Ma il mio mistero è chiuso in me

Nessun dorma, from Puccini's *Turandot*
Libretto: G. Adami and R. Simoni

1 Turán's theorem

How many edges can we place among n vertices in such a way that we make no triangle? This is perhaps the first question asked in the field of *extremal graph theory*, which is the topic of this talk. Upon some experimentation, one can make the conjecture that the best thing to do is to split the vertices into two classes, of sizes x and $n-x$, and then connect any two vertices in different classes. This will not contain a triangle, since by the pigeonhole principle, among any three vertices there will be two in the same class, which will not be adjacent. This construction will have $x(n-x)$ edges, and by the AM-GM inequality, we see that this is maximized when x and $n-x$ are as close as possible; namely, our classes should have sizes $\lfloor \frac{n}{2} \rfloor$ and $\lceil \frac{n}{2} \rceil$. Thus, this graph will have $\lfloor \frac{n}{2} \rfloor \lceil \frac{n}{2} \rceil \approx \frac{n^2}{4} \approx \frac{1}{2}\binom{n}{2}$ edges.

Indeed, this is the unique optimal construction, as was proved by Mantel in 1907. In 1941, Turán considered a natural extension of this problem, where we forbid not a triangle, but instead a larger complete graph K_{r+1}. As before, a natural construction is to split the vertices into r almost-equal classes and connect all pairs of vertices in different classes.

Definition 1. The *Turán graph* $T(n,r)$ is the graph on n vertices gotten by splitting the vertices into r classes of sizes $\lfloor \frac{n}{r} \rfloor$ or $\lceil \frac{n}{r} \rceil$, and connecting any pair in different classes.

A simple computation shows that the number of edges in $T(n,r)$ is

$$\mathsf{e}(T(n,r)) \approx \left(1-\frac{1}{r}\right)\binom{n}{2},$$

though the precise number of edges depends on $n \mod r$. As a generalization of Mantel's theorem (which corresponds to $r=2$), Turán proved the following theorem.

Theorem 1 (Turán 1941)**.** *If a graph G on n vertices contains no K_{r+1}, then it has at most $\mathsf{e}(T(n,r))$ edges. Moreover, if $\mathsf{e}(G) = \mathsf{e}(T(n,r))$, then $G \cong T(n,r)$, i.e. $T(n,r)$ is the unique extremal graph.*

There are by now dozens of proofs of this theorem, but here is one that I find particularly nice.

Proof. Let G be a K_{r+1}-free graph with the maximum possible number of edges. It suffices to prove that $G \cong T(n,r)$. For a vertex x, let $\deg(x)$ denote its degree. Suppose there exist vertices x,y,z in G such that $x \sim z$, but $x \nsim y \nsim z$. Suppose moreover that $\deg(y) < \deg(x)$. Then we create a new graph G' by deleting y and replacing it with a copy of x, namely a new vertex x' whose neighbors are the same as the neighbors of x. Since $x \nsim x'$, we can't have created any new K_{r+1}, so G' is also K_{r+1}-free. But G' has more edges than G, since in going from G to G' we deleted $\deg(y)$ edges and added $\deg(x) > \deg(y)$ edges, contradicting the maximality of G.

So we may assume that $\deg(y) \geq \deg(x)$. By the same argument, $\deg(y) \geq \deg(z)$. In this case, we make a new graph G'' by deleting both x and z and replacing them both by copies of y. Since we assumed that $x \sim z$, we deleted $\deg(x) + \deg(z) - 1$ edges, and added in their place $2\deg(y)$ edges, which is more. Thus, $\mathsf{e}(G'') > \mathsf{e}(G)$, again contradicting the maximality of G.

