Extremal combinatorics

Zdeněk Dvořák

December 22, 2020

Extremal combinatorics studies how large structures (graphs, set systems, ...) can exist subject to various constraints (forbidden subgraphs, restricted intersections, ...). We will give a couple of examples, without much effort for a coherent theory.

1 Turn's theorem and Erdős-Stone theorem

How many edges can a graph have without creating a large clique? In this direction, you have probably already seen Mantel's theorem.

Theorem 1 (Mantel)**.** *A triangle-free graph with n vertices has at most $n^2/4$ edges.*

This theorem is tight (for even n), as shown by the bipartite graph $K_{n/2,n/2}$. We might suspect a similar result for larger cliques. The *r-partite Turn graph with n vertices* $T_r(n)$ is the complete r-partite graph where the sizes of any two of its parts differ by at most 1 (i.e., they are all as close to n/r as possible). Let $t_r(n)$ be the number of edges of $T_r(n)$; note that $t_r(n) \le (1-1/r)n^2/2$. Clearly, $T_r(n)$ does not contain a clique of size $r+1$, and it indeed turns out that $t_r(n)$ is the largest number of edges an n-vertex graph with this property can have.

Theorem 2 (Turn)**.** *Let G be a graph with n vertices and let r be a positive integer. If $\omega(G) \le r$, then $|E(G)| \le t_r(n)$. Moreover, if $|E(G)| = t_r(n)$, then G is isomorphic to $T_r(n)$.*

Proof. It suffices to prove that if G is an n-vertex graph satisfying $\omega(G) \le r$ with the largest possible number of edges, then G is isomorphic to $T_r(n)$. This is clear if $n \le r$ (as then $T_r(n) = K_n$), and thus we can assume $n > r$.

First, we claim that any non-adjacent vertices in G have the same degree. Indeed, if say $v_1, v_2 \in V(G)$ were non-adjacent vertices such that $\deg v_1 <$

$\deg v_2$, then consider the graph G' obtained from $G - v_1$ by duplicating the vertex v_2. Note that $\omega(G') \leq r$, since duplicating a vertex cannot increase the clique number (any clique in G' containing the duplicate can be turned into a clique in G by replacing the duplicate by v_2). However, $|E(G')| = |E(G)| - \deg v_1 + \deg v_2 > |E(G)|$, contradicting the maximality of $|E(G)|$.

Next, we claim that if $v_1, v_2, v_3 \in V(G)$ are distinct vertices such that $v_1v_2, v_2v_3 \notin E(G)$, then also $v_1v_3 \notin E(G)$. Suppose for a contradiction that this is not the case, and thus $v_1v_3 \in E(G)$. By the previous paragraph, v_1, v_2, and v_3 all have the same degree d. Consider the graph G'' obtained from $G - \{v_1, v_3\}$ by adding two duplicates of v_2. Clearly, $\omega(G'') \leq r$. Moreover, $|E(G'')| = |E(G)| - (\deg v_1 + \deg v_3 - 1) + 2 \deg v_2 = |E(G)| - (2d - 1) + 2d = |E(G)| + 1$, contradicting the maximality of $|E(G)|$.

Hence, the relation $u \sim v$ iff $uv \notin E(G)$ is an equivalence, and thus G is a complete multipartite graph (the classes of the equivalence are the parts of G). Since G has the maximum possible number of edges subject to the condition $\omega(G) \leq r$, observe that G is r-partite (otherwise you can split one part into two and add edges between them) and the sizes of any two parts differ by at most 1 (otherwise you can move a vertex from a larger part to a smaller one, adjusting the neighborhoods accordingly and increasing the number of edges), and thus G is isomorphic to $T_r(n)$. □

What happens if instead of a clique, we forbid another graph F? Note that the Turn graph $T_r(n)$ (and any of its subgraphs) are r-colorable, and thus if F has chromatic number at least $r + 1$, then $F \not\subseteq T_r(n)$, showing that there exist graphs with at least $t_r(n)$ edges that do not contain F as a subgraph. This is not in general tight, but you cannot go much above this bound.

Theorem 3 (Erdős-Stone)**.** *Let F be a graph of chromatic number $r + 1$. For every $\varepsilon > 0$, there exists n_0 such that any graph with $n \geq n_0$ vertices and at least $(1 - 1/r + \varepsilon)n^2/2$ edges contains F as a subgraph.*

We are not going to prove this theorem; let us just note that Erdős-Stone theorem approximates the exact maximum number of edges to arbitrary precision (for large enough graphs), as long as F is not bipartite. If F is bipartite, then rdős-Stone theorem only tells you that for every $\varepsilon > 0$ and sufficiently large n, the maximum is at most εn^2, while $T_1(n)$ has no edges and thus it does not provide any nontrivial lower bound. Indeed, for bipartite graphs, the maximum is much smaller than quadratic; e.g. you might recall that the maximum number of edges of an n-vertex graph without a 4-cycle is $\Theta(n^{3/2})$. Indeed, a similar (typically non-tight) subquadratic bound holds for all bipartite graphs.

