^n\,\mathbb{P}\left[X \leq \frac{1}{3}n\right] = 2^n\left(\frac{3}{2\cdot 2^{2/3}}\right)^n = \left(\frac{3}{2^{2/3}}\right)^n < 1.89^n.$$

As you can see from this example, the difficulty in the proof was constructing a desirable diagonal k-tensor. However, once the problem is formulated in a manner in which we can define a diagonal k-tensor, relatively simple combinatorial reasoning can be used to bound the slice rank.

3.1.2 SET!

In the family-friendly visual card game SET [3], one plays with a deck of 81 unique cards, each having 4 features with 3 possibilities each. The features

[3] To play online, visit https://www.setgame.com/set/puzzle.

are the shape (diamond, squiggle, or oval), number of shapes (one, two, or three), color (red, green, or purple), and shading (solid, striped, or open). To make a **set**, you must have a combination of three cards, where each feature is either (1) all the same, or (2) all different.

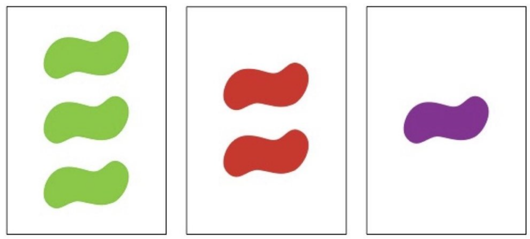

Figure 3.2 Example of a set.

Normally, one plays by laying out twelve cards, and the players try to find a set, and if that happens, they call "set!" and then the three cards they took are replaced. If there is no set from the twelve cards, three more are added, and the process is repeated until a set is found. One may ask, how many cards need to be on the table until a set is guaranteed. As mathematicians are playful by nature, you are probably not surprised to hear that this has been studied, and in fact was proven before the invention of the game by Pellegrino (1970), who showed that number is 21.

Thinking about the cards in SET mathematically, we can encode them as vectors in $\mathbb{F}_3^4$, as each card has 4 features, each taking three possible values. One may ask: what does it mean to have a SET in this sense? Taking three elements $a, b, c \in \mathbb{F}_3$, sum to 0 if they are all distinct, or all the same. However, if $a \neq b = c$, then $a + b + c = a + 2b = a - b \neq 0$ in $\mathbb{F}_3$. So, we can see that three distinct cards corresponding to the vectors $x, y, z \in \mathbb{F}_3^4$ form a set if and only if $x + y + z = 0$. So, if we had a SET-free set $S \subseteq \mathbb{F}_3^4$, we see that there would not exist three distinct elements $a, b, c \in S$ such that $a + b + c = 0$, where 0 denotes the zero vector. Before we continue with our analysis, a quick definition.

Definition 3.1.3. *Let a field $\mathbb{F}$ be given. Then, a 3-AP-free set S is a set such that no three distinct elements in S form an arithmetic progression.*

What is the relation of a 3-AP-free set and a SET-free set? Well, we can see the relation with the following lemma.

Lemma 3.1.4. *For any $a, b, c \in \mathbb{F}_3$, $a + b + c = 0$ if and only if a, b, c form an* *arithmetic progression.*

Proof. ($\longrightarrow$) Suppose that $a + b + c = 0$. From the above analysis, we see that either a, b, c are all distinct, or $a = b = c$. If $a = b = c$, then we have a trivial arithmetic progression. If a, b, c are all distinct, then we have $a = 0$, $b = 1$, and $c = 2$, as there are only 3 elements in $\mathbb{F}_3$. We see this is also an arithmetic progression.

($\longleftarrow$) Suppose that a, b, c form an arithmetic progression. Then, we claim that $b = a + n$ and $c = a + 2n$ with $n = 0$ or $n = 1$. However, then we have $a + b + c = a + (a + n) + (a + 2n) = 3a + 3n = 0 \in \mathbb{F}_3$. ◻

With this lemma, we see that a SET-free set of vectors in $\mathbb{F}_3^4$ is just a 3-AP-free set of vectors. We can generalize this to higher dimensions. It we were to add more features to our cards, say we have d features, we can still note that a SET-free set of vectors in $\mathbb{F}_3^d$ is just a 3-AP-free set of vectors. We will let $r_3(\mathbb{F}_3^d)$ denote the largest number such there is a 3-AP-free set of that size in $\mathbb{F}_3^d$. So, in order to guarantee a set, if we have d features, we need to lay out $r_3(\mathbb{F}_3^d) + 1$ cards on the table. This problem was recently solved by Ellenberg and Gijswijt, using a method that is equivalent to the slice rank polynomial method. Their motivation, perhaps unsatisfyingly, originated in finite geometry, with something known as the capset problem.

Theorem 3.1.5 (Ellenberg and Gijswijt (2017))**.** *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. From our discussion above, we see that this is equivalent to that for every distinct $a, b, c \in S$, $a + b + c \neq 0$. So, we define the following function

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

We see that $T(a, b, c) = 0$ for any distinct $a, b, c \in S$. We also see that the same also holds for any $a \neq b = c$, as $a + 2\beta \neq 0$ for any $a \neq \beta$ in $\mathbb{F}_3$. Thus, T is a diagonal 3-tensor on S, and since $T(a, a, a) = 2^d \neq 0$ for any $a \in \mathbb{F}_3^d$, Lemma 2.2.9 applies, and we have $|S| = \text{srk}(T)$.

We now bound $\text{srk}(T)$. In doing so, we note that T is a polynomial in the $3d$ variables $x_1, \ldots, x_d, y_1, \ldots, y_d, z_1, \ldots, z_d$. The total degree is $2d$ and the degree in each variable is 2. So, we have

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

where $c_{\alpha,\beta,\gamma}$ is some constant depending on α,β,γ. We can now classify the terms according to which of x, y, or z has the smallest total degree. As the overall total degree is $2d$, we see that at least one of $\sum\alpha_i$, $\sum\beta_i$ and $\sum\gamma_i$ is at most $2d/3$. From this, we see

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

for some 2-tensors $f_\alpha, g_\beta, h_\gamma$. We note that as all these terms are slices, we see that the slice rank of T 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 $2d/3$. Letting a,b,c represent the number of $0,1,2$ coordinates of α respectively, we see that this is equal to

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

To estimate this multinomial coefficient, consider the expression

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

which is true for every real x. To make the terms of interest, $b + 2c \leq 2d/3$, the dominating ones, we first divide through by $x^{2d/3}$ and then estimate when $0 < x < 1$. To this end, we obtain

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

In the first inequality, we used that $x > 0$, and in the second, we used that $x < 1$. To obtain the best upper bound, we minimize the function $f(x) = x^{-2/3} + x^{1/3} + x^{4/3}$ on the interval $0 < x < 1$. We see that $f'(x) = (4x^2 + x - 2)/3x^{5/3}$. From this, we note that $f'(x) = 0$ when $4x^2 + x - 2 = 0$, which is if $x = -1/8 \pm \sqrt{33}/8$. However, we desire $0 < x < 1$, so this implies we have $x = (\sqrt{33} - 1)/8$, which gives us $f((\sqrt{33} - 1)/8) \approx 2.755 < 2.76$. This gives us the bound of 2.76^d.

