Extremal Set Theory and the Linear Algebra Method

Tibor Szabó

Extremal Combinatorics, FU Berlin, WiSe 2017–18

In extremal set theory we consider families $\mathcal{F} \subseteq 2^V$ of subsets of an underlying ground set V and study their properties. The typical question, not entirely unexpectedly, is how large can a set family be if it does not contain a given forbidden configuration.
Extremal set theory is full of beautiful theorems, surprising connections, and tantalizing open problems and conjectures. Since the notion of set family is a very flexible one, there are also many applications, for instance to geometry. In terms of methods we will of course see an abundance of combinatorial techniques and a couple of probabilistic proofs, but the main emphasis in this chapter will be on applications of linear algebra.
There is no formal difference between the term *set family* (or *set system*) and the term *hypergraph*. The difference in the various uses of the terms is more philisophical, and in fact there is no formal obstacle in exchanging the two. In a hypergraph the focus is more on vertices, vertex subsets being in "relation", and a subset of vertices hosting or not hosting a specific configuration of relations. In a set family the focus is more on set theoretic properties of the sets, the family having or not having sets in a certain set theoretic relation.
What do humans do? Well, they think and they love. What do sets do? They intersect and they contain. We will classify our highlight tour of extremal set theory along these two main themes. In the first section we will sample four of the classic results of the field, each problem being more than half a century old. First we will feature hypergraph Turán numbers, which form a natural bridge from the previous chapter to this one. We continue with a couple of classic problems focussing on intersections of different kinds. Finally, we present Sperner's Theorem where containment is the main concern.

1 The Classics

1.1 Hypergraph Turán problems

You might have wondered that why is it that in Turán theory, unlike in Ramsey-theory, we did not go into the hypergraph generalization. There is a simple reason for that: not much is known and even much less is presentable in a lecture within the scope of this course. Therefore, in this section we will confine ourselves to constructions, conjectures, and theorems without proof.

Definition 1.1. *For an integer $n \in \mathbb{N}$ and a k-uniform hypergraph $\mathcal{H}$, the k*-uniform Turán number $\mathrm{ex}(n, \mathcal{H})$ *is the largest integer m such that there exists an $\mathcal{H}$-free k-uniform hypergraph on n vertices with m edges.*

1.1.1 Turán's $(3, 4)$-Conjecture

The smallest meaningful case of generalizing Mantel's Theorem about dense K_3-free graphs to hypergraphs is the Turán number of the complete 3-uniform clique $\mathcal{K}_4^{(3)}$ on four vertices. In other words: how many triples can one cram into n vertices but still avoid having all four triples on some four vertices.

Construction For $n \in \mathbb{N}$ let $\mathcal{G}_n$ be the 3-graph on vertex set $V_0 \cup V_1 \cup V_2 = [n]$, where $||V_i| - |V_j|| \leq 1$, for every $0 \leq i < j \leq 2$, and

$$\mathcal{G}_n = \left\{ T \in \binom{[n]}{3} : |T \cap V_i| = 1 \ \forall i \in \{0,1,2\} \text{ or } \exists i \in \{0,1,2\},\ |T \cap V_i| = 2, |T \cap V_{i+1}| = 1 \right\},$$

where we assume $V_0 = V_{2+1}$.

Proposition 1.2. *$\mathcal{G}_n$ contains no copy of $\mathcal{K}_4^{(3)}$.*

Corollary 1.3. *$ex(n, K_4^{(3)}) \geq \left(\frac{5}{9} + o(1)\right)\binom{n}{3}$*

Conjecture 1.4 (Turán's Conjecture ($1000 dollar question))**.**

$$ex(n, K_4^{(3)}) = |E(\mathcal{G}_n)|.$$

Remark 1.5. *If the conjecture is true, then there are exponentially many extremal constrcutions (Kostochka). In a Homework, you will verify some of these constructions.*

This conjecture motivated one of the recent advances in extremal combinatorics, the *flag algebra method* of Razborov. The method enables one to find upper bound proofs for extremal combinatorial problems using a computerized search. What earlier was proved using an ingenous and insghtful double-counting idea, now is verified by the brute force of compter through an exhaustive search for what to double-count. Such proofs are often less pleasing, but certainly effective. The current best known upper bound for the Turán number of $\mathcal{K}_4^{(3)}$ is $0.562\binom{n}{3}$. It was found by the flag algebra method and is very close to the above lower bound, but still a positive constant factor times $\binom{n}{3}$ away.