Lemma 4. *Let $F \subseteq K_{a,b}$ be a bipartite graph, where $a \le b$. There exists a constant c such that every n-vertex graph G with at least $cn^{2-1/a}$ edges contains F as a subgraph.*

Proof. It suffices to prove the claim for $F = K_{a,b}$, since if G contains as a subgraph $K_{a,b}$, it also contains all subgraphs of $K_{a,b}$. Let m be the number of $(a+1)$-tuples $(x, y_1, y_2, \ldots, y_a)$ of vertices of G such that $xy_1, \ldots, xy_a \in E(G)$. On one hand, we can choose x arbitrarily and then choose $y_1, \ldots, y_a$ among its neighbors, giving

$$m = \sum_{x \in V(G)} \deg^a x \ge \frac{\left(\sum_{x \in V(G)} \deg x\right)^a}{n^{a-1}} = \frac{(2|E(G)|)^a}{n^{a-1}}.$$

On the other hand, we can start by choosing $y_1, \ldots, y_a$. If these vertices are pairwise distinct, then since $K_{a,b} \not\subseteq G$, they have at most $b-1$ common neighbors that can play the role of x. On the other hand, if they are not pairwise distinct, there might be up to n choices for x, but there are at most $(a-1)an^{a-1}$ ways how to select $y_1, \ldots, y_a$ (choose $a-1$ vertices, then choose one of them and a position in $\{1, \ldots, a\}$ to which it is copied). Therefore,

$$m \le (b-1)n^a + (a-1)an^a.$$

Combining these inequalities, we obtain

$$|E(G)| \le \tfrac{1}{2}(b-1+(a-1)a)^{1/a} n^{2-1/a}.$$

□

2 Set systems

How many subsets of size r can you select from $\{1, \ldots, n\}$ so that they all pairwise intersect? If $r > n/2$, you can choose all $\binom{n}{r}$ subsets, which is not very interesting. If $r \le n/2$, you can select all subsets of size r that contain the element n; there are $\binom{n-1}{r-1}$ of them. Can you do better?

Theorem 5 (Erdős-Ko-Rado)**.** *If $A_1, \ldots, A_m$ are distinct pairwise intersecting subsets of $\{1, \ldots, n\}$ of size $r \le n/2$, then*

$$m \le \binom{n-1}{r-1}.$$

Proof. Let c be the number of pairs (C, A), where C is a directed cycle on vertices $\{1, \ldots, n\}$, A is a subpath of C with r vertices, and $V(A)$ is equal to one $A_1, \ldots, A_m$. On one hand, we can start by forming vertices of one of the sets into a path (this can be done in $mr!$ ways), then forming the rest of the vertices into a path (in $(n-r)!$ ways), then concatenating the paths to form the cycle C. This gives

$$c = mr!(n-r)!.$$

Conversely, we can start by forming the cycle C (in $(n-1)!$ ways), and observing that from C, we can select at most r pairwise intersecting r-vertex paths (once we select one such path $a_1 a_2 \ldots a_r$, any other such path starts either immediately to the left or immediately to the right from an edge $a_i a_{i+1}$ for $i \in \{1, \ldots, r-1\}$, and we cannot have both a path starting to the left and to the right from the same edge, as they would be disjoint). Hence,

$$c \leq (n-1)!r.$$

Together, the inequalities give

$$m \leq \frac{(n-1)!r}{r!(n-r)!} = \binom{n-1}{r-1}.$$

□

3 Points in convex position

It is easy to see that among any 5-points in the plane in general position (no three on the same line), you can choose four that are in the convex position. Can we find a larger convex set?

Theorem 6 (Erdős-Szekeres)**.** *Let $n \geq 2$ be an integer. Any set of at least $\binom{2n-4}{n-2} + 1$ points in the plane in general position contains r that are in the convex position.*

Proof. Without loss of generality (by rotating the set if necessary), we can assume no two of the points have the same x-coordinate. An *m-cup* is a sequence $p_1, p_2, \ldots, p_m$ of points with increasing x-coordinates in convex position such that all of them are at or below the line passing through p_1 and p_m. An *m-cap* is such a sequence where all the points are at or above this line. We will prove that any set of at least $\binom{a+b-4}{a-2} + 1$ points in the plane in general position and with pairwise distinct x-coordinates contains either an a-cup or a b-cap.

We prove this by induction on $a+b$. The claim is trivial if $\min(a,b)=2$, and thus we can assume $a,b \geq 3$. Suppose that X contains neither an a-cup nor a b-cap. Let $A \subseteq X$ consist of the rightmost points of all $(a-1)$-cups and $B \subseteq X$ of the leftmost points of all $(b-1)$-cups in X. Note that $X \setminus A$ does not contain any $(a-1)$-cup, and thus by the induction hypothesis, we have $|X \setminus A| \leq \binom{a+b-5}{a-3}$. Similarly, $|X \setminus B| \leq \binom{a+b-5}{a-2}$. Consequently,

$$|X|-|A\cap B| = |(X\setminus A)\cup(X\setminus B)| \leq \binom{a+b-5}{a-3}+\binom{a+b-5}{a-2} = \binom{a+b-4}{a-2} < |X|,$$

and thus there exists a point $p \in A \cap B$. Let $p_1,\ldots,p_{a-1}$ be a cup and $q_1,\ldots,q_{b-1}$ a cap with $p_{a-1}=p=q_1$. If q_2 is above the line $p_{a-2}p_{a-1}$, then $p_1,\ldots,p_{a-1},q_2$ is an a-cup. Otherwise, p_{a-2} is below the line q_1q_2, and thus $p_{a-1},q_1,\ldots,q_{b-1}$ is a b-cap. $\square$