1 Introduction

A simple and well-known fact in graph theory is that every n-vertex tree has $n-1$ edges. This immediately implies that if an n-vertex graph G has no cycles, then G has at most $n-1$ edges. Another well-known result in graph theory, following quickly from Euler's formula, is that an n-vertex planar graph G with $n \geq 3$ has at most $3n-6$ edges.

This class is about extremal graph theory, the study of results of this type. How many edges can an n-vertex graph have, given that it satisfies some natural constraint? Our ultimate goal is to prove the Erdős–Stone–Simonovits theorem, sometimes called the Fundamental Theorem of Extremal Graph Theory, which answers this question more or less completely for a very wide range of constraints.

The question that will occupy us for the majority of the class is what happens when the constraint is excluding a single "forbidden subgraph".

Definition 1.1. Let H and G be graphs. We say that G is *H-free* if H is not a subgraph of G (or, more formally, if G has no subgraph isomorphic to H). We will often also say that G has *no copy of H*.

The basic question we will be attempting to answer is "how many edges can an n-vertex H-free graph have?". Because we will be using this notion over and over again, it's best to just give it a name. We use $e(G)$ to denote the number of edges of a graph G.

Definition 1.2. The *extremal number* of H is defined as

$$\mathrm{ex}(n, H) = \max\{e(G) \mid G \text{ is an } n\text{-vertex } H\text{-free graph}\}.$$

In other words, $\mathrm{ex}(n, H)$ is simply the most number of edges that an H-free graph on n vertices can have. Note that this quantity is well-defined, since there are only finitely many n-vertex graphs.

In this class, we will attempt to understand how the function $\mathrm{ex}(n, H)$ behaves when H is some fixed graph, and when n tends to infinity. Additionally, we will often try to understand which graphs G are the maximizers in the definition of $\mathrm{ex}(n, H)$; that is, which graphs G have the most edges among all n-vertex H-free graphs.

Before getting into specific examples, let's briefly think about what it means to prove upper and lower bounds on $\mathrm{ex}(n, H)$. Since $\mathrm{ex}(n, H)$ is defined as the maximum of something, to prove a *lower* bound on $\mathrm{ex}(n, H)$, it suffices to exhibit an n-vertex graph G with no copy of H; such a G gives us the lower bound $\mathrm{ex}(n, H) \geq e(G)$. On the other hand, to prove an *upper* bound on $\mathrm{ex}(n, H)$, we need to prove that *every* n-vertex graph G with m edges has a copy of H; this yields the upper bound $\mathrm{ex}(n, H) < m$.

2 Forbidden cliques: Mantel's and Turán's theorems

The earliest result in extremal graph theory is due to Mantel, from more than 100 years ago. Mantel studied (though not in this language) the extremal number of the triangle graph, K_3. Let's begin by coming up with a lower bound on $\mathrm{ex}(n, K_3)$.

After playing around with it a bit, it's pretty natural to come up with the following construction. Let $G = K_{a,b}$ be a complete bipartite graph, where $a + b = n$. Then G is certainly triangle-free, since K_3 is not bipartite. Moreover, the number of edges in G is simply ab. So we find that

$$\mathrm{ex}(n, K_3) \geq ab \qquad \text{for all integers } a, b \text{ with } a + b = n.$$

Since we want as good a lower bound as possible, we want to pick a, b so that ab is maximized, subject to the constraint that $a + b = n$. Using the AM-GM inequality, we see that

$$ab \leq \left(\frac{a+b}{2}\right)^2 = \frac{n^2}{4}.$$

Moreover, equality holds if and only if $a = b = n/2$. If n is odd, then we can't have $a = b = n/2$ if a and b are both integers; the product ab is maximized when $a = \lfloor n/2 \rfloor, b = \lceil n/2 \rceil$. But in any case, we find that

$$\mathrm{ex}(n, K_3) \geq \left\lfloor \frac{n^2}{4} \right\rfloor,$$

with the example of a K_3-free n-vertex graph with $\lfloor n^2/4 \rfloor$ edges given by the complete bipartite graph $K_{\lfloor n/2 \rfloor, \lceil n/2 \rceil}$.

Mantel's theorem says that this is the best we can do.

Theorem 2.1 (Mantel 1907)**.** $\mathrm{ex}(n, K_3) = \lfloor n^2/4 \rfloor$. *Moreover, the unique n-vertex triangle-free graph with $\lfloor n^2/4 \rfloor$ edges is $K_{\lfloor n/2 \rfloor, \lceil n/2 \rceil}$.*

We won't prove this right now. Instead, we'll first generalize Mantel's theorem, and then prove the generalization.

Almost 40 years after Mantel, Turán started thinking about similar questions, and it is thanks to his work that the field of extremal graph theory exists at all. Turán was studying what happens when, rather than excluding a triangle, we exclude some larger complete graph (also known as a *clique*). Namely, he was studying $\mathrm{ex}(n, K_r)$ for $r \geq 3$.

Again, there is a natural type of example we can come up with to lower-bound $\mathrm{ex}(n, K_r)$. Namely, let G be a complete $(r-1)$-partite graph on n vertices, namely a graph obtained by splitting the n vertices into $r - 1$ parts, then putting all edges between pairs of vertices in different parts and no edges within a part. Then G certainly will not have a copy of K_r: by the pigeonhole principle, if we take any r vertices in G, two of them must lie in the same part, and thus there cannot be an edge between them. Moreover, another simple application of the AM-GM inequality (or Jensen's inequality) shows that the way to do this in order to maximize the number of edges of G is to make all the parts have as equal sizes as possible, namely to make each part have size either $\lfloor n/(r-1) \rfloor$ or $\lceil n/(r-1) \rceil$. This motivates the following definition.

Definition 2.2. The *Turán graph* $T_{r-1}(n)$ is the n-vertex complete $(r-1)$-partite graph with all parts of size either $\lfloor n/(r-1) \rfloor$ or $\lceil n/(r-1) \rceil$. We denote its number of edges by

$$t_{r-1}(n) \coloneqq e(T_{r-1}(n)).$$

Remark. In case n is divisible by $r-1$, then every part of the Turán graph $T_{r-1}(n)$ has exactly $n/(r-1)$ vertices in each part, so

$$t_{r-1}(n) = \binom{r-1}{2} \cdot \left(\frac{n}{r-1}\right)^2 = \left(\frac{r-2}{r-1}\right)\frac{n^2}{2} = \left(1 - \frac{1}{r-1}\right)\frac{n^2}{2}.$$

In case n is not divisible by $r-1$, the formula is a little messier, involving the remainder of n when divided by $r-1$. However, we still have that for all fixed r and $n \to \infty$,

$$t_{r-1} = \left(1 - \frac{1}{r-1} + o(1)\right)\frac{n^2}{2},$$

where $o(1)$ represents a quantity that tends to 0 as n tends to infinity. In other words, if we fix $\varepsilon > 0$, then for any sufficiently large n, we have that

$$\left(1 - \frac{1}{r-1} - \varepsilon\right)\frac{n^2}{2} \le t_{r-1}(n) \le \left(1 - \frac{1}{r-1} + \varepsilon\right)\frac{n^2}{2}.$$

One other useful observation is that for any n, if we delete one vertex from each of the $r-1$ parts of $T_{r-1}(n)$, we obtain a copy of $T_{r-1}(n-r+1)$. Moreover, each non-deleted vertex is adjacent to exactly $r-2$ deleted vertices. So we delete $(r-2)(n-r+1) + \binom{r-1}{2}$ edges to obtain $T_{r-1}(n-r+1)$ from $T_{r-1}(n)$. This shows that

$$t_{r-1}(n) = t_{r-1}(n-r+1) + (r-2)(n-r+1) + \binom{r-1}{2}. \tag{1}$$

Note that $T_2(n) = K_{\lfloor n/2 \rfloor, \lceil n/2 \rceil}$, so Mantel's theorem can be rephrased as saying that $\mathrm{ex}(n, K_3) = t_2(n)$ and that $T_2(n)$ is the unique n-vertex K_3-free graph with $t_2(n)$ edges. Turán's theorem generalizes this to $\mathrm{ex}(n, K_r)$ for all $r \ge 3$.

Theorem 2.3 (Turán 1941)**.** *For every $r \ge 3$, we have $\mathrm{ex}(n, K_r) = t_{r-1}(n)$. Moreover, the unique n-vertex K_r-free graph with $t_{r-1}(n)$ edges is $T_{r-1}(n)$.*

Proof. We proceed by induction, with steps of size $r-1$. So we need $r-1$ base cases, corresponding to $n = 1, 2, \ldots, r-1$. But the theorem holds for such n, because for such n, any n-vertex graph has no K_r subgraph. So $\mathrm{ex}(n, K_r) = \binom{n}{2}$ for $1 \le n \le r-1$. Moreover, $T_{r-1}(n)$ is exactly K_n in these cases. This proves the base cases of the induction.

