

INDUCED SUBGRAPH DENSITY. VI. BOUNDED VC-DIMENSION

TUNG NGUYEN, ALEX SCOTT, AND PAUL SEYMOUR

ABSTRACT. We confirm a conjecture of Fox, Pach, and Suk, that for every $d > 0$, there exists $c > 0$ such that every n -vertex graph of VC-dimension at most d has a clique or stable set of size at least n^c . This implies that, in the language of model theory, every graph definable in NIP structures has a clique or anti-clique of polynomial size, settling a conjecture of Chernikov, Starchenko, and Thomas.

Our result also implies that every two-colourable tournament satisfies the tournament version of the Erdős-Hajnal conjecture, which completes the verification of the conjecture for six-vertex tournaments. The result extends to uniform hypergraphs of bounded VC-dimension as well.

The proof method uses the ultra-strong regularity lemma for graphs of bounded VC-dimension proved by Lovász and Szegedy and the method of iterative sparsification introduced by the authors in an earlier paper.

1. INTRODUCTION

All (hyper)graphs in this paper are finite and simple. A graph H is an *induced subgraph* of a graph G if H can be obtained from G by removing vertices. A class \mathcal{C} of graphs is *hereditary* if it is closed under taking induced subgraphs and under isomorphism; and a hereditary class \mathcal{C} is *proper* if it is not the class of all graphs. We say that \mathcal{C} has the *Erdős-Hajnal property* if there exists $c > 0$ such that every graph $G \in \mathcal{C}$ has a clique or stable set of size at least $|G|^c$, where $|G|$ denotes the number of vertices of G . A conjecture of Erdős and Hajnal [18, 19] (see [14, 29] for surveys and [9, 16, 39, 41] for some recent partial results) asserts that:

Conjecture 1.1. *Every proper hereditary class of graphs has the Erdős-Hajnal property.*

For a set \mathcal{F} of subsets of a set V , a subset S of V is *shattered* by \mathcal{F} if for every $A \subseteq S$ there exists $B \in \mathcal{F}$ with $B \cap S = A$. The *VC-dimension* of \mathcal{F} (introduced by Vapnik and Chervonenkis in [50]) is the largest cardinality of a subset of V that is shattered by \mathcal{F} . Since its introduction in 1971, the notion of VC-dimension has proved to be relevant in a number of areas of pure and applied mathematics. The *VC-dimension* of a graph G is the VC-dimension of the set $\{N_G(v) : v \in V(G)\}$ of subsets of $V(G)$, where $N_G(v)$ denotes the set of all neighbours of v (not including v itself). It is not hard to see that for every $d \geq 1$, the class of graphs of VC-dimension at most d is a proper hereditary class. The aim of this paper is to confirm a conjecture of Fox, Pach, and Suk [22] that every hereditary class of graphs of bounded VC-dimension has the Erdős-Hajnal property; in their paper, they came close to settling this by proving a bound of $2^{(\log n)^{1-o(1)}}$ where the constant depending on the VC-dimension is hidden in the $o(1)$ term. (In this paper \log denotes the binary logarithm.) Our result is:

Theorem 1.2. *For every $d \geq 1$, there exists $c_d > 0$ such that every graph G of VC-dimension at most d contains a clique or stable set of size at least $|G|^{c_d}$.*

(A careful inspection of our proof yields $c_d \geq c_{d-1}^K$ for every $d \geq 2$ where K is a universal constant; and so one can take $c_d \geq 2^{-2^{O(d)}}$. We omit the details.)

As a result, this paper can be viewed as part of a growing body of results [22–24, 30] that, when VC-dimension is bounded, completely settle or significantly improve bounds for well-known open problems in extremal combinatorics.

The story behind Theorem 1.2, perhaps, began in geometric graph theory with the result of Larman, Matoušek, Pach, and Töröcsik [31] that the class of intersection graphs of line segments in general position in the plane has the Erdős-Hajnal property (in fact, they proved more, for intersection graphs of convex compact sets in the plane). Alon, Pach, Pinchasi, Radoičić, and Sharir [2] generalized this result to classes of semi-algebraic graphs of bounded description complexity. Fox, Pach, and Tóth [25] provided another extension of [31] by verifying the Erdős-Hajnal property of the class of string graphs

Date: November 13, 2023; revised July 18, 2025.

The first and the third authors were partially supported by AFOSR grant FA9550-22-1-0234 and NSF grant DMS-2154169. The second author was supported by EPSRC grant EP/X013642/1.

where every two curves cross a bounded number of times (and a recent result by Tomon [49] shows that the condition “every two curves cross a bounded number of times” can be dropped). In yet another direction, Sudakov and Tomon [48] recently proved Conjecture 1.1 for the classes of algebraic graphs of bounded description complexity, which is an analogue of the result of [2]. All of these hereditary classes (except intersection graphs of convex compact sets, and more generally, string graphs) turn out to have bounded VC-dimension. Indeed,

- for the classes of semi-algebraic graphs of bounded description complexity, this is true by the classical Milnor–Thom theorem in real algebraic geometry (see [36]);
- for the classes of algebraic graphs of bounded description complexity, this is a consequence of a theorem of Rónyai, Babai, and Ganapathy [45] on the number of zero-patterns of polynomials; and
- for the classes of string graphs where every two curves intersect at a bounded number of points, this follows from a result of Pach and Tóth [43, Lemma 4.2] together with the standard fact that a hereditary class has bounded VC-dimension if and only if it does not contain a bipartite graph, the complement of a bipartite graph, and a split graph (see Lemma 5.1).