1.1.2 The Fano plane

An alternative way to view the triangle (instead of being the smallest clique larger than the uniformity) is that it is the smallest non-2-colorable graph.
Recall Property B and our Homework exercise where we showed that a smallest non-2-colorable 3-uniform hypergraph is the Fano plane $\mathcal{F}$, where

$$V(\mathcal{F}) = [7] \quad \text{and} \quad \mathcal{F} = \{123, 345, 561, 174, 376, 572, 246\}.$$

Construction Let $\mathcal{B}$ be the 2-colorable hypergraph with the most edges: on vertex set $V_1 \cup V_2 = [n]$ with $|V_1| = \lfloor \frac{n}{2} \rfloor$ and $|V_2| = \lceil \frac{n}{2} \rceil$, take

$$\mathcal{B} = \left\{ T \in \binom{[n]}{3} : T \cap V_i \neq \emptyset \text{ for both } i = 1, 2 \right\}.$$

Since $\mathcal{B}$ is 2-colorable it cannot contain any copy of $\mathcal{F}$, consequently $ex(n, \mathcal{F}) \geq \left(\frac{3}{4} + o(1)\right)\binom{n}{3}$. It was one of the (very few) success stories of hypergraph Turán theory when this construction was proved to be best possible.

Theorem 1.6 (De Caen-Füredi, Keevash-Sudakov, Füredi-Simonovits, 2006)**.**

$$ex(n, \mathcal{F}) = |E(\mathcal{B})| = \left(\frac{3}{4} + o(1)\right)\binom{n}{3}.$$

1.2 Intersecting families

The following is a funcamental property of set families.

Definition 1.7. *A set system $\mathcal{F} \subseteq 2^{[n]}$ is called* intersecting*, if $F_1 \cap F_2 \neq \emptyset$ for every $F_1, F_2 \in \mathcal{F}$.*

The natural extremal problem asks how large an intersecting family on n vertices can be. This is the question we answer in this section.

1.2.1 Arbitrary families

We start by constructing what we think are large intersecting families.

Examples 1.8. (1) *Take all sets contain a fixed element $x \in [n]$, i.e., a* ***star****. The resulting set system has size 2^{n-1} and it is clearly intersecting, because every set contains x.*

(2) *Let $\mathcal{F} = \{F \in 2^{[n]} \ : \ |F| > n/2\}$ be the family of large sets. Since $|F_i| + |F_j| > n$, any two sets must have a non empty intersection. If n is odd, then $|\mathcal{F}| = 2^{n-1}$. When n is even, then the $\frac{n}{2}$-element sets form the middle layer of the poset and do not belong to our family. We could extend our $\mathcal{F}$ by taking those sets from the middle layer that contain a fixed element.*

These two constructions are very different and give the same bound 2^{n-1}. Is this best possible? Is it true that no matter how we select $2^{n-1}+1$ subsets of $[n]$, there will be two disjoint among them? In the next proposition we show that the answer is YES.

Proposition 1.9. *If $\mathcal{F} \subseteq 2^{[n]}$ is intersecting, then $|\mathcal{F}| \leq 2^{n-1}$.*

Note that this implies that our constructions were best possible.

