With this lemma, we are ready to prove the supersaturation theorem, Theorem 7.1.

Proof of Theorem 7.1. 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

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

Let this m be fixed, and let $n \geqslant m$. Suppose that G is an n-vertex graph with at least $(1 - \frac{1}{k-1} + \varepsilon)\binom{n}{2}$ edges. We apply Lemma 7.2 with $\beta = 1 - \frac{1}{k-1} + \varepsilon$ and $\alpha = 1 - \frac{1}{k-1} + \frac{\varepsilon}{2}$. Then Lemma 7.2 tells us that the number of m-sets $M \subseteq V(G)$ with $e(M) \geqslant (1 - \frac{1}{k-1} + \frac{\varepsilon}{2})$ is at least $\frac{\varepsilon}{2}\binom{n}{m}$.

Every such m-set M has strictly more than $t_{k-1}(m)$ 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 7.2, 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. ☐

8 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 8.1. *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

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

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

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

by assumption. □

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

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

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

By the supersaturation theorem, Theorem 7.1, 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

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

Recall that by Theorem 6.5, we have that

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

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. □

9 Beyond the Erdős–Stone–Simonovits theorem

The Erdős–Stone–Simonovits theorem gives a very satisfactory asymptotic answer to the question of how large $\mathrm{ex}(n, H)$ is for any non-bipartite H. However, we could still ask about more precise information. For example, when H is a complete graph, Turán's theorem gives us the *exact* value of $\mathrm{ex}(n, H)$ for all n, as well as a description of the unique extremal graph. Can we get something like this for more general graphs H?

Unfortunately, the answer is "no" in general. And the reason, somewhat surprisingly, is again that we don't understand bipartite graphs that well! In fact, the extremal theory of bipartite graphs is crucial to understanding the extremal theory of general graphs.

To understand this connection, let's begin with a simple example that you already saw on the homework. Let H_0 be the following graph:

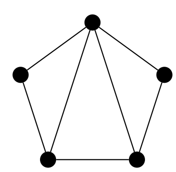

It is not hard to verify that $\chi(H_0) = 3$, so the Erdős–Stone–Simonovits theorem implies that $\mathrm{ex}(n, H_0) \leqslant t_2(n) + o(n^2)$. However, unlike the case of triangles, where we know that $\mathrm{ex}(n, K_3) = t_2(n)$ exactly, for this graph H_0 we do not have an equality. Indeed, let us begin with the Turán graph $T_2(n)$, and call its two parts A and B. We now add a perfect matching inside A, and let G be the resulting graph. Then $e(G) = t_2(n) + \lfloor n/4 \rfloor$, and we claim that G is H_0-free. This essentially boils down to a case check: we need to show that no matter how we try to embed the five vertices of H_0 into the two parts of G, we will fail. More precisely, no matter how we assign the letters A and B to the vertices of H_0, either we will label two adjacent vertices by B, or we will label three vertices in a path by the letter A. In either case, we see that this would not be a valid embedding of H_0 into G, as B has no edges in G, and A has no two-edge paths.

How can we generalize this simple example? One first natural thing to try is to more generally understand for which graphs H, we do not have the equality $\mathrm{ex}(n, H) = t_{\chi(H)-1}(n)$, that is, for which graphs H is the Turán graph not extremal. Thinking about the example above, one can come up with the following definition.

Definition 9.1. A graph H is called *color-critical* if there is an edge $e \in E(H)$ for which

$$\chi(H \setminus e) < \chi(H).$$

That is, H is color-critical if we can decrease its chromatic number by deleting a single edge.

The example of H_0 discussed above readily extends to the following simple fact.

Proposition 9.2. *Let H be a graph with $\chi(H) \geqslant 3$. If H is not color-critical, then*

$$\mathrm{ex}(n, H) > t_{\chi(H)-1}(n)$$

for all n.

Proof. Let G be obtained from the Turán graph $T_{\chi(H)-1}(n)$ by adding a single edge in one of the parts (say the first part, for concreteness). We claim that G is H-free. Note that this suffices to prove the proposition, as it implies that

$$\mathrm{ex}(n, H) \geqslant e(G) = t_{\chi(H)-1}(n) + 1 > t_{\chi(H)-1}(n).$$

So suppose for contradiction that there were some copy of H in G. We use this to define a function $f : V(H) \to [\![\chi(H) - 1]\!]$ by mapping each vertex of H to the label of the part of G containing that vertex. Note that f cannot be a proper coloring of H, as this would contradict the definition of the chromatic number. Nonetheless, f is "almost" a proper coloring: at most one pair of adjacent vertices receive the same value under f, and this value must be 1. Indeed, since G contains no edges inside any part except the first one, the only way we can get an edge both of whose endpoints have the same value is if $f(u) = f(v) = 1$, and the copy of H in G uses the edge we inserted as the edge between u and v.

