New Bounds for Induced Turán Problems

Nathan S. Sheffield

Abstract

In a recent paper, Hunter, Milojević, Sudakov and Tomon consider $\mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\})$, the maximum number of edges in an n-vertex graph containing no copy of the complete bipartite graph $K_{s,s}$ and no *induced* copy of a pattern graph H. They conjecture that this quantity can be at most a constant factor larger than the standard extremal number of H; that is, that $\mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\}) \leq O(\mathrm{ex}(n, H))$. We make progress towards resolving this conjecture. In particular, we show the following:

i) As originally stated, the conjecture has a simple counterexample, where H is the graph with 3 vertices and 1 edge. However, one may conjecture that the statement holds for all connected H; we show that this would also rule out disconnected counterexamples with more than 1 edge.

ii) The induced extremal number of a graph differs by at most $O(n)$ from that of its 2-core; by a result of Hunter, Milojević, Sudakov and Tomon, this allows us to recover in the induced setting known extremal number upper bounds for graphs with a single cycle.

iii) If H is r-degenerate, $\mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\}) \leq O(n^{2-1/(20r^4)})$. Since the degeneracy of a graph is known to control the standard extremal number, this establishes a nontrivial relationship between $\mathrm{ex}(n, H)$ and $\mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\})$.

We also propose potential connected counterexamples from incidence geometry but are unable to determine whether the conjecture holds in these cases due to the difficulty of bounding standard extremal numbers.

1 Introduction

The extremal number of a graph H, denoted $\mathrm{ex}(n, H)$, is the maximum number of edges in an n-vertex graph with no copy of H as a subgraph. It is known due to Erdős, Stone and Simonovits that $\mathrm{ex}(n, H) = n^2\left(1 - \frac{1}{\chi(H)-1}\right) + o(n^2)$ for all H, where $\chi(H)$ is the chromatic number [ES46; ES66]. However, for bipartite graphs, all this gives is $\mathrm{ex}(n, H) \leq o(n^2)$; understanding the extremal exponents of bipartite graphs is a major ongoing area of research.

What if, instead of just finding H as a subgraph, we are interested in finding an *induced* copy of H? We could define $\mathrm{ex}(n, H\text{-ind})$ to be the maximum number of edges in an n-vertex graph with no induced copy of H. However, this notion is rather uninteresting: unless H is itself complete, the complete graph K_n avoids it as an induced subgraph, so $\mathrm{ex}(n, H\text{-ind}) = \binom{n}{2}$. One might ask, though, whether this is in some sense the *only* way to avoid induced copies of H without avoiding H altogether. That is, perhaps if a graph has enough edges to force many copies of H, but contains no induced copy of H, there must exist some portion of the graph with very high density. Hunter, Milojević, Sudakov and Tomon propose the following conjecture:

Conjecture 1 ([Hun+24])**.** For any $s \in \mathbb{N}$, and any bipartite H,

$$\mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\}) \leq O(\mathrm{ex}(n, H)).$$

We always have $\mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\}) \geq \mathrm{ex}(n, \{K_{s,s}, H\})$, and if $s \geq |V(H)|$ it is clear that $\mathrm{ex}(n, \{K_{s,s}, H\}) = \mathrm{ex}(n, H)$. So, this conjecture is equivalent to the claim that, for sufficiently large s, $\mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\})$ and $\mathrm{ex}(n, H)$ differ by at most a constant factor (where the constant depends on s and H). As a heuristic for why such a conjecture might be reasonable, one could consider searching for copies of H in the Erdős–Rényi random graph $G(n, p)$: if $p \geq \Omega(1)$, there will be $\Theta(n^{2s})$ many copies of $K_{s,s}$, but if $p \leq o(1)$, then all but a subconstant fraction of the copies of H will be induced copies.

Hunter et al provide some evidence for their conjecture by directly reproducing many of the best known asymptotic upper bounds on standard extremal numbers. In particular, they show $\mathrm{ex}(n, \{K_{s,s}, C_{2k}\text{-ind}\}) \leq O(n^{1+1/k})$ for C_{2k} a length-$2k$ cycle, $\mathrm{ex}(n, \{K_{s,s}, Q_8\text{-ind}\}) \leq O(n^{8/5})$ for Q_8 the skeleton of a 3-dimensional cube, and $\mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\}) \leq O(n^{2-1/r})$ for any bipartite H with maximum degree r on one side, all of which match the corresponding bounds known for $\mathrm{ex}(n, H)$ [Hun+24]. Fox, Nenadov and Pham show that this last result can be extended to give $\mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\}) \leq O(n^{2-1/r})$ whenever one side of H has at most r vertices that are complete to the other side, and all other vertices on that side have degree at most r [FNP24]. Axenovich and Zimmerman show that, if H is additionally required to contain no copy of $K_{r,r}$, and the host graph is bipartite, the host graph can have only $o(n^{2-1/r})$ edges [AZ24], providing an induced analogue of a result of Sudakov and Tomon [ST20]. It is also known due to Scott, Seymour and Spirkl that $\mathrm{ex}(n, \{K_{s,s}, T\text{-ind}\}) \leq O(n)$ when T is a tree [SSS23]; Hunter et al show that the constant factor in that result can be made polynomial in s.

1.1 Main results

Our first observation is a simple counterexample to Conjecture 1:

Proposition 1. There exists a bipartite graph H such that $\mathrm{ex}(n, \{K_{2,2}, H\text{-ind}\}) \geq \omega(\mathrm{ex}(n, H))$.

The graph H serving as a counterexample consists of 3 vertices, with 2 joined by a single edge. This seems like a rather exceptional case, which may not preclude the conjecture from being morally true. To rule out such examples, one could modify Conjecture 1 to require H to be connected, or to have more than a single edge — our next observation shows that these two modified statements are equivalent.

Proposition 2. For any $s \in \mathbb{N}$, if H is the disjoint union of two subgraphs H_1 and H_2, then $\mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\}) \leq O(\mathrm{ex}(n, \{K_{s,s}, H_1\text{-ind}\}) + \mathrm{ex}(n, \{K_{s,s}, H_2\text{-ind}\}) + n)$.

One particular consequence of Proposition 2 is that, as long as $\mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\}) \geq \omega(n)$, this value remains unchanged up to constant factors when we remove all isolated vertices from H. For the standard extremal number, it is easy to see that such a statement holds even if we also remove all vertices of degree 1. Our next result reproduces this fact in the induced setting.

Definition 1. For $k \geq 1$, k-core(G) is the largest induced subgraph of G with minimum degree at least k.

Theorem 1. For any $s \in \mathbb{N}$ and any graph H, $\mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\}) \leq O(\mathrm{ex}(n, \{K_{s,s}, 2\text{-core}(H)\text{-ind}\}) + n)$.

Together with the results of [Hun+24], Proposition 2 and Theorem 1 imply that $\mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\}) \leq O(n^{1+1/k})$ for any bipartite H with girth at least $2k$ as long as every connected component of H contains at most a single cycle.

With the goal of determining a relationship between $\mathrm{ex}(n, H)$ and $\mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\})$ in general, we give an upper bound in terms of the degeneracy of H.

Definition 2. For $k \geq 1$, a graph G is called k-degenerate if $(k+1)$-core(G) is empty. We define the degeneracy of G to be the minimum k such that G is k-degenerate.

Although there is no general method known for determining extremal numbers of bipartite graphs, it's known that the degeneracy of the pattern graph always offers a somewhat reasonable proxy: for any bipartite H of degeneracy r, $\Omega(n^{2-2/r}) \leq \mathrm{ex}(n, H) \leq O(n^{2-1/4r})$ [AKS03], and it's conjectured that the upper bound can be strengthened to $O(n^{2-1/r})$ [Erd97]. We show that degeneracy also controls the induced extremal number, although we achieve weaker dependency on r than is known for standard extremal numbers:

Theorem 2. For any $s \in \mathbb{N}$, and any bipartite H of degeneracy r, $\mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\}) \leq O(n^{2-1/(20r^4)})$.

The fact that degeneracy controls both immediately establishes some nontrivial relationship between the standard and induced extremal numbers: if $\mathrm{ex}(n, H) = O(n^\alpha)$ and $\mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\}) = \Omega(n^\beta)$, we can show unconditionally that $\beta \leq 2 - \frac{(2-\alpha)^4}{320}$ (whereas Conjecture 1 would hold that $\beta \leq \alpha$). It would be interesting to quantitatively improve this relationship by showing an upper bound of the form $\mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\}) \leq O(n^{2-1/cr})$ for some constant c independent of r; we make some progress towards this goal by finding copies of H which, while not necessary fully induced, avoid some particular subset of H's non-edges.

Finally, in Section 5, we discuss the possibility that Conjecture 1 could be morally false, as well as false for the rather silly reason of Proposition 1. We note several graphs H where $\mathrm{ex}(n, \{K_{2,2}, H\text{-ind}) = \Theta(n^{3/2})$, but such that there may be heuristic evidence to believe $\mathrm{ex}(n, H) \leq o(n^{3/2})$. However, due to the difficulty of determining standard extremal numbers, we are unable to prove such upper bounds.