Now let $n > r-1$, and assume the theorem is true for $n-r+1$. Let G be an n-vertex graph with no copy of K_r and as many edges as possible. G must contain a copy of K_{r-1}, for otherwise we could add an edge and get a K_r-free graph with strictly more edges. Let K be some such K_{r-1} subgraph, and let $F \subseteq G$ be the subgraph obtained by deleting K. We know that $e(K) = \binom{r-1}{2}$. By induction, we know that

$$e(F) \le t_{r-1}(n-r+1).$$

Finally, each vertex of F cannot be adjacent to every vertex of K, for otherwise we would get a K_r. So the number of edges between F and K is at most $(r-2)(n-r+1)$. So

$$e(G) \le \binom{r-1}{2} + t_{r-1}(n-r+1) + (r-2)(n-r+1) = t_{r-1}(n),$$

by (1).

If $e(G) = t_{r-1}(n)$, then every inequality above must be an equality. In particular, the induction hypothesis implies that $F \cong T_{r-1}(n-r+1)$. Moreover, each vertex in F must be adjacent to exactly $r-2$ vertices in K, since we assume we have equality in the number of edges. Moreover, given two adjacnet vertices in F, they cannot be non-adjacent to the same vertex of K, for otherwise we could take the remaining $r-2$ vertices and these two to get a K_r. So this implies that each part of F is associated to exactly one missed vertex. So by adding this missed vertex to its part, we see that $G \cong T_{r-1}(n)$. □

On the homework, you will see many different proofs of Turán's theorem. It is one of those amazing mathematical theorems with dozens of different, and differently informative, proofs. It is also extremely useful, as you'll see on the homework!

Before moving on, let me just mention one convenient way to think about Turán's theorem is as follows. Note that

$$\binom{n}{2} = \frac{n^2}{2} - \frac{n}{2} = (1+o(1))\frac{n^2}{2}.$$

This shows that

$$t_{r-1}(n) = \left(1 - \frac{1}{r-1} + o(1)\right)\frac{n^2}{2} = \left(1 - \frac{1}{r-1} + o(1)\right)\binom{n}{2}.$$

Note that an n-vertex graph can have anywhere between 0 and $\binom{n}{2}$ edges. So Turán's theorem implies that a K_r-free n-vertex graph can have at most, asymptotically, a $1-1/(r-1)$ fraction of all possible edges.

3 Beyond Turán's theorem

Turán's theorem is great, and tells us exactly what $\mathrm{ex}(n, K_r)$ is for any r. But we started this class by asking about $\mathrm{ex}(n, H)$ for general H; what can we say about that? In general, we'd probably expect this problem to be really hard, and the answer should depend in complicated ways on the fixed graph H.

But it turns out that's not the case! Kind of amazingly, the answer depends, essentially, on a single parameter of the graph H—its chromatic number.

Theorem 3.1 (Erdős–Stone(–Simonovits) 1946 (1966)). *For any graph H,*

$$\mathrm{ex}(n, H) = \left(1 - \frac{1}{\chi(H)-1} + o(1)\right)\binom{n}{2}.$$

Remark.

- This is sometimes called the Fundamental Theorem of Extremal Graph Theory, for hopefully obvious reasons: it more or less completely resolves the main question that we started with.

- The history (and naming) of this theorem is a bit confusing. Erdős and Stone proved a special case of it in 1946. In 1966, Erdős and Simonovits realized that the special case actually implies (with a one-line implication) the general case, which had not been really studied before. We will soon see the special case, and how it implies the general case.

- Notice that if H is bipartite (i.e. if $\chi(H) = 2$), then $1 - 1/(\chi(H) - 1) = 0$. So the theorem simply says that if H is bipartite, then

  $$\text{ex}(n, H) = o(1) \cdot \binom{n}{2},$$

  which we usually write as $\text{ex}(n, H) = o(n^2)$. In other words, if G is an n-vertex graph containing no copy of some fixed bipartite graph H, then G must have *very few* edges—its number of edges grows sub-quadratically in n. Said differently, the fraction of all possible edges that we can put in such a graph is a vanishingly small fraction; the fraction tends to 0 as $n \to \infty$.

  Already this statement is far from obvious, and we'll soon prove it. In fact, as we'll see, proving the statement for bipartite H implies, in a certain sense, the full Erdős–Stone–Simonovits theorem.

Most of the rest of the class will be spent on proving the Erdős–Stone–Simonovits theorem. To do so, we'll prove upper and lower bounds on $\text{ex}(n, H)$ of the form $(1 - 1/(\chi(H) - 1) + o(1))\binom{n}{2}$. In fact, we can easily dispense with the lower bound.

Proposition 3.2. *For any fixed graph H and integer n,*

$$\text{ex}(n, H) \geq t_{\chi(H)-1}(n) = \left(1 - \frac{1}{\chi(H) - 1} + o(1)\right)\binom{n}{2}.$$

Proof. We claim that the Turán graph $T_{\chi(H)-1}(n)$ has no copy of H. Indeed, suppose we had some vertices in $T_{\chi(H)-1}(n)$ that defined a copy of H. Give the parts of $T_{\chi(H)-1}(n)$ names, say $V_1, \ldots, V_{\chi(H)-1}$. Then note that any two vertices of H that lie in the same part V_i cannot be adjacent in H, since $T_{\chi(H)-1}(n)$ has no edges inside any part V_i. Said differently, if we assign to any vertex v of H the number i so that $v \in V_i$, then two adjacent vertices are assigned different numbers. In other words, this yields a proper coloring of H with $\chi(H) - 1$ colors. But this contradicts the definition of the chromatic number. $\square$

4 Extremal numbers of bipartite graphs

4.1 Upper bounds

Let H be a bipartite graph. Recall that the Erdős–Stone–Simonovits theorem implies that in this case, $\mathrm{ex}(n, H) = o(n^2)$, or equivalently that

$$\lim_{n\to\infty} \frac{\mathrm{ex}(n, H)}{n^2} = 0.$$

This is pretty surprising! For example, the complete bipartite graph $K_{\lfloor n/2\rfloor,\lceil n/2\rceil}$ has $\lfloor n^2/4\rfloor$ edges and no copy of any odd cycle $C_{2\ell+1}$. Thus, $\mathrm{ex}(n, C_{2\ell+1}) \geq \lfloor n^2/4\rfloor$ for all ℓ. But the four-cycle (or any other even cycle) is bipartite, so the Erdős–Stone–Simonovits theorem implies that $\mathrm{ex}(n, C_4) = o(n^2)$. What's up with that?

We will shortly prove that in fact, $\mathrm{ex}(n, C_4) \leq O(n^{3/2})$. In case you haven't seen it before, the big-O notation means that $\mathrm{ex}(n, C_4) \leq Cn^{3/2}$ for some absolute constant C, which we won't specify. In other words, we will shortly prove that if G is an n-vertex graph with at least $Cn^{3/2}$ edges, then G has a copy of C_4, assuming C is an appropriately chosen constant. As a warm-up, we will begin with an easier special case of this result, which is the case when G is d-regular (i.e. every vertex in G has degree d). Recall that in any graph, the sum of the degrees of all the vertices equals twice the number of edges, so if G is d-regular then it has $dn/2$ edges. Thus, if G has $Cn^{3/2}$ edges and is d-regular, then $d = 2C\sqrt{n}$.

Proposition 4.1. *Let G be a d-regular n-vertex graph. If $d \geq 2\sqrt{n}$, then G contains a copy of C_4.*

Proof. Suppose for contradiction that G is C_4-free. We count the number of copies of $K_{1,2}$ in G, where $K_{1,2} =$ ●—●—● consists of one central vertex adjacent to two outer vertices. On the one hand, if we sum over all possibilities for the central vertex, we see that

$$\#(K_{1,2} \text{ in } G) = \sum_{v\in V(G)} \#(K_{1,2} \text{ with central vertex } v) = \sum_{v\in V(G)} \binom{\deg(v)}{2} = n\binom{d}{2}.$$

On the other hand, suppose we fix some $u, w \in V(G)$. We claim that they can be the outer vertices of at most one copy of $K_{1,2}$. Indeed, if not, then we would have two $K_{1,2}$s agreeing on the outer vertices, which yields a copy of C_4, a contradiction. So we conclude that

$$\#(K_{1,2} \text{ in } G) = \sum_{\substack{u,w\in V(G)\\ \text{distinct}}} \#(K_{1,2} \text{ with outer vertices } u, w) \leq \sum_{\substack{u,w\in V(G)\\ \text{distinct}}} 1 = \binom{n}{2}.$$

Rearranging, we see that

$$n\binom{d}{2} \leq \binom{n}{2} \qquad\Longleftrightarrow\qquad \binom{d}{2} \leq \frac{n-1}{2} \qquad\Longleftrightarrow\qquad d(d-1) \leq n-1.$$

But if $d \geq 2\sqrt{n}$ and $n \geq 0$, then this is a contradiction. $\square$

To prove the real result, we will need one extraordinarily useful analytic tool, called *Jensen's inequality.* We will actually only need the following special case. For a real number x and a positive integer r, we extend the definition of the binomial coefficient as

$$\binom{x}{r} = \frac{x(x-1)(x-2)\cdots(x-r+1)}{r!}.$$

Lemma 4.2 ((Consequence of) Jensen's inequality)**.** *Let $r \geq 1$ be a positive integer, and let $x_1, \dots, x_n$ be non-negative integers. Suppose that $\frac{1}{n}\sum_{i=1}^n x_i \geq r$. Then*

$$\sum_{i=1}^n \binom{x_i}{r} \geq n\binom{\frac{1}{n}\sum_{i=1}^n x_i}{r}.$$

The point of this is that if we add up terms of the form $\binom{x_i}{r}$, we can only decrease the sum if we replace each x_i by the *average* of all the x_i. One says that the function $x \mapsto \binom{x}{r}$ is *convex*: the sum of its values is minimized when all the variables are equal (to their average).