In either case we get a contradiction, so we find that there is no triple of vertices x, y, z such that $x \sim z$ but $x \nsim y \nsim z$. This means that non-adjacency is a transitive relation, and thus an equivalence relation. So we may partition the vertices of G into equivalence classes, i.e. partition G into some number k of blocks such that each block has no internal edges, but all pairs of vertices in different blocks are adjacent. Since G is K_{r+1}-free, we must have that $k < r + 1$. At this point, a simple argument using the AM-GM inequality (or Jensen's inequality) shows that the construction that maximizes edges is to take $k = r$ and to make the classes as equal as possible in size. Thus, $G \cong T(n, r)$. □

Despite its simple statement and proof, Turán's theorem is surprisingly powerful and useful. In practice, it is often most useful as a "global-to-local" principle: it asserts that if a graph is globally dense (namely having more than $\approx (1 - \frac{1}{r})\binom{n}{2}$ edges), then it has a locally very dense part, namely a complete subgraph K_{r+1}. This is useful because in many proofs, you can construct a graph which you know is dense for some unrelated reason, and Turán's theorem allows you to then restrict to a complete subgraph. As a different cool application, I leave to you to prove the following simple result.

Theorem 2 (Katona, 1969)**.** *Let X, Y be independent random vectors drawn from some discrete distribution on $\mathbb{R}^d$. Then*

$$\Pr(\|X + Y\| \geq 1) \geq \frac{1}{2}\Pr(\|X\| \geq 1)^2.$$

2 The Erdős–Stone theorem

Turán generalized Mantel's theorem by forbidding not just a triangle, but an arbitrary complete subgraph. However, there is no reason to restrict our attention to complete graphs, and we can make the following definition.

Definition 2. The *extremal number* of a graph H and an integer n, denoted $\mathrm{ex}(n, H)$, is the maximum number of edges among all graphs on n vertices that do not contain H as a subgraph.

In this language, Turán's theorem says that $\mathrm{ex}(n, K_{r+1}) \sim (1 - \frac{1}{r})\binom{n}{2}$. However, it tells us nothing about e.g. $\mathrm{ex}(n, C_4)$ and $\mathrm{ex}(n, C_5)$, where C_k denotes the cycle graph of length k; note that since C_3 is the triangle, these are also natural generalizations of Mantel's theorem.

It turns out that the order of $\mathrm{ex}(n, H)$ is closely related to the chromatic number $\chi(H)$. Recall that a *proper k-coloring* of H is a function $f : V(H) \to \{1, \dots, k\}$ such that $f(x) \neq f(y)$ whenever $x \sim y$. The *chromatic number* $\chi(H)$ of H is defined as the minimum k for which a proper k-coloring exists. The connection between $\chi(H)$ and $\mathrm{ex}(n, H)$ can be first seen by the following simple observation.

Proposition 1. *Let $r = \chi(H) - 1$. Then $T(n, r)$ contains no copy of H.*

Proof. Suppose we had a copy of H in $T(n, r)$. We can view this as an injective map $\iota : V(H) \to V(T(n, r))$. If we let the classes of $T(n, r)$ be $V_1, \dots, V_r$, then we can define $f : V(H) \to \{1, \dots, r\}$ by letting $f(v)$ be the index of $\iota(v)$, namely $f(v) = s$ if $\iota(v) \in V_s$. Then f is a proper r-coloring of H, since two vertices with the same label must be mapped into the same class of $T(n, r)$, and must therefore be non-adjacent. But this contradicts that $\chi(H) > r$, proving the claim. $\square$

This proposition immediately implies that

$$\mathrm{ex}(n, H) \geq \mathsf{e}(T(n, \chi(H) - 1)) \approx \left(1 - \frac{1}{\chi(H) - 1}\right) \binom{n}{2}.$$

Astonishingly, this lower bound is asymptotically correct, as shown by the following fundamental result.

Theorem 3 (Erdős–Stone 1946)**.** *For any graph H,*

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

where the $o(1)$ tends to 0 as $n \to \infty$.