In a model-theoretic setting, Theorem 1.2 states that every class of graphs definable in NIP (non-independence property) structures has the Erdős–Hajnal property (see [46] for a general reference on NIP theories), which was formally stated as a conjecture by Chernikov, Starchenko, and Thomas [12]. Two notable special cases of NIP graphs include distal graphs and stable graphs.¹ Malliaris and Shelah [33, 34] implicitly proved Conjecture 1.1 for stable graphs (which contains the result of [48]) by developing regularity lemmas for these graphs (see [10] for a short proof using pseudo-finite model theory). In the case of distal graphs, Basu [4] proved the Erdős–Hajnal property for graphs definable by o-minimal structures (which in fact extends [2]), before Chernikov and Starchenko [11] made use of the theory of Keisler measures in NIP to formulate regularity lemmas for distal graphs and settle the general problem in this direction (see also [47] for a short and pure model-theoretic proof). Recently, Fu [27] also combined tools from model theory and a result from [16] to prove the Erdős–Hajnal property of the class of graphs of VC-dimension at most two.

In the combinatorial setting, Theorem 1.2 turns out to be equivalent to the fact that every tournament not containing a fixed two-colourable subtournament has a transitive vertex subset of polynomial size; we will discuss this in more detail in Section 6. Theorem 1.2 also extends to the setting of uniform hypergraphs, allowing us to give tighter asymptotics on Ramsey numbers for k -uniform hypergraphs with bounded VC-dimension; we will discuss this in Section 4.

As in previous papers of the series, we need to work in a more general context, and will prove a stronger result: our main result says that graphs of bounded VC-dimension satisfy the polynomial form of a theorem of Rödl. For a graph G , let \overline{G} denote its complement; and for $\varepsilon \in (0, \frac{1}{2})$, we say that G is ε -sparse if G has maximum degree at most $\varepsilon|G|$, and ε -restricted if one of G, \overline{G} is ε -sparse. We recall that Rödl’s theorem [44] then states the following:

Theorem 1.3. *For every proper hereditary class \mathcal{C} of graphs and every $\varepsilon \in (0, \frac{1}{2})$, there exists $\delta > 0$ such that every graph $G \in \mathcal{C}$ contains an ε -restricted induced subgraph on at least $\delta|G|$ vertices.*

The dependence of δ on ε here is an important question. Rödl’s proof used Szemerédi’s regularity lemma, and gave tower-type bounds. A better dependence was obtained by Fox and Sudakov [26], who gave a proof with $\delta = 2^{-c \log^2(1/\varepsilon)}$; this was improved to $\delta = 2^{-c \log^2(1/\varepsilon) / \log \log(1/\varepsilon)}$ in [9]. Fox and Sudakov [26] conjectured that in fact the dependence can be taken to be polynomial, which would imply the Erdős–Hajnal conjecture.

Let us say a hereditary class \mathcal{C} has the *polynomial Rödl property* if there exists $C > 0$ such that for every $\varepsilon \in (0, \frac{1}{2})$, every graph $G \in \mathcal{C}$ contains an ε -restricted induced subgraph on at least $\varepsilon^C|G|$ vertices. The Fox–Sudakov conjecture is then the following strengthening of Conjecture 1.1:

Conjecture 1.4. *Every proper hereditary class of graphs has the polynomial Rödl property.*

Our main result says that every hereditary class of graphs of bounded VC-dimension does indeed satisfy Conjecture 1.4.

Theorem 1.5. *For every $d \geq 1$, the class of graphs of VC-dimension at most d has the polynomial Rödl property.*

¹This is different from the notion of stable vertex subsets in graphs; in graph-theoretic language, it means graphs not containing a fixed half-graph bipartite pattern. See also [35].

Our proof of Theorem 1.5 uses the ultra-strong regularity lemma for graphs of bounded VC-dimension proved by Lovász and Szegedy, and builds on the method of iterative sparsification introduced in earlier papers of the series [39, 40]. The method involves passing through a sequence of induced subgraphs that are successively more restricted (see Section 3), and naturally leads to the extra generality of Theorem 1.5.

We would like to remark that Bucić, Fox, and Pham [8] recently proved that the Erdős-Hajnal property is in fact equivalent to the polynomial Rödl property for any hereditary class of graphs, which implies that Theorems 1.2 and 1.5 are equivalent. Nevertheless, for our proof method means that it is more convenient to prove the result in the stronger polynomial Rödl form; and in the case of graphs of bounded VC-dimension, such an equivalence can also be obtained via an application of the ultra-strong regularity lemma (see Theorem 2.4 and the argument at the end of Section 3).

2. ULTRA-STRONG REGULARITY, BIGRAPHS AND NEARLY PURE BLOCKADES

A natural approach to proving Erdős-Hajnal-type results is to attempt to find induced subgraphs that are highly structured; the challenge is to control the size and structure of these subgraphs. In this section, we work with large subgraphs that are ‘approximately’ nicely structured: up to some noise, they look like a blowup of some much smaller graph. These will be helpful in the next section, where we look for subgraphs where the structure is more tightly controlled.