3.2 What can go Awry?

Given the previous examples, the slice rank polynomial method is promising; however, like everything, it is not perfect. In this subsection, we provide two examples in which the slice rank polynomial method, at least in the manner used, provide worse bounds than the trivial bounds. We hope this illuminates some of the possible drawbacks and difficulties that may arise when solving an extremal problem using the slice rank polynomial method.

3.2.1 Sunflower Generalization

As we saw in Section 3.1.1, Naslund and Sawin were able to prove Erdös and Szemerédi's conjecture in the special case of a 3-sunflower. In this subsection, we (perhaps naïvely) attempt to generalize their proof for all k-sunflowers, in hopes to prove the Erdös and Szemerédi conjecture. In the end, the bounds we attain are not better than the trivial bounds, even for

$k = 4$, however, let us explore the failed attempt and learn more about the slice rank polynomial method.

Proof Attempt of the Erdös and Szemerédi's conjecture. Let $\mathcal{S} \subseteq 2^{[n]}$ be a family with no k-sunflowers. In order to create our k-tensor, we wish to find a characterization of when k distinct sets do not form a sunflower. We see that k distinct sets, $A_1, A_2, \ldots, A_k$, do not form a sunflower if and only if an element of $[n]$ is shared with $2, 3, ..., k-1$ of $A_1, A_2, \ldots, A_k$. Reframing this, k distinct sets $A_1, A_2, \ldots, A_k \in \mathcal{S}$ do not form a k-sunflower if and only if there exists $i \in [n]$ such that $(v_{A_1} + v_{A_2} + \cdots + v_{A_k})_i \in \{2, 3, \ldots, k-1\}$. We are now closer to defining our diagonal k-tensor. We define the function

$$T_k(x_1, x_2, \ldots, x_k) = \prod_{i=1}^{n} \prod_{j=2}^{k-1} (j - (x_1 + x_2 + \cdots + x_k)_i),$$

where $x_1, x_2, \ldots, x_k \in \mathbb{F}^{[n]}$. From this definition, our tensor equals zero for any k distinct members $A_1, A_2, \ldots, A_k \in \mathcal{S}$. However, we still need to take care of the cases where some of these sets are equal in order to make T_k diagonal. To this end, we partition $\mathcal{S}$ into antichains, classifying the elements according to their size, say $\mathcal{S}_j = \{S \in \mathcal{S} \mid |S| = j\}$. Now, in any case that is not $A_1 = A_2 = \cdots = A_k$ or $A_1, A_2, \ldots, A_k$ all distinct, there must exist an i such that $(v_{A_1} + v_{A_2} + \cdots + v_{A_k})_i \in \{2, 3, \ldots, k-1\}$. From this, we see that T_k is a diagonal k-tensor on $\mathcal{S}_j$.

We must finally ensure that our diagonal entries are non-zero. However, working in $\mathbb{R}$, ensures this as $(v_A + v_A + \cdots + v_A)_i \in \{0, k\}$ for all i. We can now apply Lemma 2.2.9 to our k-tensor T_k on $\mathcal{S}_j$, we obtain that $|\mathcal{S}_j| = \mathrm{srk}(T_k)$. What is left is bounding $\mathrm{srk}(T_k)$.

We note that T_k is polynomial of total degree $(k-2)n$ in the kn variables $x_{1,1}, \ldots, x_{1,n}, x_{2,1}, \ldots, x_{2,n}, \ldots, x_{k,1}, \ldots, x_{k,n}$. Expanding, we obtain

$$T_k(x_1, \ldots, x_k) = \sum_{\alpha_1, \ldots, \alpha_k \in \{0,1,\ldots,k-2\}^n} c_{\alpha_1,\ldots,\alpha_k} \left(\prod_{i=1}^{d} x_{1,i}^{\alpha_{1,i}} \right) \cdots \left(\prod_{i=1}^{d} x_{k,i}^{\alpha_{k,i}} \right),$$

where $c_{\alpha_1,\ldots,\alpha_k}$ is some constant depending on $\alpha_1, \ldots, \alpha_k$. We can now classify the terms according to which of $x_1, \ldots, x_k$ has the smallest total degree. As the overall total degree is $(k-2)n$, we see that at least one of

$\sum \alpha_{1,i}, \ldots, \sum \alpha_{k,i}$ is at most $(k-2)n/k$. From this, we see

$$
\begin{aligned}
T_k(x_1, \ldots, x_k) = &\sum_{\substack{\alpha_1,\ldots,\alpha_k \in \{0,1,\ldots,k-2\}^n \\ \sum \alpha_{1,i} \le (k-2)n/k}} c_{\alpha_1,\ldots,\alpha_k} \left(\prod_{i=1}^{d} x_{1,i}^{\alpha_{1,i}} \right) \ldots \left(\prod_{i=1}^{d} x_{k,i}^{\alpha_{k,i}} \right) \\
&+ \ldots + \sum_{\substack{\alpha_1,\ldots,\alpha_k \in \{0,1,\ldots,k-2\}^n \\ \sum \alpha_{k,i} \le (k-2)n/k}} c_{\alpha_1,\ldots,\alpha_k} \left(\prod_{i=1}^{d} x_{1,i}^{\alpha_{1,i}} \right) \ldots \left(\prod_{i=1}^{d} x_{k,i}^{\alpha_{k,i}} \right) \\
= &\sum_{\substack{\alpha_1,\ldots,\alpha_k \in \{0,1,\ldots,k-2\}^n \\ \sum \alpha_{1,i} \le (k-2)n/k}} c_{\alpha_1,\ldots,\alpha_k} \left(\prod_{i=1}^{d} x_{1,i}^{\alpha_{1,i}} \right) f_{\alpha_1}(x_2, \ldots, x_k) \\
&+ \ldots + \sum_{\substack{\alpha_1,\ldots,\alpha_k \in \{0,1,\ldots,k-2\}^n \\ \sum \alpha_{k,i} \le (k-2)n/k}} c_{\alpha_1,\ldots,\alpha_k} \left(\prod_{i=1}^{d} x_{k,i}^{\alpha_{k,i}} \right) f_{\alpha_k}(x_1, \ldots, x_{k-1})
\end{aligned}
$$

for some $(k-1)$-tensors $f_{\alpha_1}, f_{\alpha_2}, \ldots, f_{\alpha_k}$. We note that as all these terms are slices, we see that the slice rank of T_k is bounded by k times the number of ways to select a vector $\alpha \in \{0,1,\ldots,k-2\}^n$ such that the sum of its coordinates is at most $(k-2)n/k$. Letting $a_1, a_2, \ldots, a_{k-1}$ represent the number of $0, 1, \ldots, k-2$ coordinates of α, we see that this is equal to

$$k \cdot \sum_{\substack{a_1+a_2+\cdots+a_{k-1}=n \\ a_2+2a_3+\cdots+(k-2)a_{k-1} \le (k-2)n/k}} \frac{n!}{a_1!a_2!\ldots a_{k-1}!}.$$