Note that the Erdős–Stone theorem does not assert that the Turán graph is the extremal graph, simply that it has almost as many edges as the extremal graph. In fact, we can in general do better than the Turán graph. The simplest example is the following graph:

$H =$ 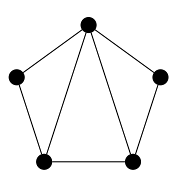

It has chromatic number 3, but we have that $\mathrm{ex}(n, H) \geq \frac{n^2}{4} + \frac{n}{4}$, which is more by the linear factor $n/4$ than the number of edges in the Turán graph $T(n, 2)$. Indeed, to see that $\mathrm{ex}(n, H) \geq \frac{n^2}{4} + \frac{n}{4}$, consider the graph gotten by adding $n/4$ disjoint edges to one part of the Turán graph $T(n, 2)$:

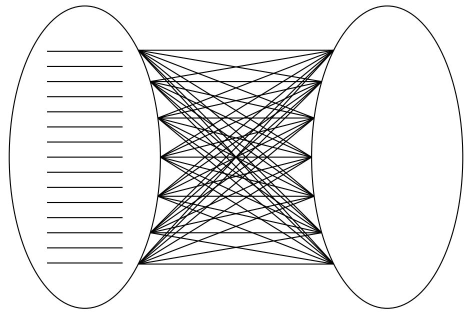

It is not too hard to verify that this graph contains no copy of this H, and we see that it has more edges than $T(n, 2)$. However, the Erdős–Stone theorem says that we can never do *much* better than $T(n, \chi(H)-1)$; namely, we can only add a sub-quadratic number of edges.

The proof of the Erdős–Stone theorem is a bit too complicated for this talk, but by the end we'll actually have most of the necessary tools to prove it.

3 The problem of Zarankiewicz

If we accept the $o(n^2)$ error term, then the Erdős–Stone theorem completely answers the question of how large $\mathrm{ex}(n, H)$ is. However, in one crucial case, we should not accept this error term. Namely, if H is bipartite, then $\chi(H) = 2$, and the Erdős–Stone theorem simply says that $\mathrm{ex}(n, H) = o(n^2)$. In other words, if H is bipartite, the main term vanishes, and the error term becomes the new main term. Thus, the Erdős–Stone theorem says almost nothing about bipartite graphs.

For historical reasons, the study of extremal numbers for bipartite graphs is often called the problem of Zarankiewicz. To get a taste for how this theory goes, let's consider $\mathrm{ex}(n, C_4)$.

Theorem 4 (Kővari–Sós–Turán 1954)**.** $\mathrm{ex}(n, C_4) = O(n^{3/2})$.

Proof. Let G be a graph with no C_4. We need to prove that G has at most $Cn^{3/2}$ edges, for some constant C. Let X denote the number of paths of length 2 in G, namely the number of copies of $\bullet\!\!-\!\!\!-\!\!\bullet\!\!-\!\!\!-\!\!\bullet$ in G. We can count X in two ways. First, if we fix the two endpoints of this path, then the fact that G is C_4-free means that there is at most one choice for the middle vertex, implying that $X \le \binom{n}{2}$. On the other hand, if we fix the middle vertex to be some vertex v, then there are exactly $\binom{\deg(v)}{2}$ paths with v as their middle vertex, implying that

$$X = \sum_{v \in V} \binom{\deg(v)}{2} \approx \frac{1}{2} \sum_{v \in V} \deg(v)^2 \ge \frac{1}{2n} \left(\sum_{v \in V} \deg(v) \right)^2,$$

where the last inequality follows from Cauchy–Schwarz. If we observe that $\sum_v \deg(v) = 2\mathsf{e}(G)$, since every edge in G is counted exactly twice in this sum (once for each endpoint), then we find that

$$\binom{n}{2} \ge X \gtrsim \frac{1}{2n} \left(\sum_{v \in V} \deg(v) \right)^2 = \frac{2\mathsf{e}(G)^2}{n}.$$

Rearranging shows that $e(G) \leq Cn^{3/2}$, as claimed. $\square$

The key idea in this proof was to treat the cycle C_4 as $K_{2,2}$, namely the complete bipartite graph where both classes have two vertices. Indeed, we used that G was $K_{2,2}$-free to conclude that every pair of endpoints contributed at most one path to the count X. In fact, Kővari, Sós, and Turán used the same idea to upper-bound $\mathrm{ex}(n, K_{s,t})$ for all s and t.