A *blockade* in a graph G is a sequence $\mathcal{B} = (B_1, \dots, B_k)$ of pairwise disjoint (and possibly empty) subsets of $V(G)$, called *blocks*; the *length* of \mathcal{B} is k and the *width* is $\min_{i \in [k]} |B_i|$. For $\ell, w \geq 0$, we say that \mathcal{B} is an (ℓ, w) -*blockade* if it has length at least ℓ and width at least w .

We will be interested in blockades where the pairs (B_i, B_j) satisfy certain density conditions. Two disjoint subsets A, B of vertices are *anticomplete* in G if G has no edge between A, B ; *complete* in G if A, B are anticomplete in \overline{G} ; and *pure* if A, B are either complete or anticomplete in G . A blockade $\mathcal{B} = (B_1, \dots, B_k)$ is *anticomplete* in G if B_i, B_j are anticomplete in G for all distinct $i, j \in [k]$, *complete* in G if \mathcal{B} is anticomplete in \overline{G} and *pure* if, for all distinct $i, j \in [k]$, B_i, B_j are either complete or anticomplete. Note that for a blockade, being pure is a far weaker condition than being a complete or anticomplete blockade.

In general, complete or anticomplete blockades will be very desirable; but we will first need to work with blockades that are only pure, and that only satisfy approximate versions of the conditions. For $\varepsilon > 0$ and a graph G , two disjoint subsets A, B are *weakly ε -sparse* in G if G has at most $\varepsilon|A||B|$ edges between A, B , and *weakly ε -pure* if they are weakly ε -sparse in G or \overline{G} . We say that B is *ε -sparse to A* in G if every vertex of B is adjacent in G to at most $\varepsilon|A|$ vertices in A ; the pair (A, B) is *ε -pure* if each is ε -sparse to the other in G or \overline{G} . A blockade \mathcal{B} is *ε -pure* or *weakly ε -pure* if every pair B_i, B_j satisfies the corresponding condition.

The first step in the proof of Theorem 3.1 is to find a large ε -pure blockade. For this we use the ultra-strong regularity lemma of Lovász and Szegedy [32, Section 4] which was reproved by Fox, Pach, and Suk [22, Theorem 1.3] via a short proof that gives better bounds (a weaker version for bipartite graphs was also developed by Alon, Fisher, and Newman [1]). The ultra-strong regularity lemma says:

Theorem 2.1. *For every $d \geq 1$, there exists $K \geq 2$ such that the following holds. For every $\varepsilon \in (0, \frac{1}{2})$ and every graph G with VC-dimension at most d , there exist $L \in [\varepsilon^{-1}, \varepsilon^{-K}]$ and an equipartition $V(G) = V_1 \cup \dots \cup V_L$ such that all but at most an ε fraction of the pairs (V_i, V_j) are weakly ε -pure in G .*

Here, an *equipartition* of a set S is a partition $S_1 \cup \dots \cup S_k$ of S such that $|S_i|, |S_j|$ differ by at most one for all $i, j \in [k] = \{1, \dots, k\}$.

It will be helpful to work with bipartite patterns in G . Let us make this more formal. A *bigraph* is a graph H together with a bipartition $(V_1(H), V_2(H))$ of $V(H)$. If H is a bigraph and $v \in V(H)$, we denote by $H \setminus v$ the bigraph H' obtained by deleting v , where for $i = 1, 2$, $V_i(H') = V_i(H) \setminus \{v\}$ if $v \in V_i(H)$, and $V_i(H') = V_i(H)$ otherwise. For a bigraph H , a *bi-induced* copy of H in a graph G is an injective map $\varphi: V(H) \rightarrow V(G)$ such that, for all $u \in V_1(H)$ and $v \in V_2(H)$, $uv \in E(H)$ if and only if $\varphi(u)\varphi(v) \in E(G)$. We recall the following standard fact (see [7, Lemma 3.3]).

Lemma 2.2. *For every bigraph H , there exists $d \geq 1$ such that every graph of VC-dimension at least d contains a bi-induced copy of H . Conversely, for every $d \geq 1$, there is a bigraph H such that every graph containing a bi-induced copy of H has VC-dimension at least d .*

In view of this, Theorem 2.1 can be restated as follows.

Theorem 2.3. *For every bigraph H , there exists $K \geq 2$ such that for every $\varepsilon \in (0, \frac{1}{2})$ and every graph G with no bi-induced copy of H , there exists $L \in [\varepsilon^{-1}, \varepsilon^{-K}]$ for which there is an equipartition $V(G) = V_1 \cup \dots \cup V_L$ such that all but at most an ε fraction of the pairs (V_i, V_j) are weakly ε -pure.*

We use Theorem 2.1 to obtain a blockade that is ε -pure and large (more specifically, we want the length and width to have polynomial dependence on ε).