We won't prove Jensen's inequality in class, but its proof is on the homework if you're interested. Once we have Jensen's inequality, we can easily prove the full result that $\mathrm{ex}(n, H) \leq O(n^{3/2})$. In fact, we will prove the following much more general result.

Theorem 4.3 (Kővári–Sós–Turán 1954)**.** *For positive integers $s \leq t$, we have*

$$\mathrm{ex}(n, K_{s,t}) \leq O(n^{2-1/s}).$$

Here, the implicit constant may depend on s and t (which we think of as fixed).

Proof. We proceed much as in the proof of Proposition 4.1. Let G be an n-vertex graph with at least $Cn^{2-1/s}$ edges, where C is some large constant we will pick later. Let d be the average degree in G, so that $d = \frac{2}{n}e(G) \geq 2Cn^{1-1/s}$. Suppose for contradiction that G is $K_{s,t}$-free. We count the number of copies of $K_{1,s}$ in G in two ways. First, by summing over the options for the central vertex, we have that

$$\#(K_{1,s} \text{ in } G) = \sum_{v \in V(G)} \binom{\deg(v)}{s} \geq n\binom{d}{s},$$

using Lemma 4.2, as well as the fact that $d \geq s$ by picking C sufficiently large. On the other hand, by counting over the s outer vertices of $K_{1,s}$, we have that every $u_1, \dots, u_s \in V(G)$ can be the outer vertices of at most $t-1$ copies of $K_{1,s}$. So

$$\#(K_{1,s} \text{ in } G) \leq \sum_{\substack{u_1,\dots,u_s \in V(G) \\ \text{distinct}}} (t-1) = (t-1)\binom{n}{s}.$$

Combining these, we see that

$$(t-1)\binom{n}{s} \geq n\binom{d}{s} \qquad \iff \qquad (t-1)(n-1)(n-2)\cdots(n-s+1) \geq d(d-1)\cdots(d-s+1).$$

Now, if n is very large (which is the regime we care about anyway), all this subtracting stuff doesn't matter. So this is roughly equivalent to

$$(t-1)n^{s-1} \geq d^s \qquad \Longleftrightarrow \qquad d \leq (t-1)^{1/s} n^{1-1/s}.$$

If C is sufficiently large, then this is a contradiction. Moreover, if C is sufficiently large, then the slightly sketchy step above where we dropped the subtractions is also OK, and we get the desired contradiction. $\square$

Note that $C_4 = K_{2,2}$, so in the case $s = t = 2$, we indeed get the claimed bound of $\mathrm{ex}(n, C_4) \leq O(n^{3/2})$.

There are a number of important consequences of the Kővári–Sós–Turán theorem. The first is that it immediately gives us a bound on $\mathrm{ex}(n, H)$ for all bipartite H. Indeed, note that if H_1 is a subgraph of H_2, then

$$\mathrm{ex}(n, H_1) \leq \mathrm{ex}(n, H_2)$$

for all n, as any H_1-free graph is also H_2-free. Now, if H is a bipartite graph, then H is a subgraph of $K_{s,t}$ for some $s \leq t$. So

$$\mathrm{ex}(n, H) \leq \mathrm{ex}(n, K_{s,t}) \leq O(n^{2-1/s}).$$

In particular, this proves that $\mathrm{ex}(n, H) = o(n^2) = o(1) \cdot \binom{n}{2}$ for bipartite H. Indeed,

$$\lim_{n\to\infty} \frac{\mathrm{ex}(n, H)}{n^2} \leq \lim_{n\to\infty} \frac{O(n^{2-1/s})}{n^2} = \lim_{n\to\infty} O(n^{-1/s}) = 0.$$

Recall that this was a consequence of the Erdős–Stone–Simonovits theorem.

4.2 Lower bounds

How good are the upper bounds we proved? Let's begin with the one we started with, $\mathrm{ex}(n, C_4) \leq O(n^{3/2})$. Can we construct an n-vertex C_4-free graph with roughly that many edges?

As it turns out, we can! The following construction was originally due to Eszter Klein in 1938 (as reported in a paper of Erdős). Note that this is before Turán's theorem, so before the birth of extremal graph theory! As such, no one really appreciated what this construction was or meant, and it was later rediscovered by Erdős, Rényi, and Sós (and independently Brown). These days, it is often called the "Erdős–Rényi" construction, which I find a little odd, both because they weren't the first to discover it, and because there are many other things named after Erdős and Rényi.

Theorem 4.4 (Klein 1938)**.** *For every $n \geq 1$, there is an n-vertex C_4-free graph with at least $n^{3/2}/64$ edges.*

Proof. First, suppose that $n = 2p^2$ for some prime p; we will later get rid of this assumption. Consider the integers mod p, which form a field that we denote $\mathbb{F}_p$. (If you don't know what the word "field" means, just believe me that among the integers mod p, we can use addition, multiplication, and division and have them work basically the same way they do in $\mathbb{R}$.)

Let $\mathbb{F}_p^2$ denote the two-dimensional plane over $\mathbb{F}_p$, i.e. the set of points (x, y) with $x, y \in \mathbb{F}_p$. For $m, b \in \mathbb{F}_p$, let $\ell_{m,b}$ denote the line $y = mx + b$ in $\mathbb{F}_p^2$. In other words, $\ell_{m,b}$ is the set of points $(x, y) \in \mathbb{F}_p^2$ satisfying $y = mx + b$.

We define a bipartite graph G with parts P, L, where $P = \mathbb{F}_p^2$ and $L = \{\ell_{m,b} : m, b \in \mathbb{F}_p\}$. The edges of G are given by *incidence*: we connect $(x, y) \in P$ to $\ell_{m,b} \in L$ if and only if (x, y) lies on the line $\ell_{m,b}$, i.e. if and only if $y = mx + b$.

Note that $|P| = |L| = p^2$, so G has $n = 2p^2$ vertices. Moreover, every line $\ell_{m,b} \in L$ has exactly p points on it, so every vertex in L has degree p in G. Therefore, $e(G) = p|L| = p^3 = (n/2)^{3/2}$.

Finally, we claim that G is C_4-free. To see this, note that G is bipartite, so the only way we could have a copy of C_4 in G is to have distinct $p_1, p_2 \in P$ and distinct $\ell_1\ell_2 \in L$ so that $p_1\ell_1p_2\ell_2$ forms a 4-cycle. But this means that p_1 lies on the lines ℓ_1, ℓ_2, and that p_2 also lies on both these lines. So we have two lines which intersect at two distinct points!

Using our intuition from $\mathbb{R}$, we expect this to be impossible, and it's impossible over $\mathbb{F}_p$ as well. Formally, let $p_1 = (x_1, y_1)$ and $p_2 = (x_2, y_2)$, and $\ell_1 = \ell_{m_1,b_1}, \ell_2 = \ell_{m_2,b_2}$. Then we have the equations

$$\begin{aligned} y_1 &= m_1x_1 + b_1 & \qquad y_2 &= m_1x_2 + b_1 \\ y_1 &= m_2x_1 + b_2 & \qquad y_2 &= m_2x_2 + b_2 \end{aligned}$$

Rearranging the first column, we see that $m_1x_1 + b_1 = m_2x_1 + b_2$, or equivalently that $(m_2 - m_1)x_1 = b_1 - b_2$. If $m_1 = m_2$ then this implies that $b_1 = b_2$, contradicting that ℓ_1, ℓ_2 are distinct. So we have that $m_1 \neq m_2$, so $x_1 = (b_1 - b_2)/(m_2 - m_1)$. But from the second column of equations, we conclude that $x_2 = (b_1 - b_2)/(m_2 - m_1)$ as well, so $x_1 = x_2$. But if we plug this into any of the equations, we conclude that $y_1 = y_2$, and thus that $p_1 = p_2$, a contradiction. So G is C_4-free.

The only remaining thing is to deal with the fact that n need not equal twice the square of a prime. So let n be arbitrary. There is an important result in number theory, called Bertrand's postulate, which says that there is always a prime between m and $2m$ for all positive integers m. Let $m = \lfloor \sqrt{n}/4 \rfloor$, and let p be a prime between m and $2m$, so that $n/8 \leq 2p^2 \leq n$. Using the construction above, we obtain a C_4-free graph G on $2p^2$ vertices with p^3 edges. We add to this graph $n - 2p^2$ isolated vertices, and we obtain a new C_4-free n-vertex graph with $p^3 \geq (n/16)^{3/2} = n^{3/2}/64$ edges. □

Using these finite fields and finite geometries might seem like a neat trick, but it turns out that it's essentially the only thing one can do. Indeed, all constructions we know of for C_4-free graphs with many edges use such techniques. Moreover, there is a powerful result of Füredi, which says that for those n for which such a construction (appropriately defined) exists, the *unique* C_4-free n-vertex graph with the most edges comes from such a construction.