But this exactly shows that f *is* a proper $(\chi(H) - 1)$-coloring of $H \setminus e$, where $e = uv$. That is, f witnesses that

$$\chi(H \setminus e) \leqslant \chi(H) - 1,$$

contradicting our assumption that H is not color-critical. □

A rather amazing theorem of Simonovits shows that this simple necessary condition is actually sufficient!

Theorem 9.3 (Simonovits 1968)**.** *Let H be a color-critical graph with $\chi(H) \geqslant 3$. Then for all sufficiently large n,*

$$\mathrm{ex}(n, H) = t_{\chi(H)-1}(n).$$

Morever, $T_{\chi(H)-1}(n)$ is the unique extremal graph.

For example, all complete graphs are color-critical, so this recovers Turán's theorem (at least for sufficiently large n). But it does more; for example, one can check that every odd cycle is color-critical, so Simonovits' theorem implies that

$$\mathrm{ex}(n, C_{2k+1}) = \left\lfloor \frac{n^2}{4} \right\rfloor$$

for all sufficiently large n. Note that the requirement that n be sufficiently large is necessary, since, for example, K_4 is C_5-free, hence

$$\mathrm{ex}(4, C_5) = \binom{4}{2} = 6 > \left\lfloor \frac{4^2}{4} \right\rfloor.$$

We will shortly see a proof of Theorem 9.3 (at least in the case $H = C_5$), but before we do, let us try to learn a bit more about $\mathrm{ex}(n, H)$ when H is not color-critical. One can do better than the argument we used in Proposition 9.2, as follows.

Definition 9.4. Let H be a graph with $\chi(H) \geqslant 3$. A *pseudo-coloring* of H is a function $f : V(H) \to [\![\chi(H) - 1]\!]$ with the property that for all $uv \in E(H)$, either $f(u) \neq f(v)$ or $f(u) = f(v) = 1$. That is, adjacent vertices receive different colors, or they both receive color 1.

Now, let B be a bipartite graph. We say that B is *in the decomposition family* of H if for every pseudo-coloring $f : V(H) \to [\![\chi(H) - 1]\!]$, there is a copy of B among the vertices colored 1 by f.

As an example, consider the octahedron graph O_3:

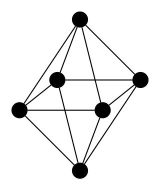

It is not hard to see that $\chi(O_3) = 3$, and with a little casework, one can verify that C_4 is in the decomposition family of O_3.

By adapting the proof of Proposition 9.2, we obtain the following.

Proposition 9.5. *Let H be a graph with $\chi(H) \geqslant 3$, and let B be a bipartite graph in the decomposition family of H. We have that*

$$\mathrm{ex}(n, H) \geqslant t_{\chi(H)-1}(n) + \mathrm{ex}\left(\left\lceil \frac{n}{\chi(H) - 1} \right\rceil, B\right).$$

Proof sketch. Let G be obtained from $T_{\chi(H)-1}$ by inserting an extremal B-free graph into the largest part. By the definition of the extremal function of B, we have that

$$e(G) = t_{\chi(H)-1}(n) + \mathrm{ex}\left(\left\lceil \frac{n}{\chi(H) - 1} \right\rceil, B\right).$$

Moreover, G is H-free, for essentially the same reason as in the proof of Proposition 9.2: if there were a copy of H in G, then we would obtain a pseudo-coloring of H in which the vertices of color 1 are B-free, contradicting the assumption that B is in the decomposition family of H. □

Concretely, this argument demonstrates that

$$\mathrm{ex}(n, O_3) \geqslant \left\lfloor \frac{n^2}{4} \right\rfloor + \Omega(n^{3/2}),$$

by our known lower bound for the extremal number of C_4. In fact, a famous theorem of Erdős and Simonovits shows that this is tight, in a very strong sense: every extremal O_3-free graph is obtained from a complete bipartite graph by putting an extremal C_4-free graph in one part, and a perfect matching in the other part. It is conjectured that this should happen more generally, namely that the extremal number of H should be obtained by inserting certain graphs into the parts of $T_{\chi(H)-1}(n)$, and ensuring that these inserted graphs avoid all members of the decomposition family of H. This conjecture remains open in general (and is probably quite difficult), but is known to hold in certain special cases. And in any case, it demonstrates why determining the exact behavior of $\mathrm{ex}(n, H)$ really requires one to understand extremal numbers of bipartite graphs, as the decomposition family of H is what ends up mattering.