1.2 Additional related work

The structure of graphs avoiding given induced subgraphs has been an active area of research for some time. Much existing work is motivated by the Erdős–Hajnal conjecture, which claims that, for any H, any n-vertex graph with no induced copy of H must contain a clique or independent set of size polynomial in n [EH89; Chu14]. Towards this conjecture, Fox and Sudakov have shown that any graph avoiding induced copies of H must contain either a complete bipartite graph or independent set of polynomial size [FS09].

To our knowledge, Hunter et al are the first to systematically consider $\mathrm{ex}(n, \{K_{s,s}, H\text{-ind})$ for general bipartite H, however the problem of forbidding one induced graph and one non-induced graph has recieved prior attention. Kühn and Osthus have shown that $\mathrm{ex}(n, \{K_{s,s}, \mathcal{H}\text{-ind}\}) \leq O(n)$, where $\mathcal{H}$ is the family of all subdivisions of a given graph H [KO04]. Loh, Tait, Timmons and Zhou showed that $\mathrm{ex}(n, \{K_r, K_{s,t}\text{-ind}\}) \leq O(n^{2-1/s})$ for all r, s, t [Loh+18], prompting further consideration of $\mathrm{ex}(n, \{F, H\text{-ind}\})$ for non-bipartite F [EGM19; Ill21a; Ill21b].

2 Disconnected counterexample

We begin by noting our counterexample to Conjecture 1.

Proposition 1. There exists a bipartite graph H such that $\mathrm{ex}(n, \{K_{2,2}, H\text{-ind}\}) \geq \omega(\mathrm{ex}(n, H))$.

Proof. Let H be the graph consisting of two vertices connected by an edge, and a third isolated vertex. Note that

$$\mathrm{ex}(n, H) = \begin{cases} 1 \text{ if } n = 2 \\ 0 \text{ otherwise} \end{cases},$$

since a subgraph isomorphic to H just entails a single edge plus some third vertex not involved in that edge. However, in the star graph $K_{1,n-1}$, any three vertices induce either an independent set or a path, depending on whether one of the three vertices is the center of the star. Since $K_{1,n-1}$ contains no cycles, we have exhibited a graph on $n-1$ many edges with no copy of $K_{2,2}$ and no induced copy of H, giving $\mathrm{ex}(n, \{K_{2,2}, H\text{-ind}\}) \geq n - 1 \geq \omega(\mathrm{ex}(n, H))$. ☐

This certainly disproves Conjecture 1, but in a rather unsatisfying way. For one, H is disconnected — it would be interesting to have a connected counterexample. Another, even more serious, objection is that H has only a single edge, so $\mathrm{ex}(n, H) \leq o(1)$. Any graph with more than one edge is avoided either by a star or a matching, so will have $\mathrm{ex}(n, H) \geq \Omega(n)$; it might well be that Conjecture 1 fails only in the trivial case when $\mathrm{ex}(n, H)$ goes to 0, and so holds whenever H has at least 2 edges. To rule out such examples, we could propose either of the following two modifications of Conjecture 1:

Conjecture 2. For any $s \in \mathbb{N}$, and any *connected* bipartite H,

$$\mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\}) \leq O(\mathrm{ex}(n, H)).$$

Conjecture 3. For any $s \in \mathbb{N}$, and any bipartite H *with at least* 2 *edges*,

$$\mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\}) \leq O(\mathrm{ex}(n, H)).$$

Since the only connected graph with a single edge is $K_{1,1}$, it is clear that Conjecture 3 implies Conjecture 2. Our next observation will show that the reverse implication also holds, so that these two modifications of Conjecture 1 are equivalent.

Proposition 2. For any $s \in \mathbb{N}$, if H is the disjoint union of two subgraphs H_1 and H_2, then $\mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\}) \leq O(\mathrm{ex}(n, \{K_{s,s}, H_1\text{-ind}\}) + \mathrm{ex}(n, \{K_{s,s}, H_2\text{-ind}\}) + n)$.

Proof. The Kővári–Sós–Turán theorem gives that $\mathrm{ex}(n, K_{s,s}) \leq n^{2-1/s}$, which means in particular that any sufficiently large graph with constant edge density must contain a copy of $K_{s,s}$. Let N be such that, for any $N' \geq N$, any N'-vertex graph with at least $\frac{N'^2}{|V(H)|^2}$ many edges must contain $K_{s,s}$ as a subgraph. We will show that, for any n, any n-vertex graph G with at least $\mathrm{ex}(n, \{K_{s,s}, H_1\text{-ind}\}) + \mathrm{ex}(n, \{K_{s,s}, H_2\text{-ind}\}) + N|V(H)|n$ many edges must contain either an induced copy of $H = H_1 \sqcup H_2$, or a copy of $K_{s,s}$.

The proof follows from a simple supersaturation argument. We claim that, if G contains no $K_{s,s}$, it must contain N induced copies of H_1 and N induced copies of H_2, all of which are vertex disjoint. This holds because, if we've already found at most N copies of H_1 and N copies of H_2, the subgraph of G induced by all vertices not included in any of these subgraphs contains at least $E(G) - N|V(H_1)|n - N|V(H_2)|n = E(G) - N|V(H)|n \geq \mathrm{ex}(n, \{K_{s,s}, H_1\text{-ind}\}) + \mathrm{ex}(n, \{K_{s,s}, H_2\text{-ind}\})$ many edges, and thus we can find another disjoint copy of either.

Now, suppose for contradiction that G contains no induced copy of H. This means that each of those copies of H_1 must have an edge to each of those copies of H_2, since otherwise the pair would induce a copy of H. Thus, the subgraph induced by all N disjoint copies of H_1 and H_2 together — which is a graph on $N|V(H)|$ many vertices — contains at least N^2 many edges. By our choice of N, this guarantees that the graph contains a copy of $K_{s,s}$. □

Corollary 1. Conjecture 2 $\iff$ Conjecture 3.

Proof. As noted, the only connected graph with a single edge is $K_{1,1}$, in which case we have $\mathrm{ex}(n, K_{1,1}) = \mathrm{ex}(n, \{K_{s,s}, K_{1,1}\text{-ind}\}) = 0$, so the Conjecture 3 $\implies$ Conjecture 2 direction is trivial. To show the reverse direction, we assume Conjecture 2, and demonstrate Conjecture 3 by contradiction: fix $s \in \mathbb{N}$, and let H be the smallest bipartite graph with at least 2 edges such that $\mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\}) \geq \omega(\mathrm{ex}(n, H))$. H cannot be connected, since otherwise Conjecture 2 would apply. So, we can split H into two disconnected components $H = H_1 \sqcup H_2$. By Proposition 2, we have $\mathrm{ex}(n, \{K_{s,s}, H_1\text{-ind}\}) + \mathrm{ex}(n, \{K_{s,s}, H_2\text{-ind}\}) + O(n) \geq \omega(\mathrm{ex}(n, H))$. Since H has at least two edges, $\mathrm{ex}(n, H) \geq \Omega(n)$, so this implies that either $\mathrm{ex}(n, \{K_{s,s}, H_1\text{-ind}\}) \geq \omega(\mathrm{ex}(n, H)) \geq \omega(\mathrm{ex}(n, H_1))$ or $\mathrm{ex}(n, \{K_{s,s}, H_2\text{-ind}\}) \geq \omega(\mathrm{ex}(n, H)) \geq \omega(\mathrm{ex}(n, H_2))$ — but either of these would contradict minimality of H. □

That is, if Conjecture 1 has any counterexample not of the form "single edge plus some number of isolated vertices", it has a connected counterexample. In the next section, we will show that any connected counterexample implies a counterexample of minimum degree at least 2.

3 Reducing to the 2-core

It is known that Conjecture 1 holds when H is restricted to be a tree — i.e., $\mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\}) \leq O(\mathrm{ex}(n, H)) \leq O(n)$ whenever H is a tree [SSS23; Hun+24]. For the standard extremal number, it's easy to show that degree-one vertices simply do not affect the extremal number: not only is $\mathrm{ex}(n, T) \leq O(n)$ for all trees T, but more generally wedging a tree to any (possibly cyclic) graph H can change H's extremal number by at most an additive $O(n)$. In this section, we will show how to recover a statement of this form in the induced setting, too.

Theorem 1. For any $s \in \mathbb{N}$ and any graph H, $\mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\}) \le O(\mathrm{ex}(n, \{K_{s,s}, 2\text{-core}(H)\text{-ind}\}) + n)$.

The first step in the proof is a standard regularization argument, originally due to Erdős and Simonovits.

Lemma 1 ([ES66]; see also [FS13]). For any $\alpha \in [0,1]$, and any n-vertex graph G with at least $n^{1+\alpha}$ many edges, there exists an induced subgraph $G' \subseteq G$ with $m \ge n^{\frac{\alpha-\alpha^2}{1+\alpha}}$ many vertices and at least $\frac{2}{5}m^{1+\alpha}$ edges, such that the maximum degree in G' is at most a $20 \cdot 2^{1/\alpha^2}$ factor larger than the minimum degree in G'.

The statement in [ES66] does not mention that G' can be taken to be induced, but as noted in [AZ24] this is immediate from the proof.

Lemma 1 will allow us to show the key technical tool in our proof of Theorem 1: a supersaturation result allowing us to find many induced copies of a subgraph that all share only one specified vertex.

Lemma 2. For any H, and any $N \in \mathbb{N}$, there exists some constant C such that for any vertex $v \in V(H)$ and any $K_{s,s}$-free graph G on $C \cdot (\mathrm{ex}(|V(G)|, \{K_{s,s}, H\text{-ind}\}) + |V(G)|)$ many edges, there exist N induced copies of H in G such that any pair of copies overlap on exactly one vertex, and that vertex is the image of v.

Proof. First, note that by Lemma 1, for any C', if C is chosen sufficiently large, we can pass to a subgraph $G' \subseteq G$ with n vertices and at least $C' \cdot (\mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\}) + n)$ many edges, such that G has maximum degree at most $\frac{\Delta \cdot E(G')}{n}$ and minimum degree at least $\frac{E(G')}{\Delta n}$, for $\Delta = 20 \cdot 2^{1/\alpha^2}$. Now, for some k to be determined later, consider the following randomized procedure:

i) Choose a uniform random $\frac{n}{k}$ vertices $R \subseteq V(G')$.
ii) Choose a uniform random vertex $u \in R$.
iii) Declare the procedure to have succeeded if there exists an induced embedding π of H in R such that $\pi(v) = u$.

We claim that, as long as C' is sufficiently large, this procedure has reasonably high success probability. The first necessary observation is that, with high probability, the graph remains nearly regular upon subsampling to R.

Claim 1. Except with probability $2^{-\mathsf{poly}(C')}$, every vertex in R has at least $\frac{|E(G')|}{2k\Delta}$ and at most $\frac{2\Delta \cdot |E(G')|}{k}$ many neighbours in R.

Proof. For any specific vertex $x \in V(G')$, consider the probability that x has fewer than $\frac{|E(G')|}{2k\Delta n}$ neighbours in R. We know that x has degree at least $\frac{|E(G')|}{\Delta n}$, so this probability is at most the chance that fewer than a $\frac{1}{2k}$ fraction of x's neighbours are chosen to belong to R. Since R is a uniform random $1/k$-fraction of all vertices, a Chernoff bound guarantees that this probability is exponentially small in the size of x's neighbourhood, which is at least $\Omega(C')$. Union bounding over all $x \in V(G')$, this means that the probability that any of them have fewer than $\frac{|E(G')|}{2k\Delta n}$ neighbours in R is exponentially small. The upper bound on neighbourhood size is identical. $\square$

We then observe that, so long as R remains nearly regular, we can find induced copies of H within R making use of a substantial fraction of the vertices.

Claim 2. If the subgraph induced by R has $|E(R)| \ge 2 \cdot \mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\})$, and maximum degree at most a $4\Delta^2$ factor larger than minimum degree, then for at least $\frac{|V(R)|}{8\Delta^2}$ many vertices $u \in R$, there exists an induced embedding π of H in R such that $\pi(v) = u$.

Proof. Suppose there are fewer than $\frac{|V(R)|}{8\Delta^2}$ many vertices of R that serve as the image of v under some embedding of H. Since the graph induced by R has maximum degree at most $\frac{4\Delta^2 \cdot |E(R)|}{|V(R)|}$, the subgraph induced by all vertices of R other than those contains at least $|E(R)| - \left(\frac{|V(R)|}{8\Delta^2}\right)\left(\frac{4\Delta^2 \cdot |E(R)|}{|V(R)|}\right) = \frac{|E(R)|}{2} \ge \mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\})$ many edges. So, it must contain an induced copy of H, which must have some vertex corresponding to v, which is a contradiction since we've removed all possible images. $\square$

Together, Claim 1 and Claim 2 guarantee that our random process has success probability at least $\frac{1}{10\Delta^2}$ if C is chosen to be sufficiently large in terms of k, since once we condition on an event with probability going to 1 in C' a random vertex in R has at least a $\frac{|V(R)|}{8\Delta^2}$ chance of leading to success.

But now, note that we could alternatively have performed the procedure in the following order:

i) Choose a uniform random vertex $u \in V(G')$.
ii) Choose a uniform random partition of all other vertices of G into k equal-sized colour classes.
iii) Choose one of the colours uniformly at random to call R.
iv) Declare the procedure to have succeeded if there exists an induced embedding π of H in $R \cup \{u\}$ such that $\pi(v) = u$.

Up to an additive difference of 1 in the number of vertices chosen in R, which affects the distribution negligibly, this process gives the same distribution over R and u as the one originally specified, and so has the same success probability of at least $\frac{1}{10\Delta^2}$. By averaging, there exists some way to perform steps i and ii such that the process still has success probability at least $\frac{1}{10\Delta^2}$ over step iii. This means that, for some vertex $u \in V(G')$ and some partition into k colours, at least $\frac{k}{10\Delta^2}$ of the colour classes contain an induced copy of H mapping v to u. Since each of these copies are (aside from u) of different vertex colours and hence disjoint, choosing $k = 10\Delta^2 N$ gives the desired statement. □

Proof of Theorem 1. Proceeding by induction on the number of edges of H, it suffices to show for all H that $\mathrm{ex}(n, \{K_{s,s}, H^+\text{-ind}) \leq O(\mathrm{ex}(n, \{K_{s,s}, H\text{-ind}) + n)$, where H^+ is obtained from H by adding a single vertex u, and a single edge (u, v) to some $v \in V(H)$. That is, we will show that adding a single degree-1 vertex to H cannot change its induced extremal number by more than a constant multiplicative factor and an additive linear factor in n.

By Lemma 2, there exists some C such that any n-vertex graph with at least $C \cdot \mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\})$ many edges must contain s many copies of H, any pair of which overlap exactly on the vertex corresponding to v. Let G be a graph with at least $2C \cdot \mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\}) + 2\,(s(|V(H)| - 1)^s + s(|V(H)| - 1))\, n$ many edges. By repeatedly removing vertices with degree less than half the average, we can find an induced subgraph $G' \subseteq G$ with minimum degree at least $\frac{C}{n}\,\mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\}) + s(|V(H)| - 1)^s + s(|V(H)| - 1)$. Now, by Lemma 2, we can find induced copies $H_1, \ldots, H_s$ of H in G' that all overlap only on $\pi(v)$. By our bound on minimum degree, we know that $\pi(v)$ has degree at least $s(|V(H)| - 1)^s + (|V(H)| - 1)$, and hence has at least $s(|V(H)| - 1)^s$ many neighbours $u_1, \ldots, u_{(s(|V(H)|-1)^s)}$ not contained in any of our identified H_i.

If G' contains no induced copy of H^+, then for every i, j, there must be an edge between u_i and some vertex of $H_j \setminus \{\pi(v)\}$; choose an arbitrary such edge for each i, j. There are only $(|V(H)| - 1)^s$ many ways to choose one element of $H_j \setminus \{\pi(v)\}$ for each j, so by pigeonhole principle there must exist a set of s many u_i for which we've chosen the exact same s-tuple of neighbours. These vertices and that s-tuple of neighbours form a copy of $K_{s,s}$. □

Corollary 2. If H is any bipartite graph of girth at least $2k$ such that each connected component contains at most a single cycle, then for any s we have $\mathrm{ex}(n, \{K_{s,s}, H\}) \leq O(n^{1+1/k})$.

Proof. Hunter et al have shown that that $\mathrm{ex}(n, \{K_{s,s}, C_{2\ell}) \leq O(n^{1+1/\ell})$ for all ℓ [Hun+24]. By Theorem 1, this gives that $\mathrm{ex}(n, \{K_{s,s}, F) \leq O(n^{1+1/k})$ for any F whose 2-core is $C_{2\ell}$, $\ell \geq k$ – in other words, any bipartite graph of girth at least $2k$ containing only a single cycle. Then, by Proposition 2, the disjoint union H of any constant number of such graphs must also have $\mathrm{ex}(n, \{K_{s,s}, H\}) \leq O(n^{1+1/k})$. □

4 Control by degeneracy

The results of Section 3 give us a new class of graphs where we have upper bounds on induced extremal numbers matching those for standard extremal numbers. However, we still have no relationship between induced and standard extremal numbers in general. Short of a proof of Conjecture 2, it would be useful to at least rule out that $\mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\})$ can be arbitrarily large in terms of $\mathrm{ex}(n, H)$. That is, to show that for every ε there exists a δ such that, for all H, if $\mathrm{ex}(n, H) \leq n^{2-\varepsilon}$, then $\mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\}) \leq n^{2-\delta}$. In this section, we will obtain such a result.

The relevant fact is that, while we know no general way of computing extremal numbers, the degeneracy of H offers a reasonable approximation, giving both lower and upper bounds. Specifically, we know that, for some constants $c > k$, for any r and any bipartite H of degeneracy r, $\Omega(n^{1-1/kr}) \leq \mathrm{ex}(n, H) \leq O(n^{1-1/cr})$. (The current best known values of c and k are $1/2$ and 4, respectively, although it is conjectured that these can be improved [AKS03; Erd97].) An upper bound on *induced* extremal numbers in terms of degeneracy would therefore allow us to constrain the induced extremal number of a graph in terms only of its non-induced extremal number. In Section 4.1, we will show a bound of the form $\mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\}) \leq O(n^{2-1/cr^c})$, which will establish some such relationship. To strengthen that relationship, it would be interesting to show a bound of the form $\mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\}) \leq O(n^{2-1/cr})$; in Section 4.2 we discuss partial progress towards such a stronger quantitative bound.

4.1 Bounding induced extremal numbers in terms of degeneracy

As in the argument for standard extremal numbers, our upper bound relies on the technique of dependent random choice. In this section, we will be able to use directly the following standard result, whereas in the next section we will have to unfold its proof to obtain some stronger guarantees:

Lemma 3 ([AKS03], [FS11])**.** For any $r, t \geq 2$, and any n-vertex graph G with at least $n^{2-1/(t^3 r)}$ many edges, there exist subsets $U_1, U_2 \subseteq V(G)$ such that every r-tuple of vertices in U_1 has at least $n^{1-1.8/t}$ many common neighbours in U_2, and likewise every r-tuple of vertices in U_2 has at least $n^{1-1.8/t}$ many common neighbours in U_1.

We will also make use of the fact that, in a $K_{s,s}$-free graph, at most a constant number of vertices are neighbours with a constant fraction of any sufficiently large vertex set.

Lemma 4. For any $\varepsilon > 0$, any $K_{s,s}$-free graph G, and any vertex subset $S \subseteq V(G)$ with $|S| \geq \frac{2s}{\varepsilon}$, there exist at most $\left(\frac{2s}{\varepsilon}\right)^s$ many vertices $v \in V(G)$ such that $|N(v) \cap S| \geq \varepsilon|S|$.

Proof. Suppose there exist $x = \left(\frac{2s}{\varepsilon}\right)^s$ many vertices $v \in V(G)$ such that $|N(v) \cap S| \geq \varepsilon|S|$. Taking these vertices along with S gives a bipartite graph with part sizes x and $|S|$ and at least $x \cdot \varepsilon|S| - s^2 = 2sx^{1-1/s}|S| - s^2 \geq s^{1/s}x^{1-1/s}|S| + sx$ many edges. The asymmetric version of the Kővári–Sós–Turán theorem guarantees that any such graph must contain a copy of $K_{s,s}$ [KST54; Hyl58]. □

In order to find an induced embedding, we will first apply Lemma 3, then use Lemma 4 to show that an appropriately-chosen random embedding of our pattern graph H in the resulting pair of subsets will be an induced copy with high probability.

Theorem 2. For any $s \in \mathbb{N}$, and any bipartite H of degeneracy r, $\mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\}) \leq O(n^{2-1/(20r^4)})$.

Proof. Let H be an r-degenerate graph, and let G be a $K_{s,s}$-free n-vertex graph with at least $n^{2-1/(20r^4)}$ many edges. Applying Lemma 3 with $t = 2.71r$, we obtain two vertex subsets $U_1, U_2 \subseteq V(G)$ such that every r-tuple in one subset has at least $n^{1-1/1.51r}$ many common neighbours in the other.

Let $v_1, \ldots, v_{|V(H)|}$ be an ordering of the vertices of H such that, for all i, v_i has at most r many neighbours v_j with $j < i$. Such an ordering is guaranteed to exist by the fact that H is r-degenerate. Now, for any tuple of numbers $w = (w_1, \ldots, w_{|V(H)|}) \in [n^{1-1/1.51r}]^{|V(H)|}$ such that $w_i \neq w_j$ for all $i \neq j$, we will define an associated embedding $\pi_w : V(H) \to V(G)$. Fix an arbitrary ordering of the vertices in G's left and right parts, respectively. For each v_i in the left (resp. right) part of H, let $\pi_w(v_i)$ be the w_ith vertex of the common neighbourhood $\bigcap_{j<i\colon (v_i,v_j)\in E(H)} N(\pi_w(v_j))$ in the ordering of the right (resp. left) part of G. Since each v_i has at most r earlier neighbours in the embedding, and any r vertices in one part have at least $n^{1-1/1.51r}$ many common neighbours, the w_ith vertex defined thus will always exist.

The image of each π_w is a homomorphic copy of H: if $(v_i, v_j) \in E(H)$, then $(\pi_w(v_i), \pi_w(v_j)) \in E(G)$, since whichever of v_i and v_j is later in the ordering will be chosen from the neighbourhood of the other. We claim that when w is chosen uniformly at random among all elements of $[n^{1-1/1.51r}]^{|V(H)|}$, the associated homomorphic copy of H has nonzero probability of being an induced subgraph.

For an r-tuple $u_1, \dots, u_r$ of vertices in one of the two parts of $V(G)$, and another vertex v, we say that v ***electrocutes*** $u_1, \dots, u_r$ if v is adjacent to a large fraction of a prefix of the common neighbourhood of the u — that is, if $|N(v) \cap A| \geq \frac{1}{100|V(H)|^2} \cdot n^{1-1/1.51r}$, where A consists of the first $n^{1-1/1.51r}$ many vertices of $\bigcap_i N(u_i)$ in the ordering of the other part of $V(G)$ from the u_i (note that our dependent random choice guarantees $|\bigcap_i N(u_i)| \geq n^{1-1/1.51r}$). Call a set $T \subseteq V(G)$ ***slippery*** if some r-tuple $(u_1, \dots, u_r) \in T^r$ of the vertices is electrocuted by another vertex $v \in T$, $v \notin (u_1, \dots, u_r)$.

Claim 3. If n is sufficiently large, and w is chosen uniformly from $[n^{1-1/1.51r}]^{|V(H)|}$, the image $\mathrm{Im}_{\pi_w}(V(H))$ is slippery with probability at most $\frac{1}{100|V(H)|^2}$.

Proof. We can describe every slippery w as follows:

i) Choose an index i, and indices $j_1, \dots, j_r$.
ii) Choose r vertices from $V(G)$ to serve as $\pi_w(v_{j_1}), \dots, \pi_w(v_{j_r})$.
iii) Choose the entries w_ℓ for $\ell \notin \{i, j_1, \dots, j_r\}$.
iv) Choose the entry w_i, ensuring that $\pi_w(v_i)$ electrocutes $\pi_w(v_{j_1}), \dots, \pi_w(v_{j_r})$.
v) Set $w_{j_1}, \dots, w_{j_r}$ to be the unique values such that $\pi_w(v_{j_1}), \dots, \pi_w(v_{j_r})$ correspond to the vertices chosen on step ii.

By upper bounding the number of available choices at each step, we can obtain an upper bound on the number of slippery w.

i) There are at most $|V(H)|^{r+1}$ many ways to choose the indices
ii) There are at most n^k many ways to choose $\pi_w(v_{j_1}), \dots, \pi_w(v_{j_r})$, where k is the number of distinct elements appearing among $j_1, \dots, j_r$ (note that the r-tuple of indices may contain repeated elements).
iii) There are at most $(n^{1-1/1.51r})^{|V(H)|-k-1}$ many ways to choose w_ℓ for $\ell \notin \{i, j_1, \dots, j_r\}$.
iv) By Lemma 4, so long as $n^{1-1/1.51r} \geq 200s|V(H)|^2$, there are at most $\left(200s|V(H)|^2\right)^s$ many vertices that electrocute $\pi_w(v_{j_1}), \dots, \pi_w(v_{j_r})$.

Overall, this means that the number of slippery w is at most

$$\begin{aligned}
&|V(H)|^{r+1} \cdot n^k \cdot (n^{1-1/1.51r})^{|V(H)|-k-1} \cdot \left(200s|V(H)|^2\right)^s \\
&= (n^{1-1/1.51r})^{\left(|V(H)|-k-1+\frac{k}{1-1/1.51r}\right)} \cdot (200s)^s\,|V(H)|^{2s+r+1} \\
&\leq (n^{1-1/1.51r})^{|V(H)|} \cdot n^{-\left(\frac{.51r-1}{1.51r-1}\right)} \cdot (200s)^s\,|V(H)|^{2s+r+1}.
\end{aligned}$$

Since we know any 1-degenerate H has extremal number $O(n)$, we can assume $r > 1$, in which case $\left(\frac{.51r-1}{1.51r-1}\right) > 0$. So, for any constant values of s, r, and $|V(H)|$, for n sufficiently large we have $n^{\left(\frac{.51r-1}{1.51r-1}\right)} > 100\,(200s)^s\,|V(H)|^{2s+r+3}$, which means that the number of slippery w is at most $\frac{(n^{1-1/1.51r})^{|V(H)|}}{100|V(H)|^2}$. Since the total number of w is exactly $(n^{1-1/1.51r})^{|V(H)|}$, the slippery tuples represent less than a $\frac{1}{100|V(H)|^2}$ fraction. □

We can now show that any particular non-edge of H is absent in $\pi_w(H)$ with good probability.

Claim 4. Let $v_i, v_j \in V(H)$ be any pair of vertices such that $(v_i, v_j) \notin E(H)$. If n is sufficiently large, and w is chosen uniformly from $[Cn^{1-1/r}]^{|V(H)|}$, then $\Pr_w[(\pi_w(v_i), \pi_w(v_j)) \in E(G)] \leq \frac{1}{50|V(H)|^2}$.

Proof. We can assume without loss of generality that $i < j$. Let the tuple $u_1, \dots, u_r$ contain the neighbours of v_i in H that appear earlier in the degeneracy ordering (repeat a vertex in the tuple if there are fewer than r such distinct vertices). If we fix random values for $w_1, \dots, w_{j-1}$, this will in particular fix the embeddings $\pi_w(u_1), \dots, \pi_w(u_r)$ for all of those neighbours, as well as the embedding $\pi_w(v_i)$.

Let A consist of the first $n^{1-1/1.51r}$ many vertices of $\bigcap_\ell N(u_\ell)$ in the ordering of the part of $V(G)$ containing $\pi_w(u_1), \dots, \pi_w(u_r)$; choosing a random value for w_j will correspond to fixing $\pi_w(v_j)$ to a uniform random element of A. By definition, unless $\pi(v_i)$ electrocutes $\pi_w(u_1), \dots, \pi_w(u_r)$, at most a $\frac{1}{100|V(H)|^2}$ fraction of the vertices of A are adjacent to $\pi_(v_i)$, so conditional on $\pi(v_i)$ not electrocuting $\pi_w(u_1), \dots, \pi_w(u_r)$ we have $\Pr_{w_j}[(\pi_w(v_i), \pi_w(v_j)) \in E(G)] \le \frac{1}{100|V(H)|^2}$. If $\pi(v_i)$ electrocutes $\pi_w(u_1), \dots, \pi_w(u_r)$, w is slippery — so Claim 3 ensures that this occurs with probability at most $\frac{1}{100|V(H)|^2}$. By union bound, this means $\Pr_w[(\pi_w(v_i), \pi_w(v_j)) \in E(G)] \le \frac{1}{100|V(H)|^2} + \frac{1}{100|V(H)|^2} = \frac{1}{50|V(H)|^2}$. ☐

Finally, we note that π_w is injective with high probability.

Claim 5. For any $i < j \le r$, we have $\Pr_w[\pi_w(v_i) = \pi_w(v_j)] \le \frac{1}{n^{1-1/1.51}}$.

Proof. When w_j is chosen, there are $n^{1-1/1.51}$ many options, at most one of which corresponds to $\pi_w(v_i)$. ☐

Now, the probability that $\mathrm{Im}_{\pi_w}(V(H))$ fails to be an induced copy of H is by union bound at most

$$\left(\sum_{i,j} \Pr_w[\pi_w(v_i) = \pi_w(v_j)]\right) + \left(\sum_{i,j\colon (v_i,v_j)\notin E(H)} \Pr[(\pi_w(v_i), \pi_w(v_j)) \in E(G)]\right)$$
$$\le |V(H)|^2 \cdot \frac{1}{n^{1-1/1.51}} + |V(H)|^2 \cdot \frac{1}{50|V(H)|^2}$$
$$\le \frac{1}{25}$$

for sufficiently large n. Since this probability is less than 1, we know in particular that G contains an induced copy of H. ☐

Corollary 3. For any constant α, if $\mathrm{ex}(n, H) \le O(n^\alpha)$, then $\mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\}) \le O\left(n^{\left(2-\frac{(2-\alpha)^4}{320}\right)}\right)$.

Proof. Let r be the degeneracy of H. Since $\Omega(n^{1-2/r}) \le \mathrm{ex}(n, H) \le O(n^\alpha)$, we must have $r \le \frac{2}{2-\alpha}$. The result now follows from $\mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\}) \le n^{2-1/(20r^4)}$. ☐

4.2 Towards better dependence on degeneracy: forbidding specific edges

The induced extremal number upper bound obtained in Theorem 2 is of the form $n^{2-1/\mathsf{poly}(r)}$, whereas for standard extremal numbers we know a bound of the form $n^{2-1/\Theta(r)}$. In the corresponding proof for standard extremal numbers, it suffices to apply the dependent random choice of Lemma 3 with $t = \Theta(1)$, guaranteeing that every r-tuple has a common neighbourhood of size larger than some constant — however, we took t much larger in order to guarantee common neighbourhoods of close to linear size.

The reason this was necessary was because our proof was counting "out-of-order". We described a process of choosing an embedding that, when followed in degeneracy order, had exactly $n^{1-1/1.51r}$ many choices at each step. However, in order to bound the number of slippery embeddings, we first had to fix the embeddings of the tuple that got electrocuted, and only then could count the number of choices for the vertex that electrocuted them. If the electrocuter appeared before the electrocutees in the degeneracy order, this meant that we couldn't just embed in order, but instead had to fix the images of the electrocutees first, allowing them n many possibilities each as opposed to $n^{1-1/1.51r}$.

One might be interested whether this technical issue can be overcome to show an upper bound of the form $n^{2-1/\Theta(r)}$. In this section, we make partial progress towards that goal, finding copies of the pattern subgraph which, while not necessarily induced, avoid particular subsets of the pattern graph's non-edges. Our first such result recovers bounds of the form $n^{2-1/\Theta(r)}$ when only a constant number of H's non-edges must be preserved.

Definition 3. For a graph H, and a subset $F \subseteq (V(H) \times V(H) \setminus E(H))$ of "forbidden" edges, let $H \setminus (F\text{-ind})$ denote the family of graphs H' on $V(H)$ such that

- $(u, v) \in E(H) \implies (u, v) \in E(H')$, and
- $(u, v) \in F \implies (u, v) \notin E(H')$.

Proposition 3. For all H of degeneracy r, and $F \subseteq (V(H) \times V(H) \setminus E(H))$ with $|F| = f$, we have $\mathrm{ex}(n, \{K_{s,s}, H \setminus (F\text{-ind})\}) \leq O(n^{2-1/(12f+6r)})$.

Proof. Let $V(F)$ be the set of vertices with an endpoint in F, noting that $|V(F)| \leq 2f$. Consider an n-vertex $K_{s,s}$-free graph G on $n^{2-1/(12f+6r)}$ many vertices. Applying Lemma 3 with $t = 1.81$, we find vertex subsets $U_1, U_2 \subseteq V(G)$ such that any $(r+2f)$-tuple of vertices in one subset has at least $n^{.001}$ many common neighbours in the other. The subgraph of H induced by $V(F)$ has at most $2f$ many vertices, and thus maximum degree at most $2f$ — so, by the techniques of [Hun+24] we can find an induced copy of that subgraph. Fix that subgraph as the embedding of $V(F)$, order the remaining vertices of H in degeneracy order, and them embed one-at-a-time. Each vertex to be embedded is neighbours with at most $r + 2f$ already-embedded vertices, so there are at least $n^{.001}$ many candidates. As long as n is large enough that $n^{.001} > |V(H)|$, this ensures that there is always an option that has not already been used, and so a copy of $H \setminus (F\text{-ind})$ can be found. □

We also observe that it is possible to forbid all edges between vertices that are close to each other in degeneracy order.

Theorem 3. Let H be an r-degenerate bipartite graph, and $v_1, \ldots, v_{|V(H)|}$ be an ordering of the vertices of H such that, for all indices i, there are at most r many indices $j < i$ with $(v_i, v_j) \in E(H)$. Then, for any $q \in \mathbb{N}$, $\mathrm{ex}(n, \{H \setminus (F\text{-ind})\}) \leq O(n^{2-1/(2000q^2r)})$, where $F = \{(v_i, v_j) \colon (v_i, v_j) \notin E(H) \text{ and } |i - j| \leq q\}$.

To prove Theorem 3, we will need to re-do the analysis of dependent random choice. We are interested in obtaining something complementary to Lemma 4: we want our sets to be such that, for any r-tuple of vertices in one set, although few vertices are neighbours with a constant fraction of the tuple's common neighbourhood, *every* vertex in that set is neighbours with at least a *somewhat-large* fraction of the tuple's common neighbourhood.

Lemma 5. For any $r, t \geq 2$, and any n-vertex graph G with at least $n^{2-1/(9t^2r)}$ many edges, there exist subsets $U_1, U_2 \subseteq V(G)$ such that both of the following conditions hold.

- Any r-tuple of vertices $(u_1, \ldots, u_r) \in U_i^r$ has a large common neighbourhood — that is, $|N(v) \cap U_{3-i}| \geq n^{1/10}$ for all $v \in U_i$.
- For any r-tuple of vertices $(u_1, \ldots, u_r) \in U_i^r$, and any other vertex $v \in U_i$, a sizeable fraction of the common neighbourhood of $(u_1, \ldots, u_r)$ is neighbours with v — that is, $\frac{|\bigcap_i N(u_i) \cap U_{3-i} \cap N(v)|}{|\bigcap_i N(u_i) \cap U_{3-i}|} \geq n^{-1/t}$.

Proof. The proof is essentially the same as that of Lemma 3, but we include it in full for completeness. First, partition $V(G)$ into two parts L and R such that at least half of the edges cross the partition. Then, choose $q_L = 3t^2r$ many vertices $\ell_1, \ldots, \ell_{q_L} \in L$ with replacement, and consider their common neighbourhood $A = \bigcap_i N(\ell_i) \bigcap R$. Let $X = |A|$, let Y be the number of tuples $(u_1, \ldots, u_p) \in A^p$ such that $|\bigcap_i N(u_i) \cap L| < n^{1/10}$, and let Z be the number of tuples $(u_1, \ldots, u_p, v) \in A^{p+1}$ such that $\frac{|\bigcap_i N(u_i) \cap L \cap N(v)|}{|\bigcap_i N(u_i) \cap L|} < n^{-1/t}$, where $p = r + tr + 2t$.

$$\mathbb{E}[X] = \sum_v \left(\frac{|N(v) \cap L|}{|L|} \right)^{q_L} \geq n \cdot \left(\frac{n^{1-1/(9t^2r)}}{n} \right)^{q_L} \geq n^{2/3}.$$

$$\mathbb{E}[Y] = \sum_{\substack{(u_1,\dots,u_r) \\ |\bigcap_i N(u_i)\cap L| \le n^{1/2}}} \Pr[u_1,\dots,u_p \in A] \le n^p \cdot \left(\frac{n^{1/10}}{n}\right)^{q_L} < 1.$$

To calculate $\mathbb{E}[Z]$, we note that $\Pr_{\ell_1,\dots,\ell_{q_L}}[u_1,\dots,u_p \in A \text{ and } v \in A] \le \Pr_{\ell_1,\dots,\ell_{q_L}}[v \in A \mid u_1,\dots,u_p \in A] = \left(\frac{|\bigcap_i N(u_i)\cap L\cap N(v)|}{|\bigcap_i N(u_i)\cap L|}\right)^{q_L}$. So, bounding the number of tuples $(u_1,\dots,u_p,v) \in L^{p+1}$ such that $\frac{|\bigcap_i N(u_i)\cap L\cap N(v)|}{|\bigcap_i N(u_i)\cap L|} < n^{-1/t}$ by n^{p+1}, we obtain

$$\mathbb{E}[Z] \le n^{p+1} \cdot \left(n^{-1/t}\right)^{q_L} < 1.$$

Since $\mathbb{E}[X-Y-Z] > n^{2/3}$, there exists some choice of $\ell_1,\dots,\ell_{q_L}$ such that $X-Y-Z > n^{2/3}$. Fix A according to this choice of ℓ, and let U_1 be obtained from A by removing one vertex from each tuple $(u_1,\dots,u_r) \in A^r$ such that $|\bigcap_i N(u_i)\cap L| < n^{1/2}$, and one vertex from each tuple $(u_1,\dots,u_r,v) \in A^{r+1}$ such that $\frac{|\bigcap_i N(u_i)\cap L\cap N(v)|}{|\bigcap_i N(u_i)\cap L|} < n^{-1/(2r)}$. We are guaranteed that $|U_1| \ge n^{2/3}$.

Now, choose $q_R = t(r+2)$ many vertices $r_1,\dots,r_{q_R} \in U_1$, and consider their common neighbourhood $B = \bigcap_i N(r_i)\cap L$. Let Y be the number of tuples $(u_1,\dots,u_r) \in B^r$ such that $|\bigcap_i N(u_i)\cap U_1| < n^{1/10}$, and let Z be the number of tuples $(u_1,\dots,u_r,v) \in B^{r+1}$ such that $\frac{|\bigcap_i N(u_i)\cap U_1\cap N(v)|}{|\bigcap_i N(u_i)\cap U_1|} < n^{-1/t}$. By the same logic as above, we have

$$\mathbb{E}[Y] \le n^r \cdot \left(\frac{n^{1/10}}{n^{2/3}}\right)^{q_R} < 1/2,$$

$$\mathbb{E}[Z] \le n^{r+1} \cdot \left(n^{-1/t}\right)^{q_R} < 1/2.$$

So, by union bound, there exists a choice of the $r_1,\dots,r_{q_R}$ such that $Y = Z = 0$. Set U_2 to be the corresponding value of B.

We claim U_1, U_2 satisfy the conditions of the statement. The fact that $Y = Z = 0$ immediately implies the conditions for vertices in U_2. Then, since $r + q_R = p$, for any $(u_1,\dots,u_r) \in U_1$ we have

$$\left|\bigcap_i u_i \cap U_2\right| = \left|\bigcap_i u_i \cap \left(\bigcap_i N(r_i)\cap L\right)\right| \ge n^{1/10},$$

and for any $(u_1,\dots,u_r,v) \in U_1^{r+1}$, we have

$$\frac{|\bigcap_i N(u_i)\cap U_2\cap N(v)|}{|\bigcap_i N(u_i)\cap U_2|} = \frac{|\bigcap_i N(u_i)\cap(\bigcap_i N(r_i)\cap L)\cap N(v)|}{|\bigcap_i N(u_i)\cap(\bigcap_i N(r_i)\cap L)\,|} \ge n^{-1/t}.$$

□

We can now run a proof strategy similar to that of Theorem 2, since Lemma 5 will let us show that the number of available embedding options doesn't grow too much when we embed *only a little bit* out-of-order.

Proof of Theorem 3. Let G be an n-vertex $K_{s,s}$-free graph with at least $n^{2-1/(2000q^2r)}$ many edges. We can apply Lemma 5 with $t = 11q$ to find two parts $U_1, U_2 \subseteq G$ such that any r-tuple in one part has at least $n^{1/10}$ many common neighbours in the other part, and any additional vertex is neighbours with at least an $n^{-1/11t}$ fraction of that common neighbourhood.

As in the proof of Theorem 2, we now define a distribution on homomorphisms $H \to G$, and show that a random homomorphism from this distribution is likely to correspond to a copy of H as a subgraph without any of the forbidden edges. Once again, we will do so by embedding vertices in degeneracy order, choosing an embedding at each step uniformly from a prefix of the list of available candidates. However, in this case instead of defining all of these prefixes in terms of fixed, arbitrary orderings of U_1 and U_2, it will be useful to choose a new, random ordering for each step of the embedding. For any tuple of numbers $w = (w_1, \dots, w_{|V(H)|}) \in [n^{1/10}]^{|V(H)|}$, and any tuple of permutations $\sigma = \big((\sigma_1^{(1)}, \sigma_1^{(2)}), \dots, (\sigma_{|V(H)|}^{(1)}, \sigma_{|V(H)|}^{(2)})\big) \in \big(S_{U_1} \times S_{U_2}\big)^{|V(H)|}$, we define $\pi_w^{(\sigma)} : V(H) \to V(G)$ such that, if v_i belongs to the left (resp. right) part of H, then $\pi_w^{(\sigma)}(v_i)$ is the w_ith vertex of $\bigcap_{j \le i\colon (v_i, v_j) \in E(H)} N(v_j) \cap U_1$ (resp. U_2) to appear in the ordering $\sigma_i^{(1)}$ (resp. $\sigma_i^{(2)}$).

We can make a slightly simpler definition of electrocution here, saying that $v \in U_i$ ***electrocutes*** $u_1, \dots, u_r \in U_i$ if $|\bigcap_i N(u_i) \cap U_j \cap N(v)| \ge \frac{1}{100|V(H)|^2} |\bigcap_i N(u_i) \cap U_j|$. We'll also define a version of slipperiness that requires indices to be close to each other: say that a tuple of vertices $Y = (y_1, \dots, y_{|Y|}) \in G^{|Y|}$ is ***slippery*** if there exist $(j_1, \dots, j_r) \in [|Y|]^r$ and $j^* \in Y \setminus \{j_1, \dots, j_r\}$ such that y_{j^*} electrocutes $(y_{j_1}, \dots, y_{j_r})$, and also $\max_i(j_i) - j^* \le q$.

Claim 6. If n is sufficiently large, w is chosen uniformly from $[n^{1/10}]^{|V(H)|}$, and σ is chosen uniformly from $\big(S_{U_1} \times S_{U_2}\big)^{|V(H)|}$, the image $\mathrm{Im}_{\pi_w^{(\sigma)}(V(H))}$ is slippery with probability at most $\frac{1}{100|V(H)|^2}$.

Proof. By union bound, it suffices to show for every particular $j_1 \le \cdots \le j_r$ and every j^* with $j_r - j^* \le q$ that the probability of $\pi_w^{(\sigma)}(v_{j^*})$ electrocuting $\left(\pi_w^{(\sigma)}(v_{j_1}), \dots, \pi_w^{(\sigma)}(v_{j_r})\right)$ is at most $\frac{1}{100|V(H)|^{r+3}}$. Fix some such $j_1, \dots, j_r, j^*$, and also fix any values for $w_1, \dots, w_{j^*-1}$ and $\sigma_1, \dots, \sigma_{j^*-1}$ — we claim that the electrocution probability is low conditioned on any such choices. Note that the probability of electrocution now depends only on the values of $w_{j^*}, \dots, w_{j_r}$ and $\sigma_{j^*}, \dots, \sigma_{j_r}$, since indices later than j_r have no effect on the embeddings of earlier vertices.

In order to count the number of embeddings that are slippery at these indices, we will divide into two types, whose counts we will bound separately. For any i such that $j^* < i \le j_r$, letting $b \in \{1, 2\}$ be the part of H to which v_i belongs, we let $A_i = \bigcap_{\ell < i\colon (v_\ell, v_i) \in E(H)) \text{ and } \ell \ne j^*} N(\pi_w^{(\sigma)}(v_\ell)) \cap U_b$ be the set of candidates for $\pi_w^{(\sigma)}(v_i)$ at the time of embedding when one ignores the potential requirement to be neighbours with $\pi_w^{(\sigma)}(v_{j^*})$.

Case 1: For all i such that $j^* < i \le j_r$, at least $n^{1/10}$ many vertices of $N(\pi_w^{(\sigma)}(v_{j^*}))$ appear among the first $100 n^{1/10 + 1/11q}$ many vertices of A_i in ordering $\sigma_i^{(b)}$.

In order to bound the probability of electrocution in this case, we can condition on any arbitrary value of σ. Now, to bound the number possible choices of $w_{j^*}, \dots, w_{j_r}$, we can do the following:

i) For each i such that $j^* < i \le j_r$ in order, choose a value for $\pi_w^{(\sigma)}$ among the first $100 n^{1/10+1/11q}$ many vertices of A_i in ordering $\sigma_i^{(b)}$.
ii) Choose a value for $\pi_w^{(\sigma)}(v_{j^*})$ among the vertices that electrocute $\pi_w^{(\sigma)}(v_{j_1}), \dots, \pi_w^{(\sigma)}(v_{j_r})$.

Note that, since σ is fixed, fixing the embeddings $\pi_w^{(\sigma)}(v_{j^*}), \dots, \pi_w^{(\sigma)}(v_{j_r})$ will uniquely determine $w_{j^*}, \dots, w_{j_r}$. In any valid embedding, v_i is chosen either from among the first $n^{1/10}$ many vertices of A (if $(v_{j^*}, v_i) \notin E(H)$), or among the first $n^{1/10}$ many vertices of $A \cap N(\pi_w^{(\sigma)}(v_{j^*})$ (if $(v_{j^*}, v_i) \in E(H)$). So, the number of ways to perform the above process is indeed an upper bound on the number of choices of $w_{j^*}, \dots, w_{j_r}$ that lead to electrocution in this case. Since, by Lemma 4, there are at most $(200s|V(H)|^2)^s$ many vertices that electrocute any fixed tuple $\pi_w^{(\sigma)}(v_{j_1}), \dots, \pi_w^{(\sigma)}(v_{j_r})$, the number of ways to perform the above process is at most $\left(100 n^{1/10+1/11q}\right)^{j_r - j^*} \cdot (200s|V(H)|^2)^s \le n^{(j_r - j^*)/10 + (j_r - j^*)/11q} \cdot (200s|V(H)|^2)^{2s} \le \left(n^{1/10}\right)^{j_r + 1 - j^*} \cdot \frac{1}{200|V(H)|^{r+3}}$ for sufficiently large n. So the probability of both belonging to case 1 and having $\pi_w^{(\sigma)}(v_{j^*})$ electrocute $\pi_w^{(\sigma)}(v_{j_1}), \dots, \pi_w^{(\sigma)}(v_{j_r})$ is at most $\frac{1}{200|V(H)|^{r+3}}$.

Case 2: For some i^* with $j^* < i^* \le j_r$, fewer than $n^{1/10}$ many vertices of $N(\pi_w^{(\sigma)}(v_{j^*}))$ appear among the first $100n^{1/10+1/11q}$ many vertices of A_{i^*} in ordering $\sigma_i^{(b)}$.

We claim that, electrocution aside, case 2 is very unlikely. Fix any values for $w_{j^*}, \dots, w_{i^*-1}$ and $\sigma_{j^*}, \dots, \sigma_{i^*-1}$. By our dependent random choice, we've guaranteed that $|A_{i^*} \cap N(\pi_w^{(\sigma)}(v_{j^*}))| \ge n^{-1/11q} \cdot |A_{i^*}|$. So, in expectation over σ_{i^*} there will be at least $100n^{1/10}$ many vertices of $N(\pi_w^{(\sigma)}(v_{j^*}))$ among the first $100n^{1/10+1/11q}$ many vertices of A_{i^*}. By a Chernoff bound, the probability of lying substantially below this expectation is extremely small. That is, the probability of having fewer than $n^{1/10}$ many vertices of $N(\pi_w^{(\sigma)}(v_{j^*}))$ among the first $100n^{1/10+1/11q}$ many vertices of A_{i^*} is at most the probability of having fewer than $n^{1/10}$ many successes in $100n^{1/10+1/11q}$ flips of an $n^{-1/11q}$-biased coin, which occurs with probability $o(2^{n^{-0.0001}})$. So, union bounding over all possible i^*, for n sufficiently large the probability of lying in case 2 is at most $\frac{1}{200|V(H)|^{r+3}}$.

Since cases 1 and 2 are exhaustive, we have upper bounded the probability of $\pi_w^{(\sigma)}(v_{j^*})$ electrocuting $\left(\pi_w^{(\sigma)}(v_{j_1}), \dots, \pi_w^{(\sigma)}(v_{j_r})\right)$ by $\frac{1}{200|V(H)|^{r+3}} + \frac{1}{200|V(H)|^{r+3}} = \frac{1}{100|V(H)|^{r+3}}$. □

Once again, we can now say that forbidden edges are unlikely unless the embedding is slippery.

Claim 7. Let $v_i, v_j \in V(H)$ be any pair of vertices such that $(v_i, v_j) \notin E(H)$, and $|i - j| \le q$. If n is sufficiently large, w is chosen uniformly from $[n^{1/10}]^{|V(H)|}$, and σ is chosen uniformly from $\left(S_{U_1} \times S_{U_2}\right)^{|V(H)|}$, then $\Pr_{w,\sigma}[(\pi_w^{(\sigma)}(v_i), \pi_w^{(\sigma)}(v_j)) \in E(G)] \le \frac{1}{50|V(H)|^2}$.

Proof. Assume without loss of generality $i < j$. Fix random values for $w_1, \dots, w_{j-1}$ and $\sigma_1, \dots, \sigma_{j-1}$. Now, choosing uniform random values for σ_j and w_j will cause $\pi_w^{(\sigma)}(v_j)$ to be chosen as a uniform random vertex of $\bigcap_{\ell<j:\,(v_\ell,v_j)\in E(H)} N(v_\ell) \cap U_b$, where $b \in \{1, 2\}$ is the index of the part of H to which v_j belongs. If the overall embedding is not slippery, at most a $\frac{1}{100|V(H)|^2}$ fraction of this set belongs to $N(\pi_w^{(\sigma)}(v_i))$. So, by Claim 6 we have $\Pr_{w,\sigma}[(\pi_w^{(\sigma)}(v_i), \pi_w^{(\sigma)}(v_j)) \in E(G)] \le \frac{1}{100|V(H)|^2} + \frac{1}{100|V(H)|^2} = \frac{1}{50|V(H)|^2}$. □

Once again, the probability of $\pi_w^{(\sigma)}$ failing to be injective goes to 0 in n, so union bounding this along with the probability of any forbidden edge existing gives overall probability strictly less than 1 for large n. Thus, there exists some copy of H in G as a subgraph that avoids inducing any edge of F. □

5 Possible connected counterexamples

In Section 4, we've shown an upper bound on $\mathrm{ex}(n, \{K_{s,s}, H\text{-ind}\})$ in terms of $\mathrm{ex}(n, H)$, which would be interesting to improve by strengthening the quantitative bounds of Theorem 2. However, even the best possible control by degeneracy one could hope for would not suffice to demonstrate Conjecture 2, because it's known the degeneracy does not completely determine standard extremal numbers. One is left with the question: is Conjecture 2 likely to be true? In this section, for the benefit of the unbelievers, we briefly discuss a potential source of counterexamples.

The simplest setting to consider would be where $s = 2$. Recall that $\mathrm{ex}(n, K_{2,2}) = \Theta(n^{3/2})$, where the lower bound is attained by the incidence graph of all points and lines over a finite projective plane $\mathrm{PG}(2, q)$ [Bro66]. These projective plane incidence graphs are highly structured; in addition to avoiding $K_{2,2}$, they may avoid many other interesting structures. A possible approach to disproving Conjecture 2 would be to find some subgraph with extremal number $o(n^{3/2})$ which is nonetheless avoided in *induced* form by some such family of incidence graphs.

Indeed, one can find examples of pattern graphs H such that the point-line incidence graph of $\mathrm{PG}(2, q)$ must always contains many copies of H, but where none of the copies are induced. Perhaps the simplest example would be the the Heawood graph with one edge deleted. The ***Heawood graph***, which we will denote Hea, is the incidence graph of the ***Fano plane*** (i.e. all points and lines over $\mathrm{PG}(2, 2)$) — let $\mathrm{Hea}^- = \mathrm{Hea} \setminus e$

be the Heawood graph with a single edge deleted (note that Hea is edge-transitive, so we need not specify which edge is removed).

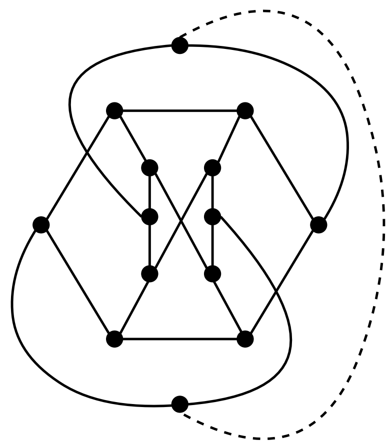

Figure 1: Hea^-, the incidence graph of the Fano plane with a single edge deleted. Deleted edge shown dashed.

Proposition 4. $\text{ex}(n, \{K_{2,2}, \text{Hea}^-\text{-ind}\}) = \Theta(n^{3/2})$.

Proof. A ***complete quadrangle*** is a set of 4 points, no 3 of which are colinear, and the 6 lines between each pair — the ***diagonals*** of a complete quadrangle are the 3 additional intersection points of those lines. For any finite field $\mathbb{F}_q$, we know the diagonals of a complete quadrangle in $\text{PG}(2, q)$ will be colinear if and only if q is a power of 2 [HP82].

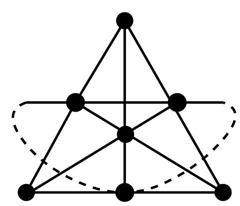

Figure 2: A complete quadrangle, with a line between 2 of the 3 diagonal points. The dashed extension of that line indicates that it passes through the third diagonal if and only if the underlying field has characteristic 2 (in which case the configuration is isomorphic to the Fano plane).

In particular, this means that Hea^- cannot appear as an induced subgraph of the point-line incidence graph of $\text{PG}(2, 2^a)$ for any $a \in \mathbb{N}$: given 7 points and 6 lines corresponding to a complete quadrangle, a line incident to two of the diagonals must also be incident to the third, thus inducing Hea. So incidence graphs of $\text{PG}(2, 2^a)$ give a family of n-vertex graphs with $\Theta(n^{3/2})$ edges avoiding both $K_{2,2}$ as a subgraph and Hea^- as an induced subgraph. □

If we knew that $\text{ex}(n, \text{Hea}^-) < o(n^{3/2})$, this would therefore be a counterexample to Conjecture 2. It seems perhaps plausible that Hea^- could have a small extremal number: it's a 2-degenerate graph of girth 6, and appears as a subgraph (although not necessarily induced) of every complete point-line incidence graph over a finite projective plane. The strongest lower bound we know on its extremal number is $\text{ex}(n, \text{Hea}^-) \geq \Omega(n^{7/5})$, obtained by considering a random host graph. It would be quite interesting to determine tight bounds on the extremal number of Hea^-, and thus determine whether or not it represents a counterexample to Conjecture 2, however this may be a difficult task: there are remarkably few graphs for which the true Turán exponent is known, with the techniques involved typically quite specialized to the

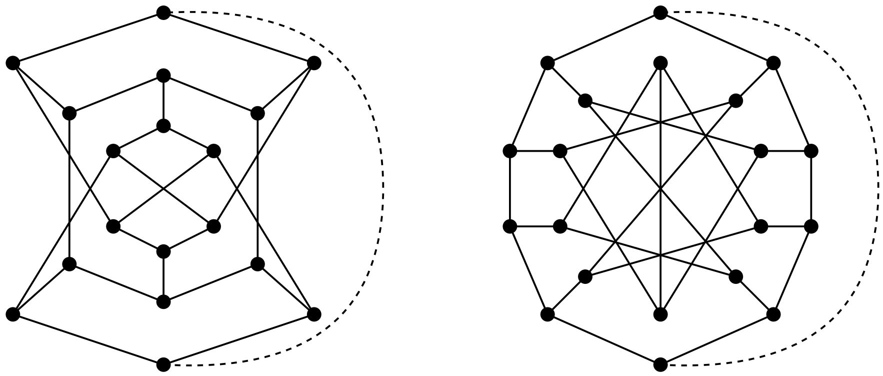

Figure 3: Two graphs H where the bound $\mathrm{ex}(n, \{K_{2,2}, H\text{-ind}\}) = \Theta(n^{3/2})$ can be obtained from incidence geometry theorems: the Pappus graph (left), and the Desargues graph (right), each with a single deleted edge (denoted by a dashed line). Pappus's theorem and Desargues's theorem, respectively, ensure that neither can appear as an induced subgraph of $\mathrm{PG}(2, q)$ for any q, since the deleted edge will always be present.

particular graph in question.

The Heawood graph is, of course, far from the only case of such a structure: there are many more complicated theorems demonstrating that, in some particular incidence configuration, 3 points must be colinear [Ric95; FP23]. Such theorems will allow us to find subgraphs which appear in all projective plane incidence graphs, but can be avoided in induced form (see Figure 3 for two additional examples). Conjecture 2 would hold that all such graphs have extremal number $\Omega(n^{3/2})$; evidence for or against that prediction could give intuition as to whether Conjecture 2 is likely to be true.

6 Acknowledgements

The results in this paper were produced as a part of the Duluth REU program. We thank the organizers, advisors, visitors and participants of the program for offering mathematical advice and providing a stimulating research environment — we are especially grateful to Maya Sankar and Carl Schildkraut for a number of helpful discussions. We also thank Joe Gallian and Colin Defant for running the program, and Jane Street Capital, Ray Sidney, Eric Wepsic and the NSF for providing funding.

References

[AKS03] Noga Alon, Michael Krivelevich, and Benny Sudakov. "Turán numbers of bipartite graphs and related Ramsey-type questions". In: *Combinatorics, Probability and Computing* 12.5-6 (2003), pp. 477–494.

[AZ24] Maria Axenovich and Jakob Zimmermann. "Induced Turán problem in bipartite graphs". In: *arXiv preprint arXiv:2401.11296* (2024).

[Bro66] William G Brown. "On graphs that do not contain a Thomsen graph". In: *Canadian Mathematical Bulletin* 9.3 (1966), pp. 281–285.

[Chu14] Maria Chudnovsky. "The Erdös–Hajnal Conjecture—A Survey". In: *Journal of Graph Theory* 75.2 (2014), pp. 178–190. DOI: https://doi.org/10.1002/jgt.21730. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/jgt.21730. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/jgt.21730.

[EGM19] Beka Ergemlidze, Ervin Győri, and Abhishek Methuku. "Turán number of an induced complete bipartite graph plus an odd cycle". In: *Combinatorics, Probability and Computing* 28.2 (2019), pp. 241–252.

[EH89] P. Erdös and A. Hajnal. "Ramsey-type theorems". In: *Discrete Applied Mathematics* 25.1 (1989), pp. 37–52. ISSN: 0166-218X. DOI: https://doi.org/10.1016/0166-218X(89)90045-0. URL: https://www.sciencedirect.com/science/article/pii/0166218X89900450.

[Erd97] Paul Erdős. "Some of my favorite problems and results". In: *The Mathematics of Paul Erdős I*. Springer, 1997, pp. 47–67.

[ES46] Paul Erdös and Arthur Stone. "On the structure of linear graphs". In: *Bulletin of the American Mathematical Society* 52.12 (1946), pp. 1087–1091.

[ES66] Paul Erdos and Miklós Simonovits. "A limit theorem in graph theory". In: *Studia Sci. Math. Hungar* 1.51-57 (1966), p. 51.

[FNP24] Jacob Fox, Rajko Nenadov, and Huy Tuan Pham. "The largest subgraph without a forbidden induced subgraph". In: *arXiv preprint arXiv:2405.05902* (2024).

[FP23] Sergey Fomin and Pavlo Pylyavskyy. "Incidences and tilings". In: *arXiv preprint arXiv:2305.07728* (2023).

[FS09] Jacob Fox and Benny Sudakov. "Density theorems for bipartite graphs and related Ramsey-type results". In: *Combinatorica* 29.2 (2009), pp. 153–196.

[FS11] Jacob Fox and Benny Sudakov. "Dependent random choice". In: *Random Structures & Algorithms* 38.1-2 (2011), pp. 68–99.

[FS13] Zoltán Füredi and Miklós Simonovits. "The history of degenerate (bipartite) extremal graph problems". In: *Erdős centennial*. Springer, 2013, pp. 169–264.

[HP82] Daniel R. Hughes and Frederick Charles Piper. *Projective planes*. Springer-Verlag, 1982.

[Hun+24] Zach Hunter et al. *Kővári-Sós-Turán theorem for hereditary families*. 2024. arXiv: 2401.10853.

[Hyl58] C Hyltén-Cavallius. "On a combinatorical problem". In: *Colloquium Mathematicae*. Vol. 6. 1. 1958, pp. 61–65.

[Ill21a] Freddie Illingworth. "A note on induced Turán numbers". In: *arXiv preprint arXiv:2105.12503* (2021).

[Ill21b] Freddie Illingworth. "Graphs with no Induced $K_{2,t}$". In: *Electronic Journal of Combinatorics* 28.1 (2021).

[KO04] Daniela Kuhn and Deryk Osthus. "Induced subdivisions in $K_{s,s}$-free graphs of large average degree". In: *Combinatorica* 24.2 (2004), pp. 287–304.

[KST54] Tamás Kővári, Vera Sós, and Pál Turán. "On a problem of K. Zarankiewicz". eng. In: *Colloquium Mathematicae* 3.1 (1954), pp. 50–57. URL: http://eudml.org/doc/210011.

[Loh+18] Po-Shen Loh et al. "Induced Turán Numbers". In: *Combinatorics, Probability and Computing* 27.2 (2018), pp. 274–288. DOI: 10.1017/S0963548317000542.

[Ric95] Jürgen Richter-Gebert. "Mechanical theorem proving in projective geometry". In: *Annals of Mathematics and Artificial Intelligence* 13 (1995), pp. 139–172.

[SSS23] Alex Scott, Paul Seymour, and Sophie Spirkl. "Polynomial bounds for chromatic number. I. Excluding a biclique and an induced tree". In: *Journal of Graph Theory* 102.3 (2023), pp. 458–471.

[ST20] Benny Sudakov and István Tomon. "Turán number of bipartite graphs with no $K_{t,t}$". In: *Proceedings of the American Mathematical Society* 148.7 (2020), pp. 2811–2818.