So we conclude that $\text{ex}(n, C_4) = \Theta(n^{3/2})$, where the big-$\Theta$ means that we have upper and lower bounds that agree up to a constant factor. Since $\text{ex}(n, K_{2,t}) \geq \text{ex}(n, C_4)$ for all $t \geq 2$, we conclude that $\text{ex}(n, K_{2,t}) = \Theta(n^{3/2})$ for all $t \geq 2$.

What about $\text{ex}(n, K_{3,3})$? We proved in Theorem 4.3 that $\text{ex}(n, K_{3,3}) \leq O(n^{5/3})$. As it turns out, this is also tight.

Theorem 4.5 (Brown 1966)**.** *For every n, there exists an n-vertex $K_{3,3}$-free graph G with $n^{5/3}/100$ edges.*

Proof sketch. I won't present the proof in detail, but will explain the big idea. Suppose that $n = p^3$. Construct a graph G with vertex set $\mathbb{F}_p^3$, where we connect two vertices (x, y, z) and (x', y', z') by an edge if and only if

$$(x - x')^2 + (y - y')^2 + (z - z')^2 = 1.$$

In other words, the neighborhood of any vertex looks like a "unit sphere" centered at that vertex, except that "spheres" don't really exist over $\mathbb{F}_p$.

Nonetheless, if we were working in $\mathbb{R}^3$, then we'd expect that any three unit spheres can intersect in at most two points: two unit spheres can intersect in a circle, and that circle can intersect a thid unit sphere in only two points. So we'd expect that G is $K_{3,3}$-free, since any three vertices have at most two common neighbors.

Since we expect a sphere to be "two-dimensional", we should expect every unit sphere to have roughly p^2 points on it, and this turns out to be true. So G has $n = p^3$ vertices, and every vertex has degree around p^2, so we expect $e(G) \approx p^5 = n^{5/3}$.

All of this intuition can be made precise, some of it with some annoyance. For example, it turns out that this only really works if $p \equiv 3 \pmod 4$. But the high-level idea is correct. $\square$

So we see that the Kővári–Sós–Turán theorem is best possible (up to the constant factor) for $s = 2$ and $s = 3$. The case of $s = 1$ is much easier, but it's also best possible there, as you saw on the homework. So it is natural to conjecture, as many have done, that the Kővári–Sós–Turán theorem is tight in general.

Conjecture 4.6 (Many people)**.** *For all $s \leq t$,*

$$\text{ex}(n, K_{s,t}) = \Theta(n^{2-1/s}).$$

Moreover, based on what I've told you so far, it is natural to expect that not only is this conjecture proved, but that the constructions look kind of the same as above. You work with the s-dimensional space $\mathbb{F}_p^s$ over the field $\mathbb{F}_p$, and use some kind of cleverly chosen polynomial or set of polynomial equations to define the adjacency condition. However, despite many people having this same idea, Conjecture 4.6 remains unproved. Moreover, many experts in the field now even question whether it is true.

Nonetheless, some other things are known about $\text{ex}(n, K_{s,t})$. Namely, it is known that $\text{ex}(n, K_{s,t}) = \Theta(n^{2-1/s})$ if t is sufficiently large compared to s. The first result of this type is due to Kollár, Rónyai, and Szabó in 1996, who proved that

$$\text{ex}(n, K_{s,t}) = \Theta(n^{2-1/s}) \qquad \text{if } t > s!.$$

To do this, they constructed a $K_{s,t}$-free graph, again using the space F_p^s, called the *norm graph*. Their construction was later modified by Alon, Rónyai, and Szabó in 1999, who defined a similar graph called the *projective norm graph* (again over $\mathbb{F}_p^s$), which implies that

$$\mathrm{ex}(n, K_{s,t}) = \Theta(n^{2-1/s}) \qquad \text{if } t > (s-1)!.$$

So, for example, we know that $\mathrm{ex}(n, K_{4,7}) = \Theta(n^{7/4})$, but have no such lower bound for $\mathrm{ex}(n, K_{4,4})$.

For about 20 years, the Alon–Rónayi–Szabó result was the best known. But very recently, Bukh proved the following theorem.

Theorem 4.7 (Bukh 2021)**.** *Suppose $s \geq 2$ and $t \geq 9^s \cdot s^{4s^{2/3}}$ are integers. Then*

$$\mathrm{ex}(n, K_{s,t}) = \Theta(n^{2-1/s}).$$

The key point is that for large s, the previous bound on t, namely $(s-1)!$, grew super-exponentially in s. But Bukh's bound, for large s, grows merely exponentially in s. The key to Bukh's construction is again to work over $\mathbb{F}_p^s$, but not to pick a *clever* polynomial. Instead, he picks a *random* polynomial, and then uses arguments from probability, combinatorics, and algebraic geometry to prove that the resulting graph is $K_{s,t}$-free with positive probability.

4.3 Non-complete bipartite graphs

So far, we have focused on complete bipartite graphs $K_{s,t}$. But it is natural to ask about $\mathrm{ex}(n, H)$ for general bipartite graphs H. What can we say about this question?

The answer is, not very much. We already saw a general-purpose upper bound, namely $\mathrm{ex}(n, H) \leq \mathrm{ex}(n, K_{s,t}) \leq O(n^{2-1/s})$ if H is a subgraph of $K_{s,t}$. But for many specific bipartite graphs, much better upper bounds are known, using arguments specific to the graph at hand. For example, the following is known about the extremal numbers of even cycles.

Theorem 4.8 (Erdős (unpublished), Bondy–Simonovits 1974)**.** *For every $\ell \geq 2$, we have*

$$\mathrm{ex}(n, C_{2\ell}) \leq O\left(n^{1+1/\ell}\right).$$

Note that in case $\ell = 2$, this matches the bound $\mathrm{ex}(n, C_4) \leq O(n^{3/2})$ that we saw earlier. It is again widely conjectured that this bound is tight, but it is only known to be tight in case $\ell \in \{2, 3, 5\}$. This is pretty remarkable: we know that

$$\mathrm{ex}(n, C_4) = \Theta(n^{3/2}) \qquad \mathrm{ex}(n, C_6) = \Theta(n^{4/3}) \qquad \mathrm{ex}(n, C_{10}) = \Theta(n^{6/5})$$

but we have no idea what the value of $\mathrm{ex}(n, C_8)$ is!

There is a general-purpose lower bound that is known. In general, algebraic techniques like the ones described above are the best techniques we have for constructing lower bounds, but they often rely on specific structures that we can exploit. The following bound holds for any bipartite graph. We use the notation big-Ω to denote the opposite of big-O—it means that the left-hand side is at least as large as the right-hand side, up to a constant factor.

Given a graph H, we define its *2-density* to be

$$m_2(H) := \max_{F \subseteq H} \frac{e(F) - 1}{v(F) - 2}.$$

Theorem 4.9. *For any bipartite H, we have*

$$\mathrm{ex}(n, H) \geq \Omega(n^{2-1/m_2(H)}).$$

The proof of Theorem 4.9 uses the probabilistic method, and I won't cover it in class. But at a high level, the idea is to pick a *random* graph G with n vertices and roughly $n^{2-1/m_2(H)}$ edges. One can then show that with positive probability, the number of copies of H in G is less than half the number of edges of G. By deleting a single edge from each copy of H, we obtain a graph with half as many edges—so still $\Omega(n^{2-1/m_2(H)})$—and no copy of H.

To end this section, let me just mention two remarkable conjectures of Erdős and Simonovits, which roughly say that the behavior of $\mathrm{ex}(n, H)$ for general bipartite H is very complicated.

Conjecture 4.10 (Erdős–Simonovits rational exponents conjecture)**.** *For every bipartite H, there exists some rational number $\alpha \in [1, 2)$ so that*

$$\mathrm{ex}(n, H) = \Theta(n^{\alpha}).$$

Moreover, the converse holds: for every rational $\alpha \in [1, 2)$, there exists some bipartite graph H so that

$$\mathrm{ex}(n, H) = \Theta(n^\alpha)$$

The first part of this conjecture is doubted by some experts, though no one has any idea how to prove or disprove it. However, the second part of the conjecture—that there exists a graph for any rational α—is widely believed to be true, and we are in fact fairly close to proving it. Every few months, a new paper appears finding a new infinite set of rational numbers that are now known to be "achievable", i.e. to be the exponent of $\mathrm{ex}(n, H)$ for some bipartite H.