Theorem 5 (Kővari–Sós–Turán 1954). *For all fixed $s \leq t$, $\mathrm{ex}(n, K_{s,t}) = O(n^{2-1/s})$.*

The proof is essentially identical to the above, except that we use Hölder's inequality (or Jensen's inequality) instead of Cauchy–Schwarz. Much like Turán's theorem, this result is much more useful than it appears at first glance; here is one cute application that I leave as another exercise.

Theorem 6. *For a set $S \subseteq \mathbb{Z}^2$ and an integer k, let*

$$d_k(S) = \max_{\substack{A,B \subset \mathbb{Z} \\ |A|=|B|=k}} \frac{|S \cap (A \times B)|}{k^2},$$

and define the upper density of S to be $d(S) = \limsup_{k\to\infty} d_k(S)$. Then for any $S \subseteq \mathbb{Z}^2$, $d(S) \in \{0, 1\}$.

The problem of Zarankiewicz becomes much more interesting once we start to search for lower bounds. Namely, can we find constructions of $K_{s,t}$-free graphs that match the Kővari–Sós–Turán upper bound? In short, the answer is somewhere between yes and no. Let's begin with $\mathrm{ex}(n, C_4)$ again.

Theorem 7 (Klein 1938, Erdős–Rényi–Sós 1966, Brown 1966). *There exists a C_4-free graph on n vertices with $\Omega(n^{3/2})$ edges.*

Proof. Let q be a prime power, and let $\mathbb{F}_q$ denote the finite field of order q. Let G be a bipartite graph with two parts P and L, where $P = \mathbb{F}_q^2$ and L consists of all lines in $\mathbb{F}_q^2$. There are roughly q^2 lines in $\mathbb{F}_q^2$, so G has $n = \Theta(q^2)$ vertices. We connect a point $p \in P$ to a line $\ell \in L$ if and only if p is incident to ℓ. Since every line contains q points, the number of edges in G is $|L|q = \Theta(q^3) = \Theta(n^{3/2})$. Moreover, we claim that G is C_4-free. Indeed, a C_4 in G would have to consist of two points p_1, p_2 and two lines ℓ_1, ℓ_2 such that p_i is incident with ℓ_j for $i, j \in \{1, 2\}$. However, two lines in $\mathbb{F}_q^2$ can intersect in at most one point, so such a configuration cannot exist, as claimed. $\square$

For $K_{3,3}$, the Kővari–Sós–Turán upper bound is $\mathrm{ex}(n, K_{3,3}) = O(n^{5/3})$. Also for this, there is a matching lower bound.

Theorem 8 (Brown 1966). *There exists a $K_{3,3}$-free graph on n vertices with $\Omega(n^{5/3})$ edges.*

Proof sketch. Let p be a prime that is $3 \bmod 4$. We define a graph G whose vertex set is $V = \mathbb{F}_p^3$, so that $n = p^3$. We connect two vertices $x = (x_1, x_2, x_3), y = (y_1, y_2, y_3) \in \mathbb{F}_p^3$ if their "Euclidean distance" is 1, namely if

$$(x_1 - y_1)^2 + (x_2 - y_2)^2 + (x_3 - y_3)^2 = 1.$$

Since this is a single polynomial constraint, the neighborhood of each vertex is a codimension-one variety, so we expect that each vertex has $\Theta(p^2)$ neighbors, which turns out to be true. Thus, the number of edges in G is $n \cdot \Theta(p^2) = \Theta(p^5) = \Theta(n^{5/3})$. Moreover, it turns out that G is $K_{3,3}$-free. Indeed, a $K_{3,3}$ in G corresponds to three "unit spheres" in $\mathbb{F}_p^3$ intersecting in at least 3 points. In $\mathbb{R}^3$, we know that three unit spheres can intersect in only two points (since any two intersect in a circle, and a circle and a sphere can only intersect in two points). It turns out that the same holds over $\mathbb{F}_p^3$ (at least when $p \equiv 3 \bmod 4$, though this construction can be modified to work when $p \equiv 1 \bmod 4$). □

Seeing these two examples, it is clear what we should do to construct $K_{s,s}$-free graphs with $\Omega(n^{2-1/s})$ edges for larger s. We should work in $\mathbb{F}_q^s$, and have our graphs be defined by some algebraic condition so that the neighborhood of each vertex is a codimension-one variety. By a result in algebraic geometry known as the Lang–Weil bound, this will imply that every vertex will have $\Theta(q^{s-1})$ neighbors (assuming some technical conditions which should hold in general), so the number of edges will indeed be $\Omega(n^{2-1/s})$. So all that remains is to pick this algebraic equation intelligently so that the resulting graph is $K_{s,s}$-free. Unfortunately, even though this technique has been around since the 1960s, no one has been able to make it work even in the case $s = 4$.[1]

However, if we recall that the same Kővari–Sós–Turán upper bound holds for $\mathrm{ex}(n, K_{s,t})$ for all $s \leq t$, then some more is known. Namely, the following holds.

Theorem 9 (Kollár–Rónyai–Szabó 1996, Alon–Rónyai–Szabó 1999)**.** *If $t \geq (s-1)! + 1$, then $\mathrm{ex}(n, K_{s,t}) = \Theta(n^{2-1/s})$.*

Thus, for example, we know that $\mathrm{ex}(n, K_{4,7}) = \Theta(n^{7/4})$, though we still have no matching lower bound for the Kővari–Sós–Turán upper bound of $\mathrm{ex}(n, K_{4,4}) = O(n^{7/4})$. The construction used in both these results is known as a *norm graph*. The simpler one, which only works when $t \geq s! + 1$, is as follows. We build a bipartite graph with parts A and B, where A and B are both equal to $\mathbb{F}_{p^s}$. If we let $N : \mathbb{F}_{p^s} \to \mathbb{F}_p$ be the norm associated to the field extension $\mathbb{F}_{p^s}/\mathbb{F}_p$, then we connect $x \in A$ to $y \in B$ whenever $N(x + y) = 1$. This is a single algebraic condition, so by the argument above, this graph will have $\Theta(n^{2-1/s})$ edges. Moreover, it turns out that it will indeed be $K_{s,t}$-free for $t \geq s! + 1$, though proving this is not so easy; it

[1] Additionally, a recent result of Blagojević, Bukh, and Karasev shows that this technique sketched above cannot work without some modification. More precisely, they show that for any polynomial equation over $\mathbb{Z}$, reducing it mod p cannot give a $K_{s,s}$-free graph as described above for $s \geq 4$. This implies that any construction along these lines must use a different equation as p varies, rather than using a single equation that we reduce mod p. Their proof uses a very interesting technique, showing that there is a *topological* obstruction to this working; the only property they use of the polynomial equation over $\mathbb{Z}$ is that it defines a continuous map on $\mathbb{R}^s$.

relies on some algebraic geometry and number theory, and somehow has to do with the fact that the condition $N(x+y) = 1$ mixes additive and multiplicative properties of the field $\mathbb{F}_{p^s}$. To obtain the stronger result that works for $t \geq (s-1)! + 1$, one has to use a modification of this construction, though the basic idea is similar.

In short, our understanding of the extremal numbers of $K_{s,t}$ is somewhat muddled. When t is much larger than s, we know the exact order of $\mathrm{ex}(n, K_{s,t})$, but even $\mathrm{ex}(n, K_{4,4})$ is unknown. Most people believe that the Kővari–Sós–Turán upper bound is tight, but no one is sure, and it seems possible that a stronger upper bound might hold.

Another natural class of bipartite graphs is the class of even cycles. For these, we have the following upper bound.

Theorem 10 (Bondy–Simonovits 1974). *Let $t \geq 2$ be fixed. Then* $\mathrm{ex}(n, C_{2t}) = O(n^{1+1/t})$.