Proof. We split $\mathcal{F}$ into 2^{n-1} pairs of subsets $\{F, F^C\}$ containing a set F and its complement. If $\mathcal{F}$ is intersecting, then it cannot contain both members of such a pair, hence $\mathcal{F}$ contains at most one of each such complementary pair. Thus $|\mathcal{F}|$ is at most the number o pairs, which is 2^{n-1}. $\square$

1.2.2 Uniform families

The same question can be asked for uniform families. It is always a good idea to start with looking for constructions, in particular how the above two constructions survive in the uniform setup. The size of a k-uniform star is now $\binom{n-1}{k-1}$: after fixing an element $x \in [n]$ there are $\binom{n-1}{k-1}$ ways to select the $k-1$ extra element to add to it out of the $n-1$ elements in $[n] \setminus \{x\}$. The relative size of this construction is $\frac{\binom{n-1}{k-1}}{\binom{n}{k}} = \frac{k}{n}$; for $k = o(n)$ it is much smaller than the relative size $\frac{2^{n-1}}{2^n} = \frac{1}{2}$ of the not-necessarily-uniform case. Naturally, we wonder whether we can do any better.
Next let us revisit the second construction above, when the intersecting property was ensured by the inequality $|F_1| + |F_2| > n$. Now this works when $|F_1| + |F_2| = 2k > n$. In this case the family $\binom{[n]}{k}$ of all k-subsets is intersecting, giving us a better construction in the range $k > \frac{n}{2}$.
We restrict our attention to the case when $k \leq \frac{n}{2}$. The next theorem says that in this range the first construction is optimal.

Theorem 1.10 (Erdős-Ko-Rado, 1961)**.** *If $n \geq 2k$, and $\mathcal{F} \subseteq \binom{[n]}{k}$ is intersecting, then*

$$|\mathcal{F}| \leq \binom{n-1}{k-1}.$$

We would like to stress that this is a fundamental result in extremal set theory. A number of different proofs, extensions and generalisations have been found since. The proof we will show here is based on a tricky double-counting, due to Katona (1974).

Proof. Fix an arbitrary intersecting family $\mathcal{F} \subseteq \binom{[n]}{k}$. We will double-count pairs of how many times it occurs that a cyclic permutation contains a member of our family as an interval. Let C_n denote the set of cyclic permutations, so $|C_n| = (n-1)!$. A set $F \subseteq [n]$, if the elements of F occur consecutively along the cycle in π. Formally we count the set

$$M = \{(\pi, F) : \pi \in C_n, F \in \mathcal{F} \text{ is an interval on } \pi\}.$$

First let us fix an arbitrary set F and count that in how many cyclic permutations it occurs as an interval. There are $|F|!$ ways to place the elements of F in positions $1, \ldots, |F|$ along the cycle.

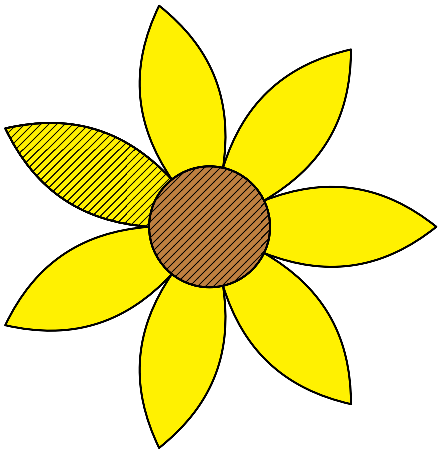

Figure 1: The central part is the kernel of the sunflower and the kernel along with a petal gives us a member of the set family (the shaded part).

Then there are $(n - |F|)$ ways to place the rest of the elements for the remaining places along the cycle. So all together there are

$$|M| = \sum_{F \in \mathcal{F}} |F|!(n - |F|)! = |\mathcal{F}|k!(n - k)!$$

elements in M.
Now let us count by fixing a cyclic permutation π first. How many members of $\mathcal{F}$ can appear on π as an interval? Fixing a k-set $F \in \mathcal{F}$ that forms an interval on π, the intersecting property ensures that every other member must start or end on this interval. This immediately gives a trivial bound of $2k - 2$ for the number of such members besdies F. But we want only $k - 1$ extra members that can form an interval on π.
Let us say that the elements of F are $x_1, x_2, \ldots, x_k$, that appear along the cycle in this clockwise order. Pair up intervals that end or start on these elements as follows. For every $i, 1 \leq i \leq k - 1$, the k-interval ending on x_i should be paired up with the k-interval starting on x_{i+1}. As $n \geq 2k$, these two intervals do not intersect, sonsequently they cannot both be a member of $\mathcal{F}$. From each of the $k - 1$ pairs there is at most one in $\mathcal{F}$, so all in all there are no more than k members of $\mathcal{F}$ that appear as interval on π.
This shows that $|M| \leq |C_n| \cdot k = (n - 1)! \cdot k$.
In conclusion

$$|\mathcal{F}|k!(n - k)! = |M| \leq (n - 1)! \cdot k,$$

which is equivalent to what we want.