When $k=3$, we obtain

$$3 \cdot \sum_{\substack{a_1+a_2=n \\ a_2 \le n/3}} \frac{n!}{a_1!a_2!} = 3 \cdot \sum_{a_2 \le n/3} \frac{n!}{(n-a_2)!a_2!} = 3 \cdot \sum_{k \le n/3} \binom{n}{k}.$$

which is the same bound as Naslund and Sawin in their proof of Theorem 3.1.2. When $k=4$, we obtain

$$4 \cdot \sum_{\substack{a+b+c=n \\ b+2c \le n/2}} \frac{n!}{a!b!c!}.$$

We also notice that for $k > 4$,

$$k \cdot \sum_{\substack{a_1+a_2+\cdots+a_{k-1}=n \\ a_2+2a_3+\cdots+(k-2)a_{k-1} \le (k-2)n/k}} \frac{n!}{a_1!a_2!\ldots a_{k-1}!} > 4 \cdot \sum_{\substack{a+b+c=n \\ b+2c \le n/2}} \frac{n!}{a!b!c!},$$

so, we first bound $\mathrm{srk}(T_4)$ using a similar method to Ellenberg and Gijswijt's slice rank analysis in Theorem 3.1.5. Consider the expression

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

which is true for every real x. To make the terms of interest, $b+2c \le n/2$, the dominating ones, we first divide through by $x^{n/2}$ and then estimate when $0 < x < 1$. To this end, we obtain

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

In the first inequality we used that $x > 0$ and in the second we used that $x < 1$. To obtain the best upper bound, we minimize the function $f(x) = x^{-1/2}+x^{1/2}+x^{3/2}$ on the interval $0 < x < 1$. We see that $f'(x) = \frac{3x^2+x-1}{2x^{3/2}}$. From this, we note that $f'(x) = 0$ when $3x^2+x-1 = 0$, which is if $x = \frac{-1\pm\sqrt{13}}{6}$. However, we desire $0 < x < 1$, so this implies we have $x = \frac{-1+\sqrt{13}}{6}$, which gives us the maximum of $f\left(\frac{-1+\sqrt{13}}{6}\right) = 2.4626... < 2.47$. So, we have $\mathrm{srk}(T_4) < 2.47^n$. □

Again, note that for $k > 4$, we have $\mathrm{srk}(T_k) > \mathrm{srk}(T_4)$, however even for $k = 4$, the bound we obtain is not as strong as the trivial bound of 2^n. There are various possible reasons for this: the approximations for $\mathrm{srk}(T_4)$ were too loose, the decomposition of T_4 into slices is not optimal, or T_4 itself has too large of a slice rank. In turn, there are various possible remedies assuming we still desire to solve this problem using the slice rank polynomial method: improve the approximations for $\mathrm{srk}(T_4)$, attempt to analyze the $\mathrm{srk}(T_4)$ via another decomposition of T_4 into slices, define a different tensor T_4 entirely using a similar framework, or attempt the problem using a different framework for k-sunflower sets and a characterization for when we do not have a k-sunflower.

3.2.2 Coloring $\mathbb{R}^n$

In this section we revisit an application of the slice rank polynomial method due to Naslund, but this time in the context of coloring problems in Euclidean spaces. We begin by introducing the nature of the problem, and some important objects.

Definition 3.2.1. *A configuration of points $S \subset \mathbb{R}^k$ is said to be* ***exponentially Ramsey*** *if, in order to guarantee no monochromatic copies of S in any coloring of $\mathbb{R}^n$, we need exponentially many colors as a function of n.*

Definition 3.2.2. *A* ***k-simplex*** *is a k-dimensional polytope which is the convex hull of its $k+1$ vertices. More intuitively, a k-simplex is the generalization of a triangle ($k = 2$) to k-dimensional space. We say that a k-simplex is* ***regular*** *if its side lengths are the same. We illustrate some examples.*

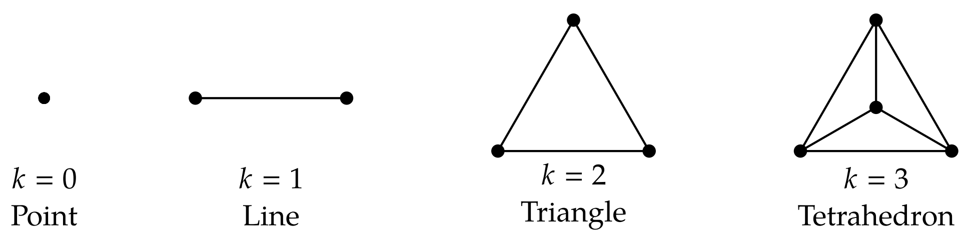

Figure 3.3 First four regular k-simplexes

Why are we so interested in exponentially Ramsey sets and k-simplexes? Frankl and Rödl (1987) proved that every k-simplex is exponentially Ramsey. Specifically, for any k, there exists $\epsilon_k > 0$ such that any coloring of $\mathbb{R}^n$ with less than $(1 + \epsilon_k + o(1))^n$ colors contains a monochromatic regular k-simplex with side-length 1.

Let us examine the case $k = 3$ and let $\chi_\Delta(\mathbb{R}^n)$ denote the minimum number of colors needed to color $\mathbb{R}^n$ so that it does not contain a monochromatic equilateral triangle of side lengths 1. Currently, the best lower bound for ϵ_3 is due to Sagdeev (2018), with $\epsilon_3 = 0.00085...$

In 2020, Naslund attempted to improve this lower bound, using the slice rank polynomial method. Unfortunately, we believe there was an error in the proof that removes the novelty of the result. We state the original theorem from Naslund, and show where we believe the error to be.

Theorem 3.2.3 (Naslund (2020a)). *We have that*

$$\chi_\Delta(\mathbb{R}^n) > (1 + c + o(1))^n,$$

where $c = 0.01446\ldots$

The modified version of Theorem 3.2.3 should instead by stated as follows.

Theorem 3.2.4. *We have that*

$$\chi_\Delta(\mathbb{R}^n) > (1 + c + o(1))^n,$$

where $c = -0.141\ldots$

In order to prove this modified theorem, we need a proposition.

Proposition 3.2.5. *For* $k \leq \frac{n}{2}$, *let* $S \subset \{0,1\}^n$ *be the set of elements with exactly* k *ones, and let* p *be the smallest odd prime such that* $p > \frac{k}{4}$. *Suppose that* $A \subset S$ *does not contain* x, y, z *with*

$$\left\|x - y\right\|_2 = \left\|y - z\right\|_2 = \|z - x\|_2 = \sqrt{2p}.$$

Let $\epsilon_0 = n^{0.525}$ *denote an error term. Then for sufficiently large* n

$$|A| \leq 3 \cdot \min_{0<t<1} \frac{(1 + t + t^2 + t^3)^n}{t^{\frac{n}{3} + \frac{k}{6} + \epsilon_0}}.$$

Proof. The bounds given by Baker et al. (2001) for the largest prime gap imply that for sufficiently large n

$$p < \frac{k}{4} + \epsilon_0.$$

For $x, y, z \in S$ consider the polynomial

$$F\colon S \times S \times S \longrightarrow \mathbb{F}_p$$

defined by

$$F(x, y, z) = \prod_{i=1}^{n} \left(x_i + y_i + z_i - 1\right).$$

If x, y, z satisfy $F(x, y, z) \neq 0$, that is if there is no i such that $x_i + y_i + z_i = 1$, then we must have $\|x - y\|_2 = \|y - z\|_2 = \|z - x\|_2$, and so they form an equilateral triangle. To see this, we do not provide a full proof, but rather a simple explanation.