Moreover, a slight weakening of the second part of the conjecture was recently proved by Bukh and Conlon. Recall from the homework that if $\mathcal{H}$ is a collection of graphs, then we say that G is $\mathcal{H}$-free if G contains no copy of any $H \in \mathcal{H}$, and we write

$$\mathrm{ex}(n, \mathcal{H}) = \max\{e(G) : G \text{ is an } n\text{-vertex } \mathcal{H}\text{-free graph}\}.$$

Theorem 4.11 (Bukh–Conlon 2018)**.** *For every rational $\alpha \in [1, 2)$, there exists some finite collection $\mathcal{H}$ of bipartite graphs for which*

$$\mathrm{ex}(n, \mathcal{H}) = \Theta(n^\alpha).$$

5 Extremal numbers of hypergraphs

It's time for everything to get more *hyper*.

If we go back to bare basics, a graph is a collection V of vertices, plus a collection E of edges, which are simply unordered pairs of vertices. Why restrict ourselves to pairs?

Definition 5.1. A *k-uniform hypergraph* (sometimes called an k-graph for short) consists of a finite collection V of *vertices*, as well as a collection E of *k-uniform hyperedges*, which are simply subsets of V of size k.

As with graphs, we say that one k-graph $\mathcal{H}$ is a *subhypergraph* (or simply *subgraph*) of another k-graph $\mathcal{G}$ if we can obtain $\mathcal{H}$ from $\mathcal{G}$ by deleting some vertices and edges. We say that $\mathcal{G}$ is *$\mathcal{H}$-free* if $\mathcal{G}$ does not contain $\mathcal{H}$ as a subgraph (and we also say that $\mathcal{G}$ has no copy of $\mathcal{H}$).

As with graphs, we define the *extremal number* of $\mathcal{H}$ as

$$\mathrm{ex}(n, \mathcal{H}) = \max\{e(\mathcal{G}) : \mathcal{G} \text{ is an } n\text{-vertex } \mathcal{H}\text{-free } k\text{-graph}\}.$$

In contrast to graphs (the case $k = 2$), we know extraordinarily little about $\mathrm{ex}(n, \mathcal{H})$ for k-graphs $\mathcal{H}$ with $k \geq 3$. For example, even the hypergraph analogue of Mantel's theorem is a famous open problem. To explain this formally, we make the following definition.

Definition 5.2. The *complete k-graph* on r vertices, denoted $K_r^{(k)}$, is the k-graph with r vertices whose edge set consists of all subsets of size k.

Then the amazing fact is that for any $k > r \geq 3$, we do not know the value of $\mathrm{ex}(n, K_r^{(k)})$. For literally no pair of (r,k)! This problem was proposed by Turán already in 1941, and he made the following conjecture, which is a natural analogue of Mantel's theorem.

Conjecture 5.3 (Turán 1941)**.**

$$\mathrm{ex}(n, K_4^{(3)}) = \left(\frac{5}{9} + o(1)\right)\binom{n}{3}.$$

The reason for 5/9 is a specific construction of an n-vertex $K_4^{(3)}$-free 3-graph, which Turán came up with, and which was the best he could come up with. You'll see Turán's construction on the homework.

Erdős offered \$500 for the resolution of Conjecture 5.3, and \$1000 for a general formula for $\mathrm{ex}(n, K_r^{(k)})$. So far, very little progress has been made on these questions. The best known bound for $\mathrm{ex}(n, K_4^{(3)})$ is due to Razborov, who proved that

$$\mathrm{ex}(n, K_4^{(3)}) \leq (0.561666 + o(1))\binom{n}{3}.$$

Note that $5/9 = 0.555\ldots$, so this is pretty close to Turán's conjecture. Unfortunately, Razborov's technique is unlikely to yield the full resolution of Conjecture 5.3, because his technique uses a computer to do complicated computations to what is essentially a "finite approximation" to the problem.

In general, the best known lower bound is due to de Caen, who proved that

$$\mathrm{ex}(n, K_r^{(k)}) \leq \left(1 - \frac{1}{\binom{r-1}{k-1}} + o(1)\right)\binom{n}{k}.$$

The best known general lower bound, due to Sidorenko, is

$$\mathrm{ex}(n, K_r^{(k)}) \geq \left(1 - \left(\frac{k-1}{r-1}\right)^{k-1} + o(1)\right)\binom{n}{k}.$$

In the case of $k = 3$, this says that

$$\mathrm{ex}(n, K_r^{(3)}) \geq \left(1 - \left(\frac{2}{r-1}\right)^{2} + o(1)\right)\binom{n}{3},$$

and this was conjectured to be optimal by Turán. You'll see the construction in the homework.

Despite not knowing the hypergraph analogues of Mantel's or Turán's theorems, the hypergraph analogue of the Kővári–Sós–Turán theorem *is* known, and is due to Erdős (1965). To state this, we need to define the hypergraph analogue of a bipartite graph.

Definition 5.4. A k-graph $\mathcal{H}$ is called *k-partite* if its vertex set can be split into k parts, so that every hyperedge of $\mathcal{H}$ contains exactly one vertex from each part.

The *complete k-partite k-graph with parts of sizes* $s_1, \ldots, s_k$, denoted $K^{(k)}_{s_1,\ldots,s_k}$, is the k-graph with parts of sizes $s_1, \ldots, s_k$, containing every edge with exactly one vertex from each part.

Note that in case $k = 2$, this simply recovers the definition of a bipartite graph and a complete bipartite graph. Because of this, the following result generalizes the Kővári–Sós–Turán theorem.

Theorem 5.5 (Erdős 1965)**.** *Let $s_1 \leq \cdots \leq s_k$ be positive integers. Then*

$$\mathrm{ex}(n, K^{(k)}_{s_1,\ldots,s_k}) \leq O\left(n^{k-\frac{1}{s_1 s_2 \cdots s_{k-1}}}\right).$$

The most important thing here is that any n-vertex k-graph has at most $\binom{n}{k} = \Theta(n^k)$ hyperedges. So this upper bound has a smaller exponent on n. This implies that if $\mathcal{H}$ is *any* k-partite k-graph, then

$$\mathrm{ex}(n, \mathcal{H}) = o(1) \cdot \binom{n}{k}.$$

The upper bound in Theorem 4.3 is still the best upper bound we have on extremal numbers of k-partite k-graphs. Moreover, as in the case of graphs, it is known that the bound in Theorem 5.5 is best possible if s_k is sufficiently large with respect to $s_1, \ldots, s_{k-1}$.

The proof of Theorem 5.5 very similar to that of Theorem 4.3, except that we combine it with an induction on k. To keep the notation from getting too crazy, we will only prove it in the case $k = 3$, which we will derive from the case $k = 2$, i.e. the Kővári–Sós–Turán theorem. Also, we will only prove it in the case $s_1 = s_2 = s_3 = s$, i.e. we will prove that

$$\mathrm{ex}(n, K^{(3)}_{s,s,s}) \leq O(n^{3-1/s^2}). \tag{2}$$

Hopefully you'll believe me (or convince yourself that it's true if you don't!) that the general result follows from the same technique, just with more bookkeeping.

Proof of (2). Let $\mathcal{G}$ be an n-vertex 3-graph with at least Cn^{3-1/s^2} hyperedges, for some constant $C > 0$ we will pick later. Suppose for contradiction that $\mathcal{G}$ is $K^{(3)}_{s,s,s}$-free. Let X be the number of copies of $K^{(3)}_{1,1,s}$ in $\mathcal{G}$. We bound X in two ways.

First, for a pair of distinct vertices v, w, let $\mathrm{codeg}(v, w)$ denote the number of hyperedges containing both v and w. Then we first claim that

$$\sum_{\substack{v,w \in V(\mathcal{G}) \\ \text{distinct}}} \mathrm{codeg}(v, w) = 3e(\mathcal{G}).$$

This is true for the same reason that the sum of the degrees in a graph equals twice the number of edges. Namely, every hyperedge of G appears exactly three times in the sum on the left-hand side.

Using this, we see that by Jensen's inequality,

$$X = \sum_{\substack{v,w \in V(\mathcal{G}) \\ \text{distinct}}} \binom{\operatorname{codeg}(v,w)}{s} \geq \binom{n}{2} \binom{\frac{1}{\binom{n}{2}} \sum_{v,w} \operatorname{codeg}(v,w)}{s} = \binom{n}{2} \binom{3e(\mathcal{G})/\binom{n}{2}}{s}.$$

Note too that since $e(\mathcal{G}) \geq Cn^{3-1/s^2}$, we have that $3e(\mathcal{G})/\binom{n}{2} \geq \Omega(n^{1-1/s^2})$. This implies that

$$X \geq \binom{n}{2} \binom{3e(\mathcal{G})/\binom{n}{2}}{s} \geq cn^2 \cdot (Cn^{1-1/s^2})^s = cC^s n^{2+s-1/s}$$

for some absolute constant $c > 0$, depending only on s.

On the other hand, we may upper-bound X by counting over s-sets of vertices which can be the "outer" vertices of the $K^{(3)}_{1,1,s}$. Namely, fix distinct $u_1, \ldots, u_s \in V(\mathcal{G})$. We define a new graph (note: not a hypergraph, a *graph*) $G(u_1, \ldots, u_s)$ as follows. The vertex set of $G(u_1, \ldots, u_s)$ is $V(\mathcal{G}) \setminus \{u_1, \ldots, u_s\}$. Moreover, given $v, w \in V(\mathcal{G}) \setminus \{u_1, \ldots, u_s\}$, we make vw an edge of $G(u_1, \ldots, u_s)$ if and only if $\{v, w, u_1, \ldots, u_s\}$ form a copy of $K^{(3)}_{1,1,s}$.

Now, for every choice of $u_1, \ldots, u_s$, we claim that $G(u_1, \ldots, u_s)$ is a $K_{s,s}$-free graph. Indeed, if we had a copy of $K_{s,s}$ in $G(u_1, \ldots, u_s)$, then we would find a copy of $K^{(3)}_{s,s,s}$ in $\mathcal{G}$, which is a contradiction. So by the Kővári–Sós–Turán theorem, we know that

$$e(G(u_1, \ldots, u_s)) \leq O(n^{2-1/s})$$

for every choice of distinct $u_1, \ldots, u_s \in V(\mathcal{G})$.

We can use this to upper-bound X, as follows. Note that $e(G(u_1, \ldots, u_s))$ is precisely the number of copies of $K^{(3)}_{1,1,s}$ that have $u_1, \ldots, u_s$ as the outer vertices. This implies that

$$X = \sum_{\substack{u_1, \ldots, u_s \in V(G) \\ \text{distinct}}} e(G(u_1, \ldots, u_s)) \leq \sum_{\substack{u_1, \ldots, u_s \in V(G) \\ \text{distinct}}} O(n^{2-1/s}) = \binom{n}{s} \cdot O(n^{2-1/s}) = O(n^{2+s-1/s}).$$

Combining our upper and lower bounds on X, we see that

$$cC^s n^{2+s-1/s} \leq O(n^{2+s-1/s}),$$

where both c and the implicit constant in the big-O depend only on s. Thus, if we pick C sufficiently large, this is a contradiction, and we conclude that $\mathcal{G}$ has a copy of $K^{(3)}_{s,s,s}$. $\square$

6 Supersaturation

In this section, we discuss a special case of a very important phenomenon in extremal combinatorics, known as *supersaturation*. Roughly speaking, extremal combinatorics proves results of the type "if some discrete structure is sufficiently 'large', then it contains at least one copy of some other structure". The example we've seen of this is Turán's theorem: if a graph (discrete structure) has sufficiently many edges (is large) then it contains a K_k subgraph (a copy of some other structure). Supersaturation, in general, boosts this to a statement of the type "if the discrete structure is just a bit larger, then it contains *very many* copies of the other structure". Specifically, we'll prove the following supersaturation version of Turán's theorem. It was first explicitly stated by Erdős and Simonovits in 1983, but it can implicitly be found in earlier works, e.g. of Erdős from 1971.

Theorem 6.1. *For every integer $k \geq 3$ and every real number $\varepsilon > 0$, there exists some $\delta > 0$ so that the following holds for all sufficiently large n. If G is an n-vertex graph with*

$$e(G) \geq \left(1 - \frac{1}{k-1} + \varepsilon\right)\binom{n}{2},$$

then G contains at least $\delta\binom{n}{k}$ copies of K_k.

Note that G has *at most* $\binom{n}{k}$ copies of K_k, so this theorem is pretty remarkable: it says that once we have just barely more edges than the Turán graph, we have not only one copy of K_k, but a constant proportion of *all possible* copies of K_k. To prove this theorem, we need the following useful lemma, which is stated in greater generality than we need.

For a graph G and a subset $M \subseteq V(G)$, we denote by $e(M)$ the number of edges entirely contained in M, or equivalently the number of edges in the induced subgraph $G[M]$.

Lemma 6.2. *Let $0 < \alpha < \beta < 1$ be real numbers, let $m \geq 2$ be an integer, and let G be an n-vertex graph with $n \geq m$. Assume that $e(G) \geq \beta\binom{n}{2}$. Then the number of sets $M \subseteq V(G)$ with $|M| = m$ and $e(M) \geq \alpha\binom{m}{2}$ is at least $(\beta - \alpha)\binom{n}{m}$.*

Proof. The key identity which underlies this proof is

$$\binom{n-2}{m-2}e(G) = \sum_{\substack{M\subseteq V(G)\\ |M|=m}} e(M).$$

This has a simple bijective proof. On the right-hand side, every edge uv is counted a number of times, and that number of times is simply the number of m-sets M which contain both u and v. But the number of such m-sets is exactly $\binom{n-2}{m-2}$, yielding the formula.

Now, let $\mathcal{M}_0$ denote the set of M with $e(M) < \alpha\binom{m}{2}$, and let $\mathcal{M}_1$ denote the set of M with $e(M) \geq \alpha\binom{m}{2}$. So our goal is to prove a lower bound on $|\mathcal{M}_1|$. Continuing the identity above, we can write

$$\begin{aligned}\binom{n-2}{m-2}e(G) &= \sum_{M\in\mathcal{M}_0} e(M) + \sum_{M\in\mathcal{M}_1} e(M)\\ &\leq \sum_{M\in\mathcal{M}_0} \alpha\binom{m}{2} + \sum_{M\in\mathcal{M}_1}\binom{m}{2}\\ &= \binom{m}{2}\left(\alpha|\mathcal{M}_0| + |\mathcal{M}_1|\right)\end{aligned}$$

since every m-set in $\mathcal{M}_0$ has at most $\alpha\binom{m}{2}$ edges, and every m-set in $\mathcal{M}_1$ has at most $\binom{m}{2}$ edges.

Note that $|\mathcal{M}_0| + |\mathcal{M}_1| = \binom{n}{m}$. Let $x = |\mathcal{M}_1|/\binom{n}{m}$, so that $1-x = |\mathcal{M}_0|/\binom{n}{m}$. Dividing by $\binom{n}{m}\binom{m}{2}$, the above inequality yields

$$\frac{\binom{n-2}{m-2}}{\binom{n}{m}\binom{m}{2}}e(G) \leq \alpha(1-x)+x = \alpha+(1-\alpha)x.$$

Now, we recall that $e(G) \geq \beta\binom{n}{2}$, so

$$\frac{\binom{n-2}{m-2}}{\binom{n}{m}\binom{m}{2}}e(G) \geq \frac{\binom{n-2}{m-2}\binom{n}{2}}{\binom{n}{m}\binom{m}{2}}\beta.$$

The final step is another magic identity, which is that $\binom{n-2}{m-2}\binom{n}{2} = \binom{n}{m}\binom{m}{2}$; in other words, the complicated fraction above is simply equal to 1. Indeed, both sides of this identity count the same object, which is the number of ways of picking an m-set out of n objects, and then picking 2 objects from the m-set.

Combining all these inequalities, we find that

$$\beta \leq \alpha + (1-\alpha)x \qquad\iff\qquad x \geq \frac{\beta-\alpha}{1-\alpha},$$

which implies that

$$|\mathcal{M}_1| = x\binom{n}{m} \geq \frac{\beta-\alpha}{1-\alpha}\binom{n}{m} \geq (\beta-\alpha)\binom{n}{m},$$

as claimed. $\square$

With this lemma, we are ready to prove the supersaturation theorem, Theorem 6.1.

Proof of Theorem 6.1. Fix $\varepsilon > 0$. Recall that $t_{k-1}(m) = (1-\frac{1}{k-1}+o(1))\binom{m}{2}$, where the $o(1)$ term tends to 0 as $m \to \infty$. This implies that there is some fixed m, depending only on ε, so that

$$t_{k-1}(m) \leq \left(1-\frac{1}{k-1}+\frac{\varepsilon}{2}\right)\binom{m}{2}.$$

Let this m be fixed, and let $n \geq m$. Suppose that G is an n-vertex graph with at least $(1 - \frac{1}{k-1} + \varepsilon)\binom{n}{2}$ edges. We apply Lemma 6.2 with $\beta = 1 - \frac{1}{k-1} + \varepsilon$ and $\alpha = 1 - \frac{1}{k-1} + \frac{\varepsilon}{2}$. Then Lemma 6.2 tells us that the number of m-sets $M \subseteq V(G)$ with $e(M) \geq (1 - \frac{1}{k-1} + \frac{\varepsilon}{2})$ is at least $\frac{\varepsilon}{2}\binom{n}{m}$.

Every such m-set M has at least $t_{k-1}(m)$ edges, so Turán's theorem implies that such an M contains a copy of K_k. In other words, we've found at least $\frac{\varepsilon}{2}\binom{n}{m}$ copies of K_k, except that we might have over-counted: each copy of K_k can be counted up to $\binom{n-k}{m-k}$ times, since the k vertices of the K_k can appear in $\binom{n-k}{m-k}$ different m-sets M.

So in total, the number of K_k in G is at least

$$\frac{\frac{\varepsilon}{2}\binom{n}{m}}{\binom{n-k}{m-k}} = \frac{\varepsilon}{2} \cdot \frac{\binom{n}{m}}{\binom{n-k}{m-k}} = \frac{\varepsilon}{2} \cdot \frac{\binom{n}{k}}{\binom{m}{k}} = \frac{\varepsilon}{2\binom{m}{k}}\binom{n}{k},$$

where the middle equality uses the same magic identity as in the proof of Lemma 6.2, namely that $\binom{n}{m}\binom{m}{k} = \binom{n}{k}\binom{n-k}{m-k}$.