Theorem 2.4. *For every bigraph H , there exists $b \geq 2$ such that for every $\varepsilon \in (0, \frac{1}{2})$ and every graph G with $|G| \geq \varepsilon^{-b}$ and no bi-induced copy of H , there is an $(\varepsilon^{-1}, \varepsilon^b|G|)$ -blockade (B_1, \dots, B_ℓ) in G such that*

- $|B_1| = \dots = |B_\ell| \leq \varepsilon^2|G|$; and
- for all distinct $i, j \in [\ell]$, B_i, B_j are ε -sparse to each other in G or \overline{G} .

Proof. Let $K \geq 2$ be given by Theorem 2.3; we claim that $b := 5K$ satisfies the theorem. To this end, by Theorem 2.3 with ε^4 in place of ε , if G is a graph with no bi-induced copy of H , then there is an equipartition $V(G) = V_1 \cup \dots \cup V_L$ with $L \in [\varepsilon^{-4}, \varepsilon^{-4K}]$ such that all but at most an ε^4 fraction of the pairs (V_i, V_j) are weakly ε^4 -pure. By Turán's theorem, there exists $J \subseteq [L]$ with $|J| \geq \frac{1}{2}\varepsilon^{-4}$ such that (V_i, V_j) is weakly ε^4 -pure for all distinct $i, j \in J$. Then there exists $I \subseteq J$ with $|I| \geq \frac{1}{2}|J| \geq \frac{1}{4}\varepsilon^{-4} \geq \varepsilon^{-1}$ such that $|V_i| = |V_j|$ for all distinct $i, j \in I$. Let $\ell := \lceil \varepsilon^{-1} \rceil$; it follows that there exists $I \subseteq J$ with $I = [\ell]$ such that $|V_i| = |V_j|$ for all distinct $i, j \in I$. For every $i \in I$, let B_i be the set of vertices v in V_i such that for every $j \in I \setminus \{i\}$,

- v has at most $\frac{1}{2}\varepsilon|V_j|$ neighbours in V_j if (V_i, V_j) is weakly ε^4 -sparse in G ; and
- v has at most $\frac{1}{2}\varepsilon|V_j|$ nonneighbours in V_j if (V_i, V_j) is weakly ε^4 -sparse in \overline{G} .

Then

$$|B_i| \geq |V_i| - (\ell - 1) \cdot 2\varepsilon^3|V_i| \geq |V_i| - \varepsilon^{-1} \cdot 2\varepsilon^3|V_i| = (1 - 2\varepsilon^2)|V_i| \geq |V_i|/2$$

and by removing vertices if necessary we may assume that $|B_i| = \lceil |V_i|/2 \rceil = \lceil m/2 \rceil$ where $m = |V_i|$. It follows that for all distinct $i, j \in I$, B_i, B_j are ε -sparse to each other in G or \overline{G} . Also, since

$$m \geq \lfloor |G|/L \rfloor \geq |G|/(2L) \geq \varepsilon^{4K+1}|G|$$

(as $|G| \geq \varepsilon^{-b} = \varepsilon^{-5K}$), it follows that for each $i \in I$,

$$|B_i| \geq m/2 \geq \varepsilon^{4K+2}|G| \geq \varepsilon^{5K}|G| = \varepsilon^b|G|$$

and

$$|B_i| \leq m \leq \lceil |G|/L \rceil \leq 2|G|/L \leq 2\varepsilon^4|G| \leq \varepsilon^2|G|.$$

This proves Theorem 2.4. ■

3. ITERATIVE SPARSIFICATION

In this section we give the proof of Theorem 1.5. We recall that a graph is ε -restricted if either the graph or its complement is ε -sparse. In view of Lemma 2.2, Theorem 1.5 is equivalent to:

Theorem 3.1. *For every bigraph H , there exists $C > 0$ such that for every $\varepsilon \in (0, \frac{1}{2})$, every graph G with no bi-induced copy of H contains an ε -restricted induced subgraph with at least $\varepsilon^C|G|$ vertices.*

The proof of Theorem 3.1 uses the framework of *iterative sparsification*, which was introduced in earlier papers from this series [39, 40] (and will be used also in [41]). The goal is to find an ε -restricted induced subgraph of size at least $\text{poly}(\varepsilon)|G|$. Rather than doing this in one step, we instead attempt to move through a sequence of induced subgraphs that are successively sparser (in G or its complement): given a y -restricted induced subgraph F , we search for an induced subgraph that is $\text{poly}(y)$ -restricted and is at most a $\text{poly}(y)$ factor smaller. Provided we can start the process, and do not get stuck on the way, the following lemma shows that this gives the required subset.

Lemma 3.2. *Let $c \in (0, 1)$, $a \geq 2$, and $t \geq 1$. Suppose that $x \in (0, c)$ and G is a graph satisfying:*

- *there is a c -restricted induced subgraph of G with at least $c^t|G|$ vertices; and*
- *for every $y \in [x, c]$ and every y -restricted induced subgraph F of G with $|F| \geq y^{2t}|G|$, there is a y^a -restricted induced subgraph of F with at least $y^{at}|F|$ vertices.*

Then G contains an x -restricted induced subgraph with at least $x^{2at}|G|$ vertices.