For each $j \in \{0,1,2,3\}$ let $a_j = \#\left\{i : x_i + y_i + z_i = j\right\}$. Note that $a_1 = 0$ since $F(x, y, z) \neq 0$. Since each vector has exactly k 1's, if we have $x_i + y_i = 2$ with $z_i = 0$ for some i, we must have $x_j + z_j = 2$ with $y_j = 0$ and $y_k + z_k = 2$ with

$x_k = 0$ for some distinct i, j, k, because there are no such i such that $x_i = 1$ but $y_i = z_i = 0$.

When considering the difference of the vectors, these are the only such i which contribute to the distance. Since these i occur the same number of times for each vector x, y, z, we must have that these distances are the same.

Furthermore, if $F(x, y, z) \neq 0$, then we can upper bound the distance

$$\|x - y\|_2^2 < 2p.$$

Since there are n coordinates, and $3k$ total entries equal to 1, we have that

$$a_0 + a_2 + a_3 = n \quad \text{and} \quad 2 \cdot a_2 + 3 \cdot a_3 = 3k.$$

Subtracting 3 times the first equation from the second, we obtain

$$a_2 = 3n - 3k - 3a_0.$$

The only coordinates that contribute to the distance are counted by a_2, and so

$$\|y - z\|_2^2 + \|z - x\|_2^2 + \|x - y\|_2^2 = 2a_2.$$

Hence

$$\|x - y\|_2^2 = 2n - 2k - 2a_0.$$

The smallest a_0 can be is if $a_3 = 0$ and all $3k$ ones are used by coordinates where the sum is 2. That is, $a_0 \geq n - \frac{3k}{2}$, and hence

$$\frac{\|x - y\|_2^2}{2} \leq \frac{k}{2} < 2p. \tag{3.1}$$

Let $G\colon S \times S \longrightarrow \mathbb{F}_p$ be given by

$$G(x, y) = \left(1 - \left(\frac{\|x - y\|_2^2}{2}\right)^{p-1}\right),$$

and note that $\frac{1}{2}\|x - y\|_2^2$ will always be an integer for $x, y \in S$. If if $x \neq y$ are such that $\frac{1}{2}\|x - y\|_2^2 < 2p$, then $G(x, y) \neq 0$ if and only if $\frac{1}{2}\|x - y\|_2^2 = p$. For $x, y, z \in S$ define

$$H(x, y, z) := F(x, y, z)G(x, y).$$

This function will be non-zero when $x = y = z$, and will be zero whenever x, y, z do not form an equilateral triangle with side length $\sqrt{2p}$. Suppose that $A \subset S$ contains no equilateral triangles of side length $\sqrt{2p}$. Then H restricted to $A \times A \times A$ will be a diagonal tensor with non-zero diagonal elements, and so by Lemma 2.2.9

$$|A| \leq \operatorname{srk}(H).$$

The polynomial H will have degree at most $n + 2p < n + \frac{k}{2} + \epsilon_0$, and we may expand it as a linear combination of monomials of the form

$$\left(x_1^{d_1} \cdots x_n^{d_n}\right)\left(y_1^{e_1} \cdots x_n^{e_n}\right)\left(z_1^{f_1} \cdots z_n^{f_n}\right)$$

where $d_i, e_i \in \{0, 1, 2, 3\}$ and $f_i \in \{0, 1\}$ for each i. This is where we differ from the proof by Naslund. Originally, it was stated that $d_i, e_i, f_i \in \{0, 1\}$, however, if we expand $G(x, y)$, we obtain

$$\begin{aligned} G(x, y) &= \left(1 - \left(\frac{\|x - y\|_2^2}{2}\right)^{p-1}\right) \\ &= \left(1 - \left(\frac{(x_1 - y_1)^2 + (x_2 - y_2)^2 + \cdots + (x_n - y_n)^2}{2}\right)^{p-1}\right). \end{aligned}$$

From this, we see that when multiplying $F(x, y, z)$ and $G(x, y)$, it is possible to obtain the term $x_1^3 f(x, y, z)$ in your monomial. For example, let $n = 4$, $p = 2$, and $k = 2$, then we have

$$\begin{aligned} &F(x, y, z)\, G(x, y) \\ &= \prod_{i=1}^{4}(x_i + y_i + z_i - 1)\left(1 - \left(\frac{(x_1 - y_1)^2 + (x_2 - y_2)^2 + (x_3 - y_3)^2 + (x_4 - y_4)^2}{2}\right)\right). \end{aligned}$$

Instead of attempting to expand this, we find that if we "chose" x_1 from the first term in our product of $F(x, y, z)$, -1 from the remaining three, and multiplied by the non-constant part of $G(x, y)$, we would obtain

$$x_1 \cdot \left(\frac{(x_1 - y_1)^2 + (x_2 - y_2)^2 + (x_3 - y_3)^2 + (x_4 - y_4)^2}{2}\right).$$

In which, when expanding into monomials, and only considering the first term in our expansion, we would obtain

$$x_1 \cdot \frac{(x_1 - y_1)^2}{2} = \frac{x_1^3 - 2x_1^2 y_1 + x_1 y_1^2}{2}.$$

With this example, it is clear that the polynomial in $G(x, y)$ may contribute up to an additional 2 powers of x_i and y_i for any i. Now, using a similar argument as that of Naslund, we find

$$\left(\sum_{i=1}^{n} d_i\right) + \left(\sum_{i=1}^{n} e_i\right) + \left(\sum_{i=1}^{n} f_i\right) \leq n + \frac{k}{2} + \epsilon_0.$$

For each monomial, one of these sums will be at most $\frac{1}{3}(n + \frac{k}{2} + \epsilon_0)$, and hence by always slicing off the lowest degree piece we have

$$\mathrm{srk}(H) \leq 3 \cdot \# \left\{ v \in \{0,1,2,3\}^n : \sum_{i=1}^{n} v_i \leq \frac{n}{3} + \frac{k}{6} + \frac{\epsilon_0}{3} \right\}.$$

For any $0 < t < 1$,

$$\begin{aligned}
\# \left\{ v \in \{0,1,2,3\}^n : \sum_{i=1}^{n} v_i \leq r \right\} &= \sum_{\substack{k_0+k_1+k_2+k_3=n \\ k_1+2k_2+3k_3 \leq r}} \binom{n}{k_0, k_1, k_2, k_3} \\
&\leq t^{-r} \sum_{k_0+k_1+k_2+k_3=n} \binom{n}{k_0, k_1, k_2, k_3} t^{k_1}(t^2)^{k_2}(t^3)^{k_3} \\
&= \frac{(1+t+t^2+t^3)^n}{t^r}
\end{aligned}$$

since the coefficient $t^{k_1+2k_2+3k_3-r}$ will be greater than 1 for $k_1 + 2k_2 + 3k_3 \leq r$. Taking the minimum over t, for $r = \frac{n}{3} + \frac{k}{6} + \frac{\epsilon_0}{3}$, we obtain the stated result. □