For $t = 2$, we recover the upper bound $\mathrm{ex}(n, C_4) = O(n^{3/2})$ we saw above, and we also know that this bound is tight. Moreover, the Bondy–Simonovits upper bound is known to be tight also for C_6 and C_{10}; namely, there are algebraic constructions of C_6- and C_{10}-free graphs with $\Omega(n^{4/3})$ and $\Omega(n^{6/5})$ edges, respectively. However, for any other value of t, it is not known whether this upper bound is tight. As with the case of complete bipartite graphs, it seems that some algebraic construction is the right approach, but no one knows what exactly this construction should be.

Let me finish by stating a famous conjecture of Erdős and Simonovits, which remains wide open.

Conjecture. *For every bipartite graph H, there exists a rational number $\alpha \in [1, 2)$ with* $\mathrm{ex}(n, H) = \Theta(n^\alpha)$. *Conversely, for every rational $\alpha \in [1, 2)$, there exists a bipartite graph H with* $\mathrm{ex}(n, H) = \Theta(n^\alpha)$.

4 Hypergraphs

Definition 3. A *ℓ-uniform hypergraph* is a pair $\mathcal{H} = (V, E)$, where V is a finite set and E consists of unordered ℓ-tuples of elements of V. Thus, a graph is just a 2-uniform hypergraph.

For topologists, it may be convenient to think of an ℓ-uniform hypergraph as just a finite $(\ell - 1)$-dimensional simplicial complex. However, note that we are only interested in the top-dimensional simplices when we speak of hypergraphs.

Definition 4. Given a fixed ℓ-uniform hypergraph $\mathcal{H}$, let $\mathrm{ex}(n, \mathcal{H})$ denote the maximum number of edges in an ℓ-uniform hypergraph on n vertices which does not contain $\mathcal{H}$ as a subhypergraph.

Note that if $\mathcal{H}$ is ℓ-uniform, then $\mathrm{ex}(n, \mathcal{H}) \leq \binom{n}{\ell} = O(n^\ell)$. As with graphs, the most basic hypergraphs whose extremal numbers we wish to understand are the complete hypergraphs.

Definition 5. Given integers $k \geq \ell \geq 2$, the *complete ℓ-uniform hypergraph on k vertices*, denoted $K_k^{(\ell)}$ is the hypergraph whose vertex set is $\{1, \ldots, k\}$, and whose edges consist of all ℓ-tuples of integers in $\{1, \ldots, k\}$.

It is pretty easy to show that there exists a real number, say $C_{k,\ell}$, such that $\mathrm{ex}(n, K_k^{(\ell)}) = (C_{k,\ell} + o(1))\binom{n}{\ell}$. However, unlike the graph case, there isn't *any* pair $k > \ell > 2$ for which we know the value of $C_{k,\ell}$. In fact, Erdős offered \$500 to evaluate this number for a single pair (k, ℓ), and \$1000 to evaluate it in general. For the simplest case, of $K_4^{(3)}$, we know

$$\frac{5}{9}\binom{n}{3} \leq \mathrm{ex}(n, K_4^{(3)}) \leq 0.561666\binom{n}{3}.$$

The lower bound is a simple construction of Turán, which he conjectured to be tight, and the upper bound is due to Razborov and follows from a lengthy computer-aided computation. There are several reasons why the hypergraph case appears to be much harder than the graph case; one main reason is that we cannot expect a result like Turán's theorem, which finds the unique extremal graph. For instance, for the case of $\mathrm{ex}(n, K_4^{(3)})$, there exist exponentially many different hypergraphs on n vertices with $(\frac{5}{9} + o(1))\binom{n}{3}$ edges and no $K_4^{(3)}$. Assuming Turán's conjecture is correct (that $\frac{5}{9}$ is the right answer), this shows that we cannot hope for a single "best" hypergraph.

There is a *lot* more to be said about Turán-type problems, both for graphs and for hypergraphs, but this is probably enough for now.