1 What is extremal combinatorics?

We begin with a magic trick. I ask the audience to give me ten integers between 1 and 100, and then I find two disjoint subsets of this set of integers which have the same sum. The following theorem guarantees that I can always do this. Here, and throughout the course, we'll denote by $[\![n]\!]$ the range $\{1, 2, \ldots, n\}$.

Theorem 1.1. *Let $a_1, \ldots, a_{10}$ be integers in $[\![100]\!]$. Then there exist two disjoint, non-empty sets $S, T \subseteq [\![10]\!]$ such that*

$$\sum_{i \in S} a_i = \sum_{j \in T} a_j.$$

Proof. For every non-empty set $S \subseteq [\![10]\!]$, let us denote by a_S the number $a_S \coloneqq \sum_{i \in S} a_i$. Note that there are $2^{10} - 1 = 1023$ non-empty subsets of $[\![10]\!]$, so we have defined 1023 numbers. Moreover, for every set S, we have that

$$a_S = \sum_{i \in S} a_i \leqslant |S| \cdot \max\{a_1, \ldots, a_{10}\} \leqslant 10 \cdot 100 = 1000.$$

That is, each of the numbers a_S lies in the interval $[1, 1000]$. As we have defined 1023 numbers, by the pigeonhole principle, two of them must be equal. That is, there must exist $S \neq T$ such that $a_S = a_T$.

We are almost done, except that S and T need not be disjoint, but this can be easily remedied: we define $S' \coloneqq S \setminus T = S \setminus (S \cap T)$, and similarly $T' \coloneqq T \setminus S$. Then S and T are disjoint. Moreover, we have that

$$a_{S'} = \sum_{i \in S'} a_i = \sum_{i \in S} a_i - \sum_{i \in S \cap T} a_i = a_S - a_{S \cap T},$$

and similarly $a_{T'} = a_T - a_{S \cap T}$, so $a_{S'} = a_{T'}$. Finally, these two sets are non-empty: this is because each of S, T was non-empty, and we cannot have $S \subseteq T$ or $T \subseteq S$, as this would contradict that they have the same sum. Therefore even when we remove $S \cap T$ from each of them, they remain non-empty. □

Note that the proof above shows that in fact, I could have allowed the audience to pick ten numbers in the range $[\![102]\!]$: in that case, each sum a_S would be at most 1020, and I could still apply the pigeonhole principle. This leads to a natural generalization of this question. Suppose I want to let my audience pick n numbers from some range $[\![M]\!]$, while still ensuring that no matter what they pick, I can find two disjoint subsets with equal sum. How large can I take M while ensuring that this is possible? Formally, we can define

$$\operatorname{magic}(n) \coloneqq \max \Bigg\{ M \in \mathbb{N} \mid \text{for all } a_1, \ldots, a_n \in [\![M]\!], \text{ there are disjoint } S, T \subseteq [\![n]\!] \text{ with } \sum_{i \in S} a_i = \sum_{j \in T} a_j \Bigg\}.$$

Note that proving a *lower bound* on magic(n) boils down to proving a variant of Theorem 1.1. Indeed, to prove a lower bound on magic(n), it suffices to exhibit some M that works, which implies that magic(n) $\geqslant M$. Thus, for example, Theorem 1.1 states that magic(10) $\geqslant 100$, and as pointed out above, the proof actually gives magic(10) $\geqslant 102$. By following the proof of Theorem 1.1, one can more generally show that

$$\text{magic}(n) \geqslant \frac{2^n - 1}{n}.$$

What about proving *upper bounds* on magic(n)? This, in turn, boils down to finding a clever set of integers $a_1, \dots, a_n$ such that all their subset sums are distinct. If we can find such integers $a_1, \dots, a_n \in [\![M]\!]$, we have proved that M does not work for the magic trick, i.e. that magic(n) $< M$. A natural idea for how to pick such numbers $a_1, \dots, a_n$ is to let them be the powers of 2, i.e. to set $a_i = 2^{i-1}$. Then it is not hard to check that all subset sums are distinct, which implies that

$$\text{magic}(n) < 2^{n-1}.$$

Our lower and upper bounds are not too far apart, but they differ by roughly a factor of n. And although obtaining these nearly-matching bounds was quite easy, closing this gap is a major open problem! Erdős offered \$500 for a proof or disproof of the following conjecture.