Proof of Theorem 3.2.3. Let $S \subset \{0,1\}^n$ be the subset of vectors with exactly k ones, for $k \leq \frac{n}{2}$, and let $A \subset S$ be the largest subset that does not contain an equilateral triangle of side length $\sqrt{2p}$. Then we need at least $\frac{|S|}{|A|}$ sets that do not contain an equilateral triangle of side lengths $\sqrt{2p}$ to cover S. Rescale every point in $\mathbb{R}^n$ by a factor of $\sqrt{2p}$ so that these points are at distance 1. As $\frac{1}{\sqrt{2p}} S \subset \mathbb{R}^n$, it follows that

$$\chi_\Delta(\mathbb{R}^n) \geq \frac{|S|}{|A|},$$

and by Proposition 3.2.5

$$\chi_\Delta(\mathbb{R}^n) \geq \frac{1}{3}\binom{n}{k} \max_{0<t<1} \frac{t^{\frac{n}{3}+\frac{k}{6}+\epsilon_0}}{(1+t+t^2+t^3)^n}.$$

Since this bound holds for any $0 \leq k \leq \frac{n}{2}$, we may take the maximum and write

$$\chi_\Delta(\mathbb{R}^n) \geq \frac{1}{3} \max_{0<t<1} \left[\left(\frac{t^{\frac{1}{3}+\frac{\epsilon_0}{n}}}{1+t+t^2+t^3} \right)^n \max_{0\leq k\leq \frac{n}{2}} \binom{n}{k} t^{\frac{k}{6}} \right].$$

Expanding $(1+x)^n$, we have that for any $0 < x < 1$

$$\frac{(1+x)^n}{n+1} < \max_{0\leq k\leq \frac{n}{2}} \binom{n}{k} x^k < (1+x)^n,$$

and hence

$$\max_{0\leq k\leq \frac{n}{2}} \binom{n}{k} t^{\frac{k}{6}} > \frac{1}{n+1} \left(1+t^{\frac{1}{6}}\right)^n.$$

We must have $t^{\frac{1}{3}} > \frac{1}{2}$, since otherwise the value of the function we are maximizing will be less than 1. Hence t will be bounded away from 0, which implies that $t^{\frac{\epsilon_0}{n}} = 1 + o(1)$. Simplifying the result, we obtain

$$\chi_\Delta(\mathbb{R}^n) > \left(\max_{0<t<1} \frac{t^{\frac{1}{3}}\left(1+t^{\frac{1}{6}}\right)}{1+t+t^2+t^3} + o(1) \right)^n$$

and the desired bound follows by computing the maximum. In fact, we get this maximum is $c = 0.859$. □

Originally, in the proof given by Naslund, we had the result

$$\chi_\Delta(\mathbb{R}^n) > \left(\max_{0<t<1} \frac{t^{\frac{1}{3}}\left(1+t^{\frac{1}{6}}\right)}{1+t} + o(1) \right)^n,$$

which would give us the maximum $c = 1.01446...$ With this result, we would indeed improve on the previous result given by Sagdeev (2018), which was that of 1.00085..., however, this does not seem to be the case with the new proof.

Assuming our modification is correct, it seems as though the slice rank polynomial method is not as precise as one would hope, the more specific method from Sagdeev seem to win out. However, not all hope is lost! It may be the case that our approximations of

$$\frac{|S|}{|A|} = \max_{0 \le k \le \frac{n}{2}} \frac{\binom{n}{k}}{\#\left\{v \in \{0,1,2,3\}^n : \sum_{i=1}^{n} v_i \le \frac{n}{3} + \frac{k}{6} + \frac{\epsilon_0}{3}\right\}},$$

are too loose. Before we set out to improve the approximations, we check computationally via Sage, whether the value above provides us with an exponential function in n, with a base strictly greater than 1. We outline how what our functions computed, followed by their outputs. There are two main functions.

function sizeOfSet(n, k)
  Create a list L of all 3-tuples (v_1, v_2, v_3) where $v_2+2v_3+3v_4 \le \frac{n}{3}+\frac{k}{6}+\frac{\epsilon_0}{3}$, and each $v_i \ge 0$.
  Sum the appropriate multinomial coefficients and store it in the integer `Sum`.
  return the nth root of $\binom{n}{k}/(3 \cdot \texttt{Sum})$.
end function

function maxSizeOfSet(n)
  Loop over all values of $0 \le k \le n/2$.
  return the maximum value of sizeOfSet(n, k).
end function

Running this, for n up to $10,000$, we never achieved a maximum greater than 0.9. We provide a table of values given by the code we ran.

n	value of k	value of $\lvert S\rvert/\lvert A\rvert$
100	47	$(0.822)^{100}$
1000	457	$(0.847)^{1000}$
5000	2258	$(0.854)^{5000}$
10000	4517	$(0.855)^{10000}$

Table 3.1 Explicit values provided by code, for n up to 10,000.

From this, we believe the error not to lie in the approximations. It simply seems to be the case that the slice rank polynomial method is not strong enough to provide a bound better than the optimized methods for this specific problem. Of course, as before, there are many possible reasons for this. It may be due to how we decomposed our tensor in slices or the definition of the tensor itself. However, with this in mind, we turn to the next section which will provide a short but meaningful discussion on what we may expect the slice rank of a tensor to be given its definition.

3.3 Expectations of the Slice Rank

In this section we discuss heuristics for bounding the slice rank. There are three key features of a tensor that govern the slice rank. These are:

1. k itself,
2. the total degree of the k-tensor,
3. and the multiplicity of each variable.

For the sake of our analysis in this chapter, we let the total degree of our k-tensor be m. We also suppose, like in many of the applications above, that the k variables in our k-tensor are elements of $\mathbb{F}^{[n]}$, for some field $\mathbb{F}$ and a fixed integer n. Typically, in these problem, we have the slice rank of our k-tensor in terms of n, something like the form $\text{srk}(T) \leq s^n$, for some s. Finally, suppose you are given the slice rank of a k-tensor T, say $\text{srk}(T) = c^n$. We would like to analyze the slice rank of related k-tensors.

Suppose we are given a k'-tensor, T', that is intimately related to T with $k' > k$. For example, let T be Naslund and Sawin's sunflower 3-tensor

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

and let T' be

$$T'(x, y, z, w) = \prod_{i=1}^{n} (2 - (x_i + y_i + z_i + w_i)).$$

In this case, the total degree of the two tensors are the same, n. The key difference between these two tensors, in the analysis of their slice rank, lies in the bound of each term into slices, which in this case is n/k or n/k' (in the examples above, we saw this bound was usually the total degree over k, and this follows that trend). When analyzing the slice rank for T, we obtained the bound $3(n+1)\sum_{k=0}^{n/3}\binom{n}{k}$, but for T', we would obtain the bound $4(n+1)\sum_{k=0}^{n/4}\binom{n}{k}$. Since we are mainly considering how this bound grows exponentially, the difference between 3 and 4 is negligible, we mainly consider the difference between summing the binomial terms to $n/4$ and summing to $n/3$.

From this, we see that the slice rank (or, at least, the bound we may attain for it) of T' is smaller than that of T. This tells us that if k increases then, in turn, the bound for each slice term would decrease, leading to a decrease of the slice rank. Of course, the notion of generalizing a k-tensor to a k'-tensor is not necessarily a well-defined notion, as there can be many choices that "make sense", but in cases like the one above, we may apply this heuristic to understand what we may expect of the slice rank.

In a similar manner, we may imagine what would happen if the total degree, m, of our k-tensor T increased in our new k-tensor T'. If this is the only change, then using the same analysis as above, the bound of each term into slices would increase, which would lead to an increase in the slice rank.

We now consider the possible multiplicity of each variable. Suppose that in our k-tensor T the possible multiplicity of each variable was at most a, and we are now given a k-tensor T' in which the possible multiplicity of each variable is $a' > a$. Then, there are two aspects to consider. Typically, when the multiplicity of each variable increases, this necessarily implies that the total degree of our tensor increases. For example, if T is Naslund and Sawin's sunflower 3-tensor and we define T' as

$$T'(x,y,z) = \prod_{i=1}^{n}(2-(x_i+y_i+z_i))(3-(x_i+y_i+z_i)),$$

then the multiplicity of each variable went from at most 1 to at most 2, and the total degree increased from n to $2n$. As we saw above, an increase in the total degree necessarily leads to an increase of the slice rank via the increase of the bounds of each term into slices. However, there is also another change when we increase the multiplicity of our variables: the types of multinomial coefficients we consider. If you take the above as an example, the Naslund and Sawin tensor considered binomial coefficients, namely the sum $\sum_{k=0}^{n/3}\binom{n}{k}$, but with our tensor T', we would have to consider trinomial coefficients, namely the sum $\sum_{\substack{k_1+k_2+k_3=n \\ k_2+2k_3\le 2n/3}} \frac{n!}{k_1!k_2!k_3!}$. The bounds we attain for T is $\mathrm{srk}(T) < 1.89^n$, but using the Ellenberg-Gijswijt method (T'

is intimately related to their tensor), we find that we obtain $\mathrm{srk}(T') < 2.76^n$. This is a big increase. Even if the total degree stayed the same and we only considered those terms in which $k_2 + 2k_3 \leq n/3$, we would obtain the bound $\mathrm{srk}(T') < 2.08^n$. From this analysis, we see that an increase in multiplicity of each variable will typically lead to an increase to the slice rank, especially if the total degree happens to increase as well.

Natural questions arise from this kind of analysis. For example, what can we expect to occur to the slice rank if the multiplicity increases, but the total degree of our polynomial (somehow) decreases? Questions of this nature are difficult to answer in full generality, but we hope that this set of heuristics helps the reader understand what they may expect from the slice rank analysis before they embark on it.

Chapter 4

The Partition Rank

This section dedicates itself to an existing generalization of the slice rank, due to Naslund (2020b), called the **partition rank**, as well as a short discussion on possible directions in the generalization of the slice rank method.

4.1 Preliminaries

Recall that we use the notation, where given variables $x_1, \dots, x_n$ and a set $S \subseteq \{1, ..., n\}$ with $S = \{s_1, \dots, s_k\}$, we use the notation x_S to denote the subset of variables $x_{s_1}, \dots, x_{s_k}$. So, for a function f of k variables, we have

$$f(x_S) = f(x_{s_1}, \dots, x_{s_k}).$$

For example, if $S = \{1, 3, 4\}$, then $f(x_S) = f(x_1, x_3, x_4)$. We now provide some important definitions necessary for introducing the partition rank.

Definition 4.1.1. *A* ***partition*** *of* $[n]$ *is a collection* P *of non-empty, pairwise disjoint, subsets of* $[n]$ *such that*

$$\bigcup_{A \in P} A = [n].$$

We say that P *is the* ***trivial partition*** *if it only consists of a single set,* $[n]$*.*

Definition 4.1.2. *Let* X *be a finite set, let* $\mathbb{F}$ *be a field and suppose that we are given a* k*-tensor* $T : X^k \longrightarrow \mathbb{F}$*. If there exists some non-trivial partition* P *such that*

$$T(x_1, \dots, x_n) = \prod_{A \in P} f_A(x_A),$$

for some functions f_A, that T is said to have ***partition rank 1****.*

Another manner of thinking about a partition rank 1 k-tensor T is if the variables can be split into disjoint non-empty sets $S_1, \dots, S_t$, with $t \geq 2$ such that

1. $S_1 \cup \dots \cup S_t = [n]$, and
2. $T(x_1, \dots, x_n) = f_1(x_{S_1}) f_2(x_{S_2}) \dots f_t(x_{S_t})$, for some tensors $f_1, \dots, f_t$.

That is, our k-tensor T will have partition rank 1 if and only if it can be written as

$$T(x_1, \dots, x_n) = f(x_I) g(x_J)$$

for some tensors f, g and some disjoint $I, J \neq \emptyset$ where $I \cup J = [n]$. Notice that T will be a slice (equivalently, have slice rank 1), if it can be written in the above form at $|S| = 1$ or $|T| = 1$. In this frame of mind, we note that a k-tensor having partition rank 1 is a "less restrictive" requirement than being a slice.

Example 7. *Let X be a finite set and $\mathbb{F}$ a field. The function $T : X^8 \to \mathbb{F}$ given by*

$$T(x_1, x_2, x_3, x_4, x_5, x_6, x_7, x_8) = f_1(x_1, x_3, x_8) f_2(x_2, x_7) f_3(x_4, x_6) f_4(x_5)$$

will have partition rank 1, with partition P given by the sets $S_1 = \{1, 3, 8\}$, $S_2 = \{2, 7\}$, $S_3 = \{4, 6\}$, and $S_4 = \{5\}$.

We can now discuss the partition rank.

4.2 The Partition Rank and Other Directions

We begin by defining the partition rank itself.

Definition 4.2.1. *Let X be a finite set, let $\mathbb{F}$ be a field and suppose that we are given a k-tensor $T : X^k \to \mathbb{F}$. The* ***partition rank*** *of T is the minimal r such that*

$$T = \sum_{i=1}^{r} g_i,$$

where each g_i has partition rank 1. For notational purposes, we denote $r = \operatorname{prk}(T)$.

The partition rank is the minimal such rank where we partition or separate the variables. However, keeping in mind that we want to find generalizations of the slice rank, it may prove useful to have notation for the rank of a k-tensor where we restrict to a specific subset of partitions $\mathcal{P}$.

Definition 4.2.2. *Let X be a finite set, let $\mathbb{F}$ be a field, let $\mathcal{P}$ be a collection of non-trivial partitions of $[n]$ and suppose that we are given a k-tensor $T : X^k \to \mathbb{F}$. We say that T has **$\mathcal{P}$-rank 1** if there exists a partition $P \in \mathcal{P}$ such that*

$$T(x_1, \ldots, x_n) = \prod_{A \in P} f_A(x_A)$$

*for some function f_A. The **$\mathcal{P}$-rank** of a function $F : X^k \to \mathbb{F}$ is the minimal r such that*

$$F = \sum_{i=1}^{r} g_i,$$

where each g_i has $\mathcal{P}$−rank 1.

Before we continue our discussion, we provide an important definition.

Definition 4.2.3. *Consider two partitions P, P' of the set $[n]$. We say that P' is a **refinement** of P if every set $S' \in P'$ is a subset of some set $S \in P$. For example, letting $P' = \{\{1, 2\}, \{3, 4\}, \{5, 6\}, \{7\}\}$ and $P = \{\{1, 2, 3, 4\}, \{5, 6, 7\}\}$, we see that P' is a refinement of P as $\{1, 2\}, \{3, 4\} \subset \{1, 2, 3, 4\}$ and $\{5, 6\}, \{7\} \subset \{5, 6, 7\}$.*

As discussed in subsection 2.2, we were able to compare the slice rank and the tensor rank of generic k-tensors. In fact, in general, we have $\text{srk}(T) \leq \text{trk}(T)$, and we are guaranteed equality when $k = 1, 2$. We would like to ask the same with the partition rank. To do so, we state and prove the following proposition, from Naslund (2020b).

Proposition 4.2.4 (Proposition 9 in Naslund (2020b))**.** *Let $\mathcal{P}, \mathcal{P}'$ be two collections of non-trivial partitions of $[n]$. Suppose that every partition $P \in \mathcal{P}$ is refined by some partition $P' \in \mathcal{P}'$. Then we have, for any k-tensor $T : X^k \to \mathbb{F}$,*

$$\mathcal{P}\text{-rank}(T) \leq \mathcal{P}'\text{-rank}(T).$$

Proof. Suppose that T has $\mathcal{P}'$-rank 1. Then, there exists $P' \in \mathcal{P}'$ and f_A such that

$$T = \prod_{A \in P'} f_A.$$

However, since P' refines some $P \in \mathcal{P}$, we may write

$$T = \prod_{B \in P} g_B,$$

where

$$g_B = \prod_{\substack{A \in P' \\ A \subset B}} f_A.$$

□

Again, we note that the partition rank is the $\mathcal{P}$-rank when $\mathcal{P}$ is the set of all non-trivial partitions. Similarly, the slice rank is the $\mathcal{P}_{\text{slice}}$-rank where $\mathcal{P}_{\text{slice}}$ denotes the set of partitions of $[n]$ into a set of size 1 and a set of size $n-1$. From this, we see that $\mathcal{P}_{\text{slice}} \subseteq \mathcal{P}$, and thus, for any k-tensor T,

$$\text{prk}(T) \leq \text{srk}(T).$$

Finally, we note that the tensor rank is the $\mathcal{P}_{\text{tensor}}$-rank where $\mathcal{P}_{\text{tensor}}$ denotes the set of partitions of $[n]$ into a n sets of size 1. However, this partition is a refinement of every partition in $\mathcal{P}_{\text{slice}}$, so by Proposition 4.2.4, we have that, for any k-tensor T,

$$\text{prk}(T) \leq \text{srk}(T) \leq \text{trk}(T).$$

When $k = 2$, using the same analysis as in subsection 2.2, we note that these three notions of rank are equivalent, since there is only one non-trivial partition of 2. When $k = 3$, the partition rank and the slice rank are equivalent, as the non-trivial partitions of 3 either split up into two sets, one of size 1 and the other of size 2, or into three sets of size 1. However, we notice that if we were to write a function as the product of three 1-tensors, we can always multiply those tensors together to create one 2-tensor. We note that, usually, the tensor rank is different when $k = 3$. When $k = 4$, these three notions of rank are usually all different, however, the partition rank can be substantially lower than the slice rank. Consider this example from Naslund (2020b).

Example 8. *Consider $k = 4$. The only partitions of $[4]$ that do not refine partitions appearing in $\mathcal{P}_{slice}$ are $\{\{1,2\},\{3,4\}\}$, $\{\{1,3\},\{2,4\}\}$, and $\{\{1,4\},\{2,3\}\}$. For a finite set X and a field $\mathbb{F}$, consider the 4-tensor $F : X^4 \to \mathbb{F}$ defined by*

$$F(x,y,z,w) = \begin{cases} 1, & x = y \text{ and } z = w, \\ 0 & \text{otherwise.} \end{cases}$$

That is,

$$F(x, y, z, w) = \delta(x, y)\delta(z, w)$$

where $\delta(x, y)$ is the function that is 1 *when $x = y$ and* 0 *otherwise. Then, by definition we note that* $\text{prk}(F) = 1$. *However, by Theorem* 4.6 *of Blasiak et al.* (2017) *and methods that are not of direct interest for this thesis, we can find that* $\text{srk}(F) = |X|$.

This new notion of the rank of a k-tensor is interesting in it of itself, however, for it to be a generalization of the slice rank of interest, we need a similar relationship between $\text{prk}(T)$ and $|X|$. From the analysis above, we note that we always have, as a result of 2.2.8, $\text{prk}(T) \leq \text{srk}(T) \leq |X|$. However, the following lemma is the relationship that we ultimately desire.

Lemma 4.2.5 (Partition Rank Lemma, Lemma 11 in Naslund (2020b).)**.** *Let $T : X^k \to \mathbb{F}$ be a diagonal k-tensor with non-zero diagonal entries. Then* $\text{prk}(T) = |X|$.

Proof. As we have already shown that the partition rank is at most $|A|$, our goal is to prove the lower bound. The proof proceeds by induction on the number of variables. When $n = 2$, this is the usual notion of rank, and so the result follows. Suppose that F has partition rank $r < |A|$, that is suppose that we can write

$$F(x_1, \ldots, x_n) = \sum_{i=1}^{r} f_i(x_{S_i}) g_i(x_{T_i})$$

for some sets S_i, T_i with $S_i \cap T_i = \emptyset$ and $S_i \cup T_i = \{1, \ldots, n\}$. Assume without loss of generality that $|S_i| \leq \frac{n}{2}$ for each i. If there is no i such that $|S_i| = 1$, then choose an arbitrary variable, say x_1, and average over that coordinate. Then

$$\sum_{x_1 \in X} F(x_1, \ldots, x_n) = \sum_{a \in A} c_a \delta_a(x_2) \cdots \delta_a(x_n) = \sum_{i=1}^{r} \tilde{f}_i(x_{S_i \setminus \{1\}}) \tilde{g}_i \left(x_{T_i \setminus \{1\}}\right),$$

for functions $\tilde{f}, \tilde{g}$ given by averaging f, g over x_1. This contradicts the inductive hypothesis since $\sum_{a \in A} c_a \delta_a(x_2) \cdots \delta_a(x_n)$ will have partition rank equal to $|A| > r$.

Suppose that there exists some S_i such that $|S_i| = 1$. Then $S_i = \{j\}$ for some $j \in \{1, \ldots, n\}$. Let U be the set of indices u for which $S_u = \{j\}$. Consider the annihilator of U, defined to be

$$V = \left\{ h\colon X \to \mathbb{F} : \sum_{x_j \in X} f_u(x_j)h(x_j) = 0 \text{ for all } u \in U \right\}.$$

This vector space has dimension at least $|X| - |U|$, and this will be positive since $|U| \leq r < |A| \leq |X|$. Let $v \in V$ have maximal support, and set $\Sigma = \{x \in X : v(x) \neq 0\}$. Then $|\Sigma| \geq \dim V \geq |X| - |U|$, since otherwise there exists non-zero $w \in V$ vanishing on Σ, and the function $v + w$ would have a larger support than v. Multiplying both sides of our expression by $v(x_j)$ and summing over x_j reduces the dimension by 1. Indeed

$$\begin{aligned} &\sum_{x_j \in X} v(x_j)F(x_1, \ldots, x_n) \\ &\qquad = \sum_{a \in A} c_a \delta_a(x_1) \cdots \delta_a(x_{j-1})\delta_a(x_{j+1}) \cdots \delta_a(x_n) \left(\sum_{x_j \in X} v(x_j)\delta_a(x_j) \right), \end{aligned}$$

and since the sum $\sum_{x_j \in X} v(x_j)\delta_a(x_j)$ will be non-zero for at least $|X| - |U|$ values of $a \in X$, the partition rank of the above must be at least $|A| - |U|$ by the inductive hypothesis. Since

$$\sum_{x_j \in X} v(x_j) f_i(x_{S_i}) = 0$$

for each $i \in U$, it follows that

$$\sum_{x_j \in X} v(x_j) \sum_{i=1}^{r} c_i f_i(x_{S_i}) g_i(x_{T_i})$$

will be a sum of at most $r - |U|$ partition rank 1 functions, and hence it has partition rank at most $r - |U|$. This implies that $|A| - |U| \leq r - |U|$, which is a contradiction, and the lemma is proven. □

As a direct result of the Partition Rank Lemma, we can use the partition rank as we have been using the slice rank in previous applications. As for a diagonal k-tensor $T : X^k \to \mathbb{F}$, we have $\text{prk}(T) = |X| = \text{srk}(T)$, it may seem that using the partition rank does not lead to an improvement as opposed to using the slice rank. However, in practice, the analysis of the partition rank allows for tighter bounds on $|X|$. In fact, with the introduction of the partition rank, Naslund (2020b) was able to generalize the work of **?**, which provided a bound on the size of a set $A \subset \mathbb{F}_q^n$ which contained no right corners, to k-right corners for fields $\mathbb{F}_q^n$, where $q = p^r$ for some prime $p > k$. The analysis used in Naslund's proof would not have been possible using the slice rank, the full generality was a necessity.

From the formulation of the partition rank, we have a generalization in the direction of how we partition the variables of our k-tensor. That is, for the slice rank, given a k-tensor, we would consider how to write our k-tensor as the sum of the product of 1- and $(k-1)$- tensors. However, with the partition rank, we are not limited to just this form in our product, we are free to choose any partition of $[n]$. It is also worth noting, however, that the $\mathcal{P}$-rank allows for control as to what types of partition we do and don't allow for. At the moment, there does not exist a criterion for equality between the $\mathcal{P}$-rank and $|X|$, however, this allows for a possible avenue for generalization.

Another possible avenue of generalization is to understand the relationship between the slice rank (or any generalization of the slice rank) of a k-tensor $T : X^k \to \mathbb{F}$ and $|X|$ if our k-tensor is not diagonal. For example, is it the case that $\text{srk}(T) = |X|$ if T is a block-diagonal k-tensor (with whatever it may mean for a k-tensor to be block-diagonal), or if T is a upper- or lower-triangular k-tensor? If we can say something to that extent, then we would be less restricted in the construction of our k-tensors before applying the slice rank method, or any of its generalizations like the partition rank. This would allow for a definition of a k-tensor with necessarily smaller slice rank and thus, a tighter bound on $|X|$. The question just becomes if a relationship of this manner exists with k-tensors that are not diagonal.

Bibliography

Alon, Noga, Amir Shpilka, and Christopher Umans. 2012. On sunflowers and matrix multiplication. In *2012 IEEE 27th Conference on Computational Complexity*, 214–223. IEEE.

Babai, László, and Péter Frankl. 1988. *Linear algebra methods in combinatorics*. University of Chicago.

Baker, Roger C., Glyn Harman, and János Pintz. 2001. The difference between consecutive primes, ii. *Proceedings of the London Mathematical Society* 83(3):532–562.

Berlekamp, Elwyn R. 1969. On subsets with intersections of even cardinality. *Canadian Mathematical Bulletin* 12(4):471–474.

Blasiak, Jonah, Thomas Church, Henry Cohn, Joshua A Grochow, Eric Naslund, William F Sawin, and Chris Umans. 2017. On cap sets and the group-theoretic approach to matrix multiplication. *Discrete Analysis* 3.

Buchanan, Calum. 2020. Eventown and oddtown. lecture notes.

Croot, Ernie, Vsevolod F. Lev, and Péter Pál Pach. 2017. Progression-free sets in $\mathbb{Z}_4^n$ are exponentially small. *Annals of Mathematics* 331–337.

Ellenberg, Jordan S., and Dion Gijswijt. 2017. On large subsets of with no three-term arithmetic progression. *Annals of Mathematics* 339–343.

Erdös, Paul, and Endre Szemerédi. 1978. Combinatorial properties of systems of sets. *Journal of Combinatorial Theory, Series A* 24(3):308–313.

Frankl, Peter, and Vojtěch Rödl. 1987. Forbidden intersections. *Transactions of the American Mathematical Society* 300(1):259–286.

Håstad, Johan. 1989. Tensor rank is NP-Complete. In *International Colloquium on Automata, Languages, and Programming*, 451–460. Springer.

Jukna, Stasys. 2011. *Extremal combinatorics: with applications in computer science*. Springer Science & Business Media.

Lickteig, Thomas. 1985. Typical tensorial rank. *Linear algebra and its applications* 69:95–120.

Naslund, Eric. 2020a. Monochromatic equilateral triangles in the unit distance graph. *Bulletin of the London Mathematical Society* 52(4):687–692.

———. 2020b. The partition rank of a tensor and k-right corners in $\mathbb{F}_q^n$. *Journal of Combinatorial Theory, Series A* 174:105,190.

Naslund, Eric, and Will Sawin. 2017. Upper bounds for sunflower-free sets. In *Forum of Mathematics, Sigma*, vol. 5. Cambridge University Press.

Pellegrino, Giuseppe. 1970. Sul massimo ordine delle calotte in $s_{4,3}$. *Matematiche (Catania)* 25(10).

Petrov, Fedor. 2016. List of counting proofs instead of linear algebra method in combinatorics. mathoverflow comment.

Ross, Sheldon. 2014. *A first course in probability*. Pearson, 9th ed.

Sagdeev, Arsenii A. 2018. Improved Frankl–Rödl theorem and some of its geometric consequences. *Problems of Information Transmission* 54(2):139–164.

Szabó, Tibor. 2017. Three-wise restrictions and the slice rank. lecture notes.

———. 2019. Linear algebra method. lecture notes.

Tao, Terence. 2016a. Notes on the slice rank of tensors. blog post.

———. 2016b. A symmetric formulation of the Croot-Lev-Pach-Ellenberg-Gijswijt capset bound. blog post.