Proof. By the first condition of the lemma, there exists $y \in [x^a, c]$ minimal such that G has a y -restricted induced subgraph F with $|F| \geq y^{2t}|G|$. If $y \geq x$, then by the second condition of the lemma and since $a \geq 2$, F has a y^a -restricted induced subgraph with at least $y^{at}|F| \geq y^{at+2t}|G| \geq y^{2at}|G|$ vertices; but this contradicts the minimality of y since $x^a \leq y^a < y$. Thus $x^a \leq y < x$; and so F is x -restricted, and $|F| \geq y^{2t}|G| \geq x^{2at}|G|$. This proves Lemma 3.2. \blacksquare

We will find a subgraph satisfying the first bullet by using Rödl's theorem (Theorem 1.3), with a suitable $t = t(c)$. However, finding a subgraph that satisfies the second bullet is more challenging, and we need to allow for an alternative 'good' outcome. We will show in Lemma 3.7 that if we get stuck then we can instead find a large complete or anticomplete blockade (note that this is much stronger than being pure; and the blockades given by Theorem 2.4 are only ε -pure). Let us show that, if we can find sufficiently large blockades, then we can obtain the Erdős-Hajnal property; we will return to the (stronger) polynomial Rödl property at the end of the section.

Lemma 3.3. *Let \mathcal{C} be a hereditary class of graphs. Suppose that there exists $d \geq 2$ such that for every $x \in (0, 2^{-d})$ and every $G \in \mathcal{C}$, either:*

- G has an x -restricted induced subgraph with at least $x^d|G|$ vertices; or
- there is a complete or anticomplete $(k, |G|/k^d)$ -blockade in G , for some $k \in [2, 1/x]$.

Then there exists $a \geq 2$ such that every n -vertex graph in \mathcal{C} has a clique or stable set of size at least $n^{1/a}$.

Proof. A *cograph* is a graph with no induced four-vertex path; and it is well-known that every k -vertex cograph has a clique or stable set of size at least $k^{1/2}$. Thus, it suffices to prove by induction that every $G \in \mathcal{C}$ contains an induced cograph of size at least $|G|^{1/(2d^2)}$. We may assume $|G| > 2^{2d^2}$. Let $x := |G|^{-1/(2d)} \in (0, 2^{-d})$. By the hypothesis, either:

- there exists $S \subseteq V(G)$ with $|S| \geq x^d|G|$ such that $G[S]$ is x -restricted; or
- there is a complete or anticomplete $(k, |G|/k^d)$ -blockade in G , for some $k \in [2, 1/x]$.

If the first bullet holds, then since $|S| \geq x^d|G| \geq x^{-1}$, Turán's theorem gives a clique or stable set in $G[S]$ of size at least $(2x)^{-1} > x^{-1/2} = |G|^{1/(4d)}$. If the second bullet holds, then since $k \geq 2$, by applying induction to each block in the blockade to obtain a cograph and combining them, we see that G contains an induced cograph of size at least

$$k(|G|/k^d)^{1/(2d^2)} = k^{1-1/(2d)}|G|^{1/(2d^2)} > |G|^{1/(2d^2)}.$$

This proves Lemma 3.3. \blacksquare

The key step in making the iterative sparsification strategy work is therefore to show that if the second bullet of Lemma 3.2 does not hold then we can find a sufficiently large complete or anticomplete blockade. We will show this in Lemma 3.6: given a y -restricted graph F with no bi-induced H , we will prove that we can either pass to the desired $\text{poly}(y)$ -restricted subgraph or find a complete or anticomplete blockade whose length and width depend polynomially on y . We will argue by induction on $|H|$, and grow the blockade one block at a time: the first step (Lemma 3.4) is to find a complete or anticomplete pair (A, B) , where A has size $\text{poly}(y)|F|$ and B contains all but a small fraction of the rest of F ; we then (Lemma 3.5) repeat the argument inside B , continuing until we obtain a blockade. Restricted graphs can either be dense or sparse: we will assume for the moment that our restricted graph is sparse, and handle the dense case later by taking complements.

Lemma 3.4. *Let H be a bigraph, and let $v \in V(H)$. Let $b \geq 2$ be given by Theorem 2.4. Assume there exists $a \geq 2$ such that every n -vertex graph with no bi-induced copy of $H \setminus v$ contains a clique or stable set of size at least $n^{1/a}$. Let $y \in (0, 1/|H|)$, and let F be a y -sparse graph with no bi-induced copy of H . Then either:*

- F has a y^{2a} -restricted induced subgraph with at least $y^{3ba^2}|F|$ vertices; or
- there is an anticomplete pair (A, B) in F with $|A| \geq y^{3ba^2}|F|$ and $|B| \geq (1 - 3y)|F|$.

Proof. We have a sparse graph, and want to find an anticomplete pair (A, B) . We will do this by first using Theorem 2.4 (a corollary of the ultra-strong regularity lemma) to find a large, $\text{poly}(y)$ -pure blockade and then looking at how the rest of the graph attaches to it. So let $\varepsilon := y^{3a^2}$, and suppose that the first outcome does not hold; then $|F| > y^{-3ba^2} = \varepsilon^{-b}$. By Theorem 2.4, F has a $(\varepsilon^{-1}, \varepsilon^b|F|)$ -blockade (B_1, \dots, B_ℓ) with $\ell = \lceil \varepsilon^{-1} \rceil$, such that:

- $|B_1| = \dots = |B_\ell| \leq 2\varepsilon^2|F|$; and
- for all distinct $i, j \in [\ell]$, B_i, B_j are ε -sparse to each other in F or \overline{F} .