Conjecture 1.2 (Erdős's distinct sums conjecture). *We have that* magic(n) $\geqslant \Omega(2^n)$. *That is, the upper bound is best possible up to a constant factor.*

Here, we recall the asymptotic notation that we'll use a lot throughout the course: the statement $f(n) \geqslant \Omega(g(n))$ means that for all sufficiently large n, we have that $f(n) \geqslant c \cdot g(n)$, where $c > 0$ is some constant independent of n.

The best known bounds on this problem are as follows. First, for the upper bound, Bohman found a construction of a set of integers beating the powers of two, which implies that

$$\text{magic}(n) \leqslant 0.22002 \cdot 2^n$$

for all sufficiently large n. In the other direction, a beautiful probabilistic[1] argument of Erdős and Moser proves that

$$\text{magic}(n) \geqslant \Omega\left(\frac{2^n}{\sqrt{n}}\right),$$

which improves the argument of Theorem 1.1 by roughly a factor of $\sqrt{n}$. The best known constant factor was obtained only recently by Dubroff, Fox, and Xu, who proved that

$$\text{magic}(n) \geqslant \left(\sqrt{\frac{2}{\pi}} - o(1)\right) \frac{2^n}{\sqrt{n}},$$

[1] You may wonder what probability has to do with any of this. Somewhat amazingly, even for problems like this that have nothing do with probability, many of our most powerful techniques are probabilistic. This introduction of random tools into non-random problems is called the *probabilistic method*, and we will see several examples of it throughout the course.

where $o(1)$ represents a quantity that tends to 0 as n tends to infinity.

This simple problem is an instance of extremal combinatorics, which is the topic of this course. In this context, "extremal" means "maximum or minimum": extremal questions are questions of the form "how large or small can an object be, subject to certain constraints?" For example, magic(n) is an extremal function: it asks how large M can be so long as $[\![M]\!]$ contains no set of n integers with distinct subset sums. Equivalently, we can view this function in the "opposite" perspective, and ask how many elements we can fit into $[\![M]\!]$ without creating two equal subset sums. The "combinatorics" in "extremal combinatorics" means that we are working with discrete structures, such as, in this case, finite sets of integers. For most of this course we will be focusing on extremal graph theory, where the objects of study are graphs (and their variants).

As in this example, whenever we are confronted with an extremal function, we would like to prove upper and lower bounds on it that are as close as possible. Generally speaking, one bound is proved by constructions of large or small objects satisfying the requisite property (like the powers of 2 above), and the other bound is proved by showing that *any* object of appropriate size must not have the requisite property (like the proof of Theorem 1.1). One of the things that I love about extremal combinatorics is that there are many natural and easy-to-state questions (such as the one discussed above), for which we can easily prove some upper and lower bounds, but for which it currently seems extremely difficult (or downright hopeless) to pin down the answer. That said, there keep being remarkable breakthroughs on some previously intractable problems, usually via the introduction of remarkable new ideas. In this course, I will try to give you a flavor of what extremal combinatorics is all about, some of the beautiful ideas that are used in the field, what some of the major open problems are, and hint at some of the recent breakthroughs on these and other questions.

2 What is extremal graph theory?

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 $n \geqslant 3$ 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 major goal, for the next few lectures, 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 some time is what happens when the constraint is excluding a single "forbidden subgraph".

Definition 2.1. 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 2.2. 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 $\mathrm{ex}(n, H) \geqslant e(G)$. 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 $\mathrm{ex}(n, H) < m$.

3 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

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

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

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

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

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

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 3.1 (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 $r \geqslant 3$.

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 3.2. 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

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

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 $r \geqslant 3$.

Theorem 3.3 (Turán 1941)**.** *For every $r \geqslant 3$, 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 $1 \leqslant n \leqslant r-1$. 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

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

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

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

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 adjacent 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 over the next few days, 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!