To conclude, we recall that m depends solely on ε and k. Therefore, if we define $\delta = \varepsilon/(2\binom{m}{k})$, then this will only depend on ε and k, and that yields the desired result. □

7 Proof of the Erdős–Stone–Simonovits theorem

We are finally ready to prove the Erdős–Stone–Simonovits theorem. We begin by observing a simple reduction, due to Erdős and Simonovits, which says that to prove the bound on $\mathrm{ex}(n, H)$ for all H, it suffices to prove it for a very special class of H. Let $K_k[s]$ denote the complete k-partite graph with parts of size s. (Note that this is the same graph as the Turán graph $T_k(ks)$.)

Proposition 7.1. *Suppose that for all positive integers k, s, we have that*

$$\mathrm{ex}(n, K_k[s]) = \left(1 - \frac{1}{k-1} + o(1)\right)\binom{n}{2}.$$

Then

$$\mathrm{ex}(n, H) = \left(1 - \frac{1}{\chi(H) - 1} + o(1)\right)\binom{n}{2}$$

for every graph H.

Proof. We already proved the lower bound in the Erdős–Stone–Simonovits theorem, namely that

$$\mathrm{ex}(n, H) \geq t_{\chi(H)-1}(n) = \left(1 - \frac{1}{\chi(H) - 1} + o(1)\right)\binom{n}{2}.$$

So it only suffices to prove the upper bound. Now, the key claim is that if H has chromatic number k, then H is a subgraph of $K_k[s]$ for some positive integer s. Indeed, if H has chromatic number k, then we may split the vertices of H into k color classes, with the property that no edge of H goes between two vertices in the same color class. If s is the maximum size of one of the color classes, this precisely means that H is a subgraph of $K_k[s]$. But in that case, we see that

$$\mathrm{ex}(n, H) \le \mathrm{ex}(n, K_k[s]) = \left(1 - \frac{1}{k-1} + o(1)\right)\binom{n}{2},$$

by assumption. □

So it suffices to prove what is often called the Erdős–Stone theorem, namely the statement that $\mathrm{ex}(n, K_k[s]) \le (1 - \frac{1}{k-1} + o(1))\binom{n}{2}$ for every k, s. This is what we now do.

Proof of the Erdős–Stone theorem. Fix some $\varepsilon > 0$. Our goal is to prove that if n is sufficiently large in terms of ε, k, and s, and if G is an n-vertex graph with

$$e(G) \ge \left(1 - \frac{1}{k-1} + \varepsilon\right)\binom{n}{2}$$

edges, then G contains a copy of $K_k[s]$.

By the supersaturation theorem, Theorem 6.1, we know that G has at least $\delta\binom{n}{k}$ copies of K_k, where $\delta > 0$ depends only on ε and k. We define a k-uniform hypergraph $\mathcal{G}$ whose vertex set is $V(G)$, and we make a k-tuple of vertices a hyperedge of $\mathcal{G}$ if and only if the k-tuple defines a copy of K_k in G. Then we have that

$$e(\mathcal{G}) = \#(\text{copies of } K_k \text{ in } G) \ge \delta\binom{n}{k}.$$

Recall that by Theorem 5.5, we have that

$$\mathrm{ex}(n, K^{(k)}_{s,s,\ldots,s}) \le Cn^{k-1/s^{k-1}}$$

for some fixed constant $C > 0$. Now, if δ is fixed (which it is, since it only depends on ε and k), and if n is sufficiently large, then

$$\delta\binom{n}{k} > Cn^{k-1/s^{k-1}}. \tag{3}$$

This is because, as we've discussed previously, $\binom{n}{k}$ grows as $\Theta(n^k)$, and on the right-hand side we have a smaller power of n. So as long as n is sufficiently large in terms of the other parameters, we have that (3) holds.

Thus, for sufficiently large n, we have that $e(\mathcal{G}) > \mathrm{ex}(n, K^{(k)}_{s,s,\ldots,s})$, which implies that $\mathcal{G}$ contains a copy of $K^{(k)}_{s,s,\ldots,s}$. In other words, inside $V(G)$, we can find k sets of s vertices each, with the property that whenever we pick one vertex from each part, they yield a copy of K_k in G. But that precisely means we have found a copy of $K_k[s]$ in G, as claimed. □

8 Topological problems in hypergraphs

The first result I mentioned in this class is that if G is an n-vertex graph with no cycle, then G has at most $n-1$ edges. There is a topological perspective on this: if we view a graph as a topological object, then a cycle is simply a topological copy of the circle which is contained in our graph.

Similarly, we can think of a 3-uniform hypergraph as a two-dimensional topological space. So a natural analogue of this question for 3-graphs is the following: how many hyperedges can we put in an n-vertex 3-graph without a topological copy of the sphere? Formally, we make the following definition. Formally, we care about what are called *triangulations of the sphere*, which are 3-graphs with the property that if we view the vertices and triangles as geometric objects, we get something homeomorphic to a sphere (if you don't like the word homeomorphic, this just means that the resulting geometric object is an "ordinary" polyhedron).

Let $\mathcal{S}$ be the family of all triangulations of the sphere. Then what we are interested in is the 3-uniform extremal function $\mathrm{ex}(n, \mathcal{S})$. This question was studied, and resolved, by Brown, Erdős, and Sós in 1973.

Theorem 8.1 (Brown–Erdős–Sos 1973)**.**

$$\mathrm{ex}(n, \mathcal{S}) = \Theta(n^{5/2}).$$

As always, we need to prove both an upper and a lower bound. In this case, somewhat surprisingly, the upper bound is a little easier than the lower bound, and closely follows the proof of Theorem 5.5.

Proof of the upper bound. Let $\mathcal{G}$ be an n-vertex 3-graph with at least $Cn^{5/2}$ edges, for some big constant $C > 0$. We want to prove that $\mathcal{G}$ contains a triangulation of the sphere. In fact, we will prove that $\mathcal{G}$ necessarily contains a *double pyramid*: this is a 3-graph consisting of vertices $x, y, z_1, \ldots, z_\ell$ and all edges of the form xz_iz_{i+1} and yz_iz_{i+1}, with addition mod ℓ. Here's a picture of the double pyramid in case $\ell = 5$, where every triangle you see is a hyperedge.

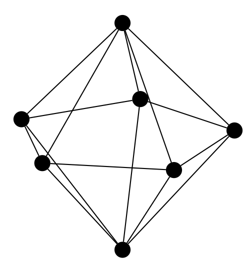

Let X be the number of copies of $K^{(3)}_{1,1,2}$ in $\mathcal{G}$, where $K^{(3)}_{1,1,2} =$ 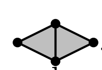. As always, we will estimate X in two ways. First, as in the proof of Theorem 5.5, we know that

$$X = \sum_{\substack{v,w \in V(\mathcal{G})\\ \text{distinct}}} \binom{\mathrm{codeg}(v,w)}{2} \geq \binom{n}{2}\binom{\frac{1}{\binom{n}{2}}\sum_{v,w}\mathrm{codeg}(v,w)}{s} = \binom{n}{2}\binom{3e(\mathcal{G})/\binom{n}{2}}{2} \geq cC^2n^3,$$

where $c > 0$ is an absolute constant. On the other hand, for any pair u_1, u_2 of distinct vertices, let $G(u_1, u_2)$ be the graph with vertex set $V(\mathcal{G}) \setminus \{u_1, u_2\}$, where a pair vw is an edge if and only if $\{v, w, u_1, u_2\}$ form a copy of $K_{1,1,2}^{(3)}$ with v, w as the vertices that lie in both hyperedges.

Now, the key claim is that if $G(u_1, u_2)$ contains a cycle for some u_1, u_2, then $\mathcal{G}$ contains a double pyramid, and thus a triangulation of the sphere. Indeed, in the definition of a double pyramid above, we can set $x = u_1, y = u_2$, and $z_1, \ldots, z_\ell$ the vertices of a cycle in $G(u_1, u_2)$. So if $\mathcal{G}$ has no triangulation of the sphere, then $G(u_1, u_2)$ has no cycle for all u_1, u_2. This implies that

$$e(G(u_1, u_2)) \le v(G(u_1, u_2)) - 1 = n - 3.$$

Adding this all up, we conclude that

$$X = \sum_{\substack{u_1, u_2 \in V(\mathcal{G}) \\ \text{distinct}}} e(G(u_1, u_2)) \le \binom{n}{2}(n-3) < n^3.$$

But if we let C be sufficiently large, this contradicts our lower bound $X \ge cC^2 n^3$. $\square$

The lower bound, just as Klein's construction of a C_4-free graph with many edges, uses a clever input from finite fields. But first, there is a key observation. Let $\mathcal{P}_k$ be the single pyramid hypergraph with a base of size k. Here are drawings of $\mathcal{P}_3, \mathcal{P}_4, \mathcal{P}_5$, where the shaded triangles are the hyperedges.