Let $D := V(F) \setminus (B_1 \cup \dots \cup B_\ell)$ and $m := |B_1|$. For $i \in [\ell]$, a vertex $v \in D$ is *mixed* on B_i if it has a neighbour and a nonneighbour in B_i .

Claim. *Every vertex in D is mixed on at most $y\ell$ of the blocks B_1, \dots, B_ℓ .*

Proof. Suppose there is a vertex $w \in D$ mixed on at least $y\ell$ of the blocks B_1, \dots, B_ℓ , say B_1, \dots, B_r where $r \geq y\ell \geq y\varepsilon^{-1} = y^{1-3a^2}$. Let J be the graph with vertex set $[r]$ where for all distinct $i, j \in [r]$, $ij \notin E(J)$ if and only if B_i, B_j are ε -sparse to each other in F .

We claim that there is no bi-induced copy of $H \setminus v$ in J . Suppose that there is; and we may assume $V(H \setminus v) \subseteq V(J)$. We assume that $v \in V_1(H)$ without loss of generality. For each $u \in V_2(H)$, let w_u be a neighbour of w in B_u if $uv \in E(H)$ and a nonneighbour of w in B_u if $uv \notin E(H)$. For each $z \in V_1(H) \setminus \{v\}$ and $u \in V_2(H)$, since $uz \notin E(H)$ if and only if $uz \notin E(J)$ if and only if B_u, B_z are ε -sparse to each other in F , w_u is adjacent in F to at most $\varepsilon|B_z|$ vertices in B_z if $uz \notin E(H)$ and nonadjacent in G to at most $\varepsilon|B_z|$ vertices in B_z if $uz \in E(H)$. Thus, for each $z \in V_1(H) \setminus \{v\}$, there are at least (note that $\varepsilon \leq y < 1/|H|$)

$$|B_z| - |V_2(H)|\varepsilon|B_z| \geq |B_z| - |H|y|B_z| > 0$$

vertices $z' \in B_z$ such that for every $u \in Y$, $w_u z' \in E(F)$ if and only if $uz \in E(H)$; let w_z be such a vertex. It follows that $\{w\} \cup \{w_z : z \in V_1(H) \setminus \{v\}\} \cup \{w_u : u \in Y\}$ forms a bi-induced copy of H in F , a contradiction.

Thus, there is no bi-induced copy of $H \setminus v$ in J . By the choice of a , J thus contains a clique or stable set I with $|I| \geq |J|^{1/a} \geq (y^{1-3a^2})^{1/a} = y^{-3a+1/a}$. Let $S := \bigcup_{i \in I} B_i$; then $|S| = |I|m$ since $|B_i| = m$ for all $i \in I$. If I is a stable set in J , then $F[S]$ has maximum degree at most

$$m + |I|\varepsilon m = (1/|I| + \varepsilon)|I|m \leq (y^{3a-1/a} + y^{3a})|S| \leq 2y^{3a-1}|S| \leq y^{2a}|S|;$$

and similarly, if I is a clique in J , then $\overline{F}[S]$ has maximum degree at most $y^{2a}|S|$. Thus $F[S]$ is y^{2a} -restricted which is the first outcome of the lemma, a contradiction. This proves the claim. \square

Now, by the claim, there exists $i \in [\ell]$ such that there are at most $y|D|$ vertices in D that are mixed on B_i . Since F is y -sparse, there are at most $y|F|$ vertices in D that are complete to B_i . Thus, since

$$|B_1| + \dots + |B_\ell| \leq \ell\varepsilon^2|F| \leq 2\varepsilon|F| = 2y^{3a^2}|F| \leq y|F|,$$

there are at least

$$|F| - y|D| - y|F| - (|B_1| + \dots + |B_\ell|) \geq (1 - 3y)|F|$$

vertices in F with no neighbour in B_i . Because $|B_i| \geq \varepsilon^b|F| = y^{3ba^2}|F|$, the second outcome of the lemma holds. This proves Lemma 3.4. \blacksquare

We now apply Lemma 3.4 repeatedly to move from an anticomplete pair to an anticomplete blockade.

Lemma 3.5. *Let H be a bigraph, and let $v \in V(H)$. Let $b \geq 2$ be given by Theorem 2.4. Assume there exists $a \geq 2$ such that every n -vertex graph with no bi-induced copy of $H \setminus v$ contains a clique or stable set of size at least $n^{1/a}$. Let $0 < y \leq 2^{-12|H|}$, and let F be a y -sparse graph with no bi-induced copy of H . Then either:*

- F has a y^a -restricted induced subgraph with at least $y^{2ba^2}|F|$ vertices; or
- there is an anticomplete $(y^{-1/2}, y^{2ba^2}|F|)$ -blockade in F .