$\square$

Remark 1.11. *If $n > 2k$, then one can show that stars are the* only *intersecting families of size $\binom{n-1}{k-1}$. This requires a more careful analysis of the proof.*

1.3 Sunflowers

In this section we will be concerned with not just the pairwise intersections, but also the intersections of arbitrary many members of the family.

A set family $\mathcal{S}$ is called a sunflower (or Δ-*system*) if $A \cap B = \cap_{F \in \mathcal{S}} F$ for every $A, B \in \mathcal{S}$. The set $\cap_{F \in \mathcal{S}} F$ is called the *kernel* of the sunflower and $A \setminus \cap_{F \in \mathcal{S}} F$, for $A \in \mathcal{S}$ are its *petal*s.

Set families can have pretty wild and scrambled intersection relations. So when we study them it often comes in handy if one can locate a large organized subfamily. Sunflowers are as organized as it gets in terms of their multiple intersections. The following theorem guarantees a sunflower subfamily of arbitrary size if the number of sets is large enough.

Theorem 1.12 (Erdős-Rado, 1960)**.** *Let $\ell \in \mathbb{N}$. If $\mathcal{F}$ is a k-uniform set family and $|\mathcal{F}| > (\ell-1)^k k!$ then $\mathcal{F}$ contains a sunflower with ℓ petals.*

Remark 1.13. *An important feature of the result is that the size of the vertex set is* not *part of the statement.*

Proof. We use induction on the uniformity k. For $k = 1$ having more than $(\ell - 1)^1 1!$ one-element sets in the family means that there are ℓ pairwise disjoint sets: a sunflower with an empty kernel. Let $k > 1$ and let $\mathcal{F}$ be a k-uniform family without a sunflower with ℓ petals. Then there exists a set X of at most $(\ell - 1)k$ elements that intersects every $F \in \mathcal{F}$. Indeed, one can take for example the union of the sets in a maximal subfamily of pairwise disjoint sets from $\mathcal{F}$. Such a subfamily is a sunflower hence has at most $\ell - 1$ members, implying the upper bound $(\ell - 1)k \geq |X|$.
We classify members of $\mathcal{F}$ according to their intersection with X. For $x \in X$, let $\mathcal{F}_x = \{F \setminus \{x\} : F \in \mathcal{F}, x \in F\}$. These families also contain no sunflower with ℓ petals, but they are $(k-1)$-uniform. By induction we have $|\mathcal{F}_x| \leq (\ell - 1)^{k-1}(k - 1)!$ for every $x \in X$.
Then
$$|\mathcal{F}| \leq \sum_{x \in X} |\mathcal{F}_x| \leq |X| \cdot ((\ell - 1)^{k-1}(k - 1)!) \leq (\ell - 1)^k k!.$$
$\square$

The bound in the theorem is known not to be tight, and it is wide open what the best possible bound would be, even if we only wanted $\ell = 3$ petals. The following simple construction shows that it has to be at least exponential in k.
Construction Let $X = \{x_1, \ldots, x_k, y_1, \ldots, y_k\}$ be the vertex set and let
$$\mathcal{F} = \{F \subseteq X : \ |F \cap \{x_i, y_i\}| = 1 \text{ for every } i\}.$$
Then $\mathcal{F}$ has no sunflower with three petals and $|\mathcal{F}| = 2^k$.
There are better constructions with c^k members where c is some constant larger than 2 (**HW**), but no super-exponential construction is known. In fact Erdős and Rado thought the lower bound should be closer to the truth (and Erdős certainly put his money where his mouth was.)

Conjecture 1.14 (Erdős-Rado ($1000 dollar question))**.** *There is a constant C such that every k-uniform family with C^k sets contains a sunflower with three petals.*

The best known upper bound, due to Kostochka, is just slightly smaller than $k!$.

1.4 Antichains

This is our classic section about the containment relation of sets. Forbidding non-containment in a set family on the vertex set $[n]$ leads to the notion of *chain* in the Boolean poset[1] $(2^{[n]}, \subseteq)$. The maximum size of a chain is $n + 1$.
Forbidding containment leads to the notion of *antichain* in the Boolean poset. The natural extremal question is how many subsets of $[n]$ can be selected if it is forbidden to select two sets such that one is subset of the other?
One can select all $\binom{n}{k}$ subsets of a given size k: they certainly satisfy this property. The choice $k = \left\lfloor \frac{n}{2} \right\rfloor$ maximizes this number.

[1]A *poset* (or *partially ordered set*) $(P, \preceq)$ is a set P together with a binary relation $\preceq$ that is reflexive, antisymmetric, and transitive. The family 2^X of all subsets of a set X is a poset with the containment relation $\subseteq$.

Theorem 1.15 (Sperner's Theorem, 1928). *If $\mathcal{F} \subseteq 2^{[n]}$ is a family of subsets such that for every $A, B \in \mathcal{F}$ we have $A \not\subseteq B$ then*

$$|\mathcal{F}| \leq \binom{n}{\lfloor n/2 \rfloor}.$$

Proof. For the proof we revisit the permutation method. We will count permutations $\pi \in S_n$ of $[n]$ which have an initial segment from $\mathcal{F}$. Formally, double-count the set

$$M = \{(\pi, F) : \pi \in S_n, F \in \mathcal{F}, F = \{\pi(1), \ldots, \pi(|F|)\}\}$$

On the one hand for every $F \in \mathcal{F}$ there are $|F|!(n-|F|)!$ permutations $\pi \in S_n$ with $\{\pi(1), \ldots, \pi(|F|)\} = F$. So

$$|M| = \sum_{F \in \mathcal{F}} |F|!(n - |F|)!.$$

On the other hand for every $\pi \in S_n$ there is *at most one* index k such that $\{\pi(1), \ldots, \pi(k)\} \in \mathcal{F}$, because otherwise there will be two members of $\mathcal{F}$ with one containing the other. So $M \leq n!$. Hence

$$\sum_{F \in \mathcal{F}} |F|!(n - |F|)! \quad \leq \quad n!$$

$$1 \geq \sum_{F \in \mathcal{F}} \frac{1}{\binom{n}{|F|}} \geq \sum_{F \in \mathcal{F}} \frac{1}{\binom{n}{\lfloor \frac{n}{2} \rfloor}} = |\mathcal{F}| \frac{1}{\binom{n}{\lfloor \frac{n}{2} \rfloor}}$$

□

We will continue on this theme in the third section of this chapter, but first discuss extensions and variants of the Erdős-Ko-Rado where the main focus is on the pairwise intersections.

2 Families with intersection restrictions

2.1 A unique permissible intersection size

In the Erdős-Ko-Rado theorem, both in its uniform and its not-necessarily-uniform versions, we forbid empty pairwise intersections. In other words the set of allowed sizes for pairwise intersections is $\{1, 2, \ldots, n-1\}$. What if we restrict us more and allow only a single intersection size? That is, let $\lambda \in \mathbb{N}$ be an integer and $\mathcal{F} \subseteq 2^{[n]}$ such that $|F_1 \cap F_2| = \lambda$ for every $F_1, F_2 \in \mathcal{F}$, $F_1 \neq F_2$. How large can then $\mathcal{F}$ be?
We will tackle this problem in a bit, but we first collect ideas by considering a modular variant (with a story [1]).

2.1.1 Eventown vs. Oddtown.

In a little town called Eventown the 32 citizens love to organize various clubs. But they have to follow the strict traditional Eventown-rules of which there are two:

E1 Every club has to have an even number of members.

E2 Every pair of clubs has to have an even number of members in common.

The city council at some point is facing an administrative nightmare due to the number of clubs getting out of control. Indeed, if the citizens for example, were to pair themselves up and would join or not join clubs only together with their pairs, then the system of all $2^{16} = 65536$ such possibilities would create a feasible club system according to the above Eventown-rules. The mayor wants to cut down the number of clubs and considers changing the century old rules whose motivation is anyway lost in the obscurity of old times. After consultation with the wise, they consider the following slight modification of the Eventown-rules.

O1 Every club has to have an odd number of members.

O2 Every pair of clubs has to have an even number of members in common.

The difference is in only one word, the conditions **E2** and **O2** are the same. How many clubs could there be now? Let us first take a look at some construction ideas, with the set of citizens being denoted by $[n]$.

1) Taking the n singletons creates n clubs of size 1 each with all pairwise intersection 0.

2) When n is even, one could take the complement of singletons. This creates an $(n-1)$-uniform family with all pairwise intersections having size $n-2$.

3) Again for n even, we could consider the so-called *two-star* construction. Let $F_i = \{i, n-1, n\}$ for $1 \leq i \leq n-2$, $F_{n-1} = \{1, 2, ..., n-2, n-1\}$, and $F_n = \{1, 2, ..., n-2, n\}$. This system is not anymore uniform, but the possible set sizes are 3 and $n-1$, both odd and for the pairwise intersections we have $|F_i \cap F_j| \in \{2, n-2\}$ for $i \neq j$.

All these construction have n sets and there are in fact many more such constructions. The next theorem shows that one cannot do better.

Theorem 2.1 (Oddtown Theorem, Berlekamp, 1969)**.** *Let $\mathcal{F} \subseteq 2^{[n]}$ be a set family satisfying both* **O1** *and* **O2***. Then $|\mathcal{F}| \leq n$.*

This theorem is pretty significant for the mayor of Eventown. With changing just a single word in the rules, the number of possible clubs is reduced from exponential to linear. The council votes to change the name of the town to Oddtown and they live happily ever after.

Before the actual proof of the Oddtown Theorem we make a few general comments in order to motivate and introduce the general method that will be of use in this chapter.
In a typical extremal combinatorial problem, the greater the number of extremal families[2], the less likely that a purely combinatorial argument will lead to a solution, since a proof eventually must consider all extremal structures, and be tight for each of them in each of the proof steps. If these families are combinatorially very different, this might necessarily lead to an unmanageable number of combinatorial case distinctions.
For some of these problems the stars are aligned and the difficulties posed by multiple extremal examples can be mitigated by realizing that the combinatorial problem, or rather its extremal structures, hides the features and concepts of another mathematical discipline in the background. In such cases, the simplest, or most efficient descriptions of extremal structures are not necessarily combinatorial, but might have to be formulated in another language, which could be algebraic, probabilistic, or, even topological. Solutions of this sort, connecting different branches of mathematics, are considered gems: they are rare and beautiful.
The Oddtown Theorem is one of those situations, where the the extremal families are far from being unique. In fact one can prove that their number is super-exponential [**?**][Exercise 1.1.14]. For the (easy) proof of Theorem 2.1 one only has to realize that the right language of the problem is the one of linear algebra. And then even though the number of extremal set-systems is superexponential and the feasibility of their combinatorial characterization is questionable at best, they have a very simple linear algebraic description as the orthonormal bases in $\mathbb{F}_2^n$.
The connection between combinatorics and algebra is provided through the characteristic vector $\mathbf{v}_F \in \{0,1\}^n$ of sets $F \subseteq [n]$, where $(v_F)_i = 1$ if and only if $i \in F$.
The key realization relevant to families with pairwise intersection restrictions is that the size of the intersection of two sets is equal to the standard inner product of the two characteristic vectors:

$$|A \cap B| = \sum_{i=1}^{n} (\mathbf{v}_A)_i (\mathbf{v}_B)_i =: \mathbf{v}_A \cdot \mathbf{v}_B.$$

[2]In an extremal combinatorial problem an *extremal family* is one with the largest (smallest) possible number of edges among those with the required property. In our problem these are the set families of size n that satisfy both **O1** and **O2**.

Indeed, when calculating a term $(\mathbf{v}_A)_i(\mathbf{v}_B)_i$ of the sum, we have a 1 if and only if $i \in A \cap B$. This connection will allow us to translate the combinatorial condition we have on a family into linear algebra and use this information to derive the linear independence of the characteristic vector. This in turns makes the dimension of the space an upper bound on their number.
The simple linear algebra fact that the size of a linearly independent set of vectors is *at most* the dimension of the ambient vector space is called the *dimension bound.* It is the simplest manifestation of the Linear Algebra method.

Proof of the Oddtown Theorem. Let $\mathcal{F} = \{C_i : 1 \leq i \leq m\}$. To each set C_i we associate its characteristic vector $\mathbf{v}_i \in \{0,1\}^n$. From rules **O1** and **O2**, using $\mathbf{v}_i \cdot \mathbf{v}_j = |C_i \cap C_j|$, we infer that $\mathbf{v}_i \cdot \mathbf{v}_i$ is odd for every $i \in [n]$ and $\mathbf{v}_i \cdot \mathbf{v}_j$ is even for every $i \neq j$.
We would like to claim the linear independence of the vectors $\mathbf{v}_1, \ldots, \mathbf{v}_m$, for which should name a field to work over. Considering our conditions on the pairwise dot products, it does not come as a great shock that we choose to show linear independence over the two-element field $\mathbb{F}_2$.
Suppose we have a linear combination $\sum_{i=1}^{m} \alpha_i \mathbf{v}_i = \mathbf{0}$, with $\alpha_i \in \mathbb{F}_2$. For every $j \in [m]$, we take the dot product of it with $\mathbf{v}_j$ and obtain

$$0 = \mathbf{0} \cdot \mathbf{v}_j = \left(\sum_{i=1}^{m} \alpha_i \mathbf{v}_i \right) \cdot \mathbf{v}_j = \sum_{i=1}^{m} \alpha_i \left(\mathbf{v}_i \cdot \mathbf{v}_j\right) = \alpha_j \left(\mathbf{v}_j \cdot \mathbf{v}_j\right) = \alpha_j.$$

Here we used that all but one of the dot products are zero over $\mathbb{F}_2$. This implies that $\alpha_j = 0$ for every j, consequently the vectors are indeed linear independent. Hence, their number cannot be more than the dimension of the space and we get $m \leq \dim \mathbb{F}_2^n = n$. $\square$

2.1.2 Fisher's Inequality

Let us return to our original problem, where we restricted the size of every pairwise intersection to a unique non-zero integer λ. It turns out that the same upper bound holds here as well.

Theorem 2.2 (Non-uniform Fisher Inequality)**.** *Let $n \in \mathbb{N}$ and suppose $1 \leq \lambda \leq n$. If $\mathcal{F} \subseteq 2^{[n]}$ satisfies $|F_1 \cap F_2| = \lambda$ for every $F_1, F_2 \in \mathcal{F}$, $F_1 \neq F_2$, then $|\mathcal{F}| \leq n$.*

Note that if $\lambda = 0$, then the sets F_i must be pairwise disjoint and hence we can have at most $n+1$ such sets (including the $\emptyset$).
Fischer's Inequality originates from statistics/experiment design, where the problem is how to divide test subject between treatments fairly. There are several other versions of the inequality, but this one is more general.
Constructions: for $\lambda = 1$ we can take near pencils (see Figure 2) and (nondegenerate) projective planes. The case $\lambda = 1$ is the De Bruijn-Erdős theorem. For $\lambda = n-2$ we can take the complements of singletons.
It is an open problem to classify all cases of equality. Can we have equality for all λ?

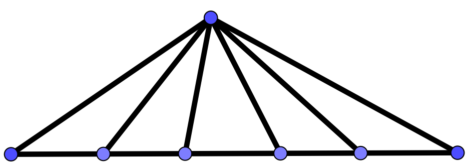

Figure 2: A near pencil, a.k.a., degenerate projective plane

References

[1] L. Babai and P. Frankl. Linear Algebra Methods in Combinatorics.