$\mathcal{P}_3 =$ 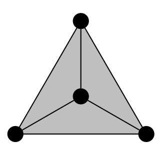 $\mathcal{P}_4 =$ 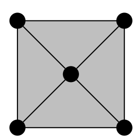 $\mathcal{P}_5 =$ 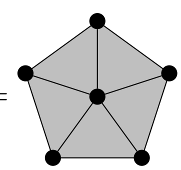

Lemma 8.2. *Let $\mathcal{H} \in \mathcal{S}$ be a triangulation of the sphere. Then $\mathcal{H}$ contains a copy of $\mathcal{P}_k$ for some $k \in \{3, 4, 5\}$.*

Proof. The key fact we need here, which is a simple consequence of Euler's formula, is that every planar graph has a vertex of degree at most 5. Said differently, in any triangulation of the sphere, one of the vertices is incident to at most 5 (two-dimensional) edges of the polyhedron. But that means that that vertex plus its neighbors forms a copy of $\mathcal{P}_k$ for some $k \in \{3, 4, 5\}$. $\square$

To construct the lower bound in Theorem 8.1, Brown, Erdős, and Sós constructed an n-vertex 3-graph with $\Omega(n^{5/2})$ edges and no copy of $\mathcal{P}_3$, $\mathcal{P}_4$, or $\mathcal{P}_5$. Thus, this 3-graph is $\mathcal{S}$-free, by Lemma 8.2.

Before describing the construction, let's briefly return to Klein's construction of a C_4-free graph. Recall that we defined it as follows. We let P be the set of points in $\mathbb{F}_p^2$, and L the set of (non-vertical) lines in $\mathbb{F}_p^2$, and the edges were given by incidence: some $p \in P$ is adjacent to some $\ell \in L$ if and only if $p \in \ell$. Said differently, we identify P with pairs (x, y) where $x, y \in \mathbb{F}_p$, and we identify L with pairs (m, b) with $m, b \in \mathbb{F}_p$, and incidence is given by

$$(x, y) \sim (m, b) \qquad \Longleftrightarrow \qquad y = mx + b.$$

By changing variables, we see that this is more or less the same construction as the following. Define a bipartite graph with parts A, B. Vertices of A are pairs (a_1, a_2) with $a_1, a_2 \in \mathbb{F}_p \setminus \{0\}$, vertices of B are pairs (b_1, b_2) with $b_1, b_2 \in \mathbb{F}_p \setminus \{0\}$, and adjacency is given by

$$(a_1, a_2) \sim (b_1, b_2) \qquad \Longleftrightarrow \qquad a_1b_1 + a_2b_2 = 1.$$

It is not hard to check that this graph, like the earlier construction we defined, has $n = 2(p-1)^2 = \Theta(p^2)$ vertices, $\Theta(p^3) = \Theta(n^{3/2})$ edges, and no copy of C_4. The proof is the same as before: assume there is a C_4, which yields four linear equations, and by tracing through these linear equations we find that we can't have two distinct vertices from A and two distinct vertices from B.

In fact, we can generalize this construction even further. If we fix non-zero $c_1, c_2 \in \mathbb{F}_p$, then we can define an alternate bipartite graph $G(c_1, c_2)$ with vertex set $A \cup B$, with adjacency given by

$$(a_1, a_2) \sim (b_1, b_2) \qquad \Longleftrightarrow \qquad c_1 \cdot a_1b_1 + c_2 \cdot a_2b_2 = 1.$$

Then again, G is a C_4-free graph with $n = 2(p-1)^2$ vertices and $\Theta(n^{3/2})$ edges. Again, the point is that we can trace through the linear equations, and the non-zero multipliers c_1, c_2 don't matter.

With this, we can pretty easily prove the lower bound in Theorem 8.1. As discussed above, thanks to Lemma 8.2, it suffices to prove the following result.

Proposition 8.3 (Brown–Erdős–Sós 1973)**.** *For all n, there exists an n-vertex 3-graph $\mathcal{G}$ with $\Theta(n^{5/2})$ edges and no copy of $\mathcal{P}_3, \mathcal{P}_4$, or $\mathcal{P}_5$.*

Proof. We assume for simplicity that $n = 3(p-1)^2$ for some prime p; as in the proof of Theorem 4.4, this assumption can be lifted thanks to Bertrand's postulate. Let $\mathcal{G}$ be the 3-partite 3-graph with parts A, B, C, where

$$\begin{aligned} A &= \{(a_1, a_2) : a_1, a_2 \in \mathbb{F}_p \setminus \{0\}\} \\ B &= \{(b_1, b_2) : b_1, b_2 \in \mathbb{F}_p \setminus \{0\}\} \\ C &= \{(c_1, c_2) : c_1, c_2 \in \mathbb{F}_p \setminus \{0\}\}, \end{aligned}$$

and we make a triple $\{(a_1, a_2), (b_1, b_2), (c_1, c_2)\}$ a hyperedge of $\mathcal{G}$ if and only if

$$a_1b_1c_1 + a_2b_2c_2 = 1.$$

Note that if we fix $(c_1, c_2) \in C$, then the graph $G(c_1, c_2)$ of pairs $(a_1, a_2) \in A, (b_1, b_2) \in B$ with $\{(a_1, a_2), (b_1, b_2), (c_1, c_2)\}$ a hyperedge is precisely the graph $G(c_1, c_2)$ defined above. In particular, we see that there are $\Theta(p^3)$ hyperedges containing every $(c_1, c_2) \in C$, which implies that there are $\Theta(p^5) = \Theta(n^{5/2})$ hyperedges in $\mathcal{G}$.

Additionally, note that for every $(c_1, c_2) \in C$, the graph $G(c_1, c_2)$ is bipartite and C_4-free. So in particular, $G(c_1, c_2)$ does not contain any copy of C_3, C_4, or C_5. This implies that there is no copy of $\mathcal{P}_3, \mathcal{P}_4$, or $\mathcal{P}_5$ centered at any $(c_1, c_2) \in C$. But the hypergraph is symmetric, so the same argument applies if we try to center the $\mathcal{P}_k$ in A or in B. So we conclude that $\mathcal{G}$ is $\{\mathcal{P}_3, \mathcal{P}_4, \mathcal{P}_5\}$-free, as claimed. □

Given the Brown–Erdős–Sós theorem, it is natural to ask about triangulations of other surfaces. Namely, for an integer $g \geq 1$, let $\mathcal{S}_g$ denote the set of all 3-graphs which are triangulations of a genus-g surface. (Recall that a genus-g surface is the same as a g-hole torus.) Then it is natural to ask for the value of $\text{ex}(n, \mathcal{S}_g)$. For many years, there was no real progress on this problem, but it was very recently resolved.

Theorem 8.4 (Kupavskii–Polyanskii–Tomon–Zakharov 2021)**.** *For every $g \geq 1$, we have that*

$$\text{ex}(n, \mathcal{S}_g) = \Theta(n^{5/2}).$$

Note that the implicit constant in the big-Θ above depends on g. To prove this, they had to prove both upper and lower bounds. The lower bound is totally different from above, essentially because there is not as simple of a "local" obstruction to the presence of a triangulation of a genus-g surface. Instead of the lower bound above, they use the probabilistic method, plus a powerful result of Gao on the number of n-vertex triangulations of a genus-g surface.

The upper bound uses two separate, pretty cool ideas. The first is a much trickier version of the double pyramid argument of Brown, Erdős, and Sós, which allows one to build a triangulated torus in a 3-graph by finding and gluing together very many double pyramids along a long cycle. The second is a simple probabilistic argument that shows that we can upper-bound $\text{ex}(n, \mathcal{S}_{g+1})$ in terms of $\text{ex}(n, \mathcal{S}_g)$. Basically, the idea is that we can randomly partition the edges of $\mathcal{G}$ into two parts. In one part, we can find a triangulation of a genus-g surface, using what we already know about $\text{ex}(n, \mathcal{S}_g)$. In the other part, we can find a triangulation of the torus, using what we know about $\text{ex}(n, \mathcal{S}_1)$. Moreover, by exploiting the randomness, we can ensure that we can glue these together to get a triangulated genus-$(g+1)$ surface. This can be used to show inductively that $\text{ex}(n, \mathcal{S}_{g+1}) \leq O(n^{5/2})$.

What about non-orientable surfaces? It turns out that every non-orientable surface has a non-orientable genus $k \geq 1$, and a non-orientable surface of genus k can be obtained by gluing together k copies of the real projective plane. Let $\mathcal{S}_k^-$ denote the set of triangulations of the genus-k non-orientable surface. Then the same probabilistic lower bound implies that $\text{ex}(n, \mathcal{S}_k^-) \geq \Omega(n^{5/2})$ for all n. Moreover, the same gluing argument for the upper bound shows that it suffices to understand the case of the real projective plane, i.e. to upper-bound $\text{ex}(n, \mathcal{S}_1^-)$. This was *extremely* recently resolved by Maya Sankar.

Theorem 8.5 (Sankar 2022+)**.** $\text{ex}(n, \mathcal{S}_1^-) = \Theta(n^{5/2})$. *As a consequence,* $\text{ex}(n, \mathcal{S}_k^-) = \Theta(n^{5/2})$ *for all $k \geq 1$.*