Proof. Suppose that the second outcome does not hold. Let $n \geq 0$ be maximal such that there is a blockade (B_0, B_1, \dots, B_n) in F with $|B_n| \geq (1 - 3y^{1/2})^n|F|$ and $|B_{i-1}| \geq y^{2ba^2}|F|$ for all $i \in [n]$. Since the second outcome does not hold, $n < y^{-1/2}$; and so, since $y \leq 2^{-12}$,

$$|B_n| \geq (1 - 3y^{1/2})^n|F| \geq 4^{-3y^{1/2n}}|F| \geq 4^{-3}|F| \geq y^{1/2}|F|.$$

Hence $F[B_n]$ has maximum degree at most $y|F| \leq y^{1/2}|B_n|$; and so Lemma 3.4 (with $y^{1/2}$ in place of y , note that $y^{1/2} \leq 2^{-6|H|} < 1/|H|$) implies that either:

- there exists $S \subseteq B_n$ with $|S| \geq y^{3ba^2/2}|B_n|$ such that $F[S]$ is y^a -restricted; or
- there is an anticomplete pair (A, B) in $F[B_n]$ with $|A| \geq y^{3ba^2/2}|B_n|$ and $|B| \geq (1 - 3y^{1/2})|B_n|$.

If the second bullet holds, then $(B_0, B_1, \dots, B_{n-1}, A, B)$ would be a blockade contradicting the maximality of n since $|A| \geq y^{3ba^2/2}|B_n| \geq y^{3ba^2/2+1/2}|F| \geq y^{2ba^2}|F|$. Thus the first bullet holds; and so $S \subseteq V(F)$ is y^a -restricted in F and satisfies $|S| \geq y^{3ba^2/2}|B_n| \geq y^{3ba^2/2+1/2}|G| \geq y^{2ba^2}|F|$. Hence the first outcome of the lemma holds. This proves Lemma 3.5. \blacksquare

We have been assuming that our restricted graphs are sparse. We now move to the general case, by handling dense and sparse graphs simultaneously.

For a bigraph H , its *bicomplement* is the bigraph $\overline{\overline{H}}$ with the same bipartition and edge set $\{uv : u \in V_1(H), v \in V_2(H), uv \notin E(H)\}$.

Lemma 3.6. *Let H be a bigraph, and let $v \in V(H)$. Let b be given by Theorem 2.4. Assume there exists $a \geq 2$ such that every n -vertex graph with no bi-induced copy of $H \setminus v$ contains a clique or stable set of size at least $n^{1/a}$. Let $0 < y \leq 2^{-12|H|}$, and let F be a y -restricted graph with no bi-induced copy of H . Then either:*

- F has a y^a -restricted induced subgraph with at least $y^{2ba^2}|F|$ vertices; or
- there is a complete or anticomplete $(y^{-1/2}, y^{2ba^2}|F|)$ -blockade in F .

Proof. Since F is y -restricted, either F or \overline{F} is y -sparse. If F is y -sparse then the result follows from Lemma 3.5 (using the assumption on $H \setminus v$). If \overline{F} is y -sparse then the result follows from applying Lemma 3.5 to the complement of F (noting that \overline{F} is $\overline{\overline{H}}$ -free, and using the assumption on $\overline{\overline{H}} \setminus v$). ■

With Lemma 3.6 in hand, we can now prove that the conditions of Lemma 3.3 are satisfied.

Lemma 3.7. *For every bigraph H , there exists $d \geq 2$ such that for every $x \in (0, 2^{-d})$ and every graph G with no bi-induced copy of H , either:*

- G has an x -restricted induced subgraph with at least $x^d|G|$ vertices; or
- there is a complete or anticomplete $(k, |G|/k^d)$ -blockade in G , for some $k \in [2, 1/x]$.

Proof. We argue by induction on $|H|$. We may assume that $|H| \geq 2$. Choose $v \in V(H)$. By Lemma 3.3 and the induction hypothesis applied to $H \setminus v$ and $\overline{\overline{H}} \setminus v$, there exists $a \geq 4$ such that every n -vertex graph with no bi-induced copy of $H \setminus v$, and every n -vertex graph with no bi-induced copy of $\overline{\overline{H}} \setminus v$, contains a clique or stable set of size at least $n^{1/a}$. Let $c := 2^{-12|H|}$, and let $b \geq 2$ be given by Theorem 2.4. By Rödl's Theorem 1.3, we can choose some $t \geq ba^2$ such that every graph G with no bi-induced copy of H contains a c -restricted induced subgraph with at least $c^t|G|$ vertices. We claim that $d := 2 \max(at, |H|) \geq 8t$ satisfies the lemma.

To show this, let $x \in (0, 2^{-d}) \subseteq (0, c)$, and suppose that G has no bi-induced copy of H . Suppose that the second outcome of the lemma does not hold; that is, there is no $k \in [2, 1/x]$ such that there is a complete or anticomplete $(k, |G|/k^d)$ -blockade in G .

Claim. *For every $y \in [x, c]$ and every y -restricted induced subgraph F or G with $|F| \geq y^{2t}|G|$, there is a y^a -restricted induced subgraph of F with at least $y^{at}|F|$ vertices.*

Proof. By the choice of t and Lemma 3.6, either:

- F has a y^a -restricted induced subgraph with at least $y^{2ba^2}|F| \geq y^{2t}|F| \geq y^{at}|F|$ vertices; or
- there is a complete or anticomplete $(y^{-1/2}, y^{2ba^2}|F|)$ -blockade in F .

If the second bullet holds, then since $y^{2ba^2}|F| \geq y^{4t}|G| \geq y^{d/2}|G|$ by the choice of d , there would be a complete or anticomplete $(y^{-1/2}, y^{d/2}|G|)$ blockade in G , which contradicts that the second outcome of the lemma does not hold (note that $y^{1/2} \leq c^{1/2} \leq \frac{1}{2}$). Thus the first bullet holds, proving the claim. □

Lemma 3.2 and the claim imply that G has an x -restricted induced subgraph with at least $x^{2at}|G| \geq x^d|G|$ vertices, which is the first outcome of the lemma. This proves Lemma 3.7. ■

Lemmas 3.3 and 3.7 show the Erdős-Hajnal property for every class of graphs with no bi-induced copy of a given bipartite graph. But much of the argument so far was about density, and that allows us to go a step further, and apply Theorem 2.4 once more to prove Theorem 3.1.

Proof of Theorem 3.1. It follows from Lemmas 3.3 and 3.7 that there exists $a \geq 2$ such that every n -vertex graph with no bi-induced copy of H has a clique or stable set of size at least $n^{1/a}$; and by increasing a if necessary we may assume $2^a > |H|$. Let $b \geq 2$ be given by Theorem 2.4. We claim that $C := 2ab$ satisfies the theorem. To see this, let $\varepsilon \in (0, \frac{1}{2})$ and G be a graph with no bi-induced copy of H . It suffices to show that G has an ε -restricted induced subgraph with at least $\varepsilon^{2ab}|G|$ vertices.

We may assume that $|G| > \varepsilon^{-2ab}$. By Theorem 2.4 with ε^{2a} in place of ε , there is an $(\varepsilon^{-2a}, \varepsilon^{2ab}|G|)$ -blockade (B_1, \dots, B_ℓ) in G where $\ell \geq \varepsilon^{-2a}$, such that $|B_1| = \dots = |B_\ell| =: m$; and for all distinct $i, j \in [\ell]$, B_i, B_j are ε^{2a} -sparse to each other in G or \overline{G} . Let J be the graph with vertex set $[\ell]$ where $ij \notin E(J)$ if and only if B_i, B_j are ε^{2a} -sparse to each other in G .

Claim. J has no bi-induced copy of H .

Proof. Suppose not; and we may assume $V(H) \subseteq V(J)$. For every $i \in V(H)$, let v_i be a uniformly random vertex in B_i , chosen independently. Then for (X, Y) the bipartition of H , the probability that $\{v_1, \dots, v_{|H|}\}$ forms a bi-induced copy of H in G is at least $1 - |X||Y|\varepsilon^{2a} > 1 - |H|2^{2-2a} > 0$ by the choice of a ; and so there is a bi-induced copy of H in G , a contradiction. This proves the claim. \square

Now, by the claim and the choice of a , J has a clique or stable set I with $|I| \geq |J|^{1/a} \geq \varepsilon^{-2}$. Let $S := \bigcup_{i \in I} B_i$; then $|S| = |I|m \geq m \geq \varepsilon^{2ab}|G|$. If I is stable in J , then $G[S]$ has maximum degree at most

$$m + |I|\varepsilon^a m = (1/|I| + \varepsilon^a)|I|m \leq 2\varepsilon^2|S| \leq \varepsilon|S|.$$

Similarly, if I is a clique in J , then $\overline{G}[S]$ has maximum degree at most $\varepsilon|S|$. This proves Theorem 3.1. \blacksquare

4. RAMSEY NUMBERS FOR HYPERGRAPHS WITH BOUNDED VC-DIMENSION

A k -uniform hypergraph is a hypergraph whose edges all have cardinality k . The notion of VC-dimension extends naturally to k -uniform hypergraphs \mathcal{H} . For a set S , let $\binom{S}{k}$ denote the family of subsets of size k of S ; and for $S \in \binom{V(\mathcal{H})}{k-1}$, let $N_{\mathcal{H}}(S) := \{v \in V(\mathcal{H}) : S \cup \{v\} \in E(\mathcal{H})\}$. The VC-dimension of \mathcal{H} is the VC-dimension of the family $\{N_{\mathcal{H}}(S) : S \in \binom{V(\mathcal{H})}{k-1}\}$ with ground set $V(\mathcal{H})$.

A *clique* in a k -uniform hypergraph \mathcal{H} is a subset of $V(\mathcal{H})$ whose subsets of size k are all in $E(\mathcal{H})$, and a *stable set* is a clique in the complementary k -uniform hypergraph. The *Ramsey number* $R_k(t)$ is the smallest integer $n \geq 1$ such that every n -vertex k -graph contains a clique or stable set of size at least t . By Ramsey's theorem, $R_k(n)$ exists for all values of n and k ; equivalently, writing $f_k(n)$ for the largest m such that every k -uniform hypergraph with n vertices has a clique or stable set of size at least m , we have $f_k(n) \rightarrow \infty$ for any fixed k . The growth rate of f_k has received considerable attention: for $k = 2$, it is well known that $f_2(n) = \Theta(\log n)$, but for $k \geq 3$ the bounds are much farther apart. For $k = 3$, Erdős, Hajnal and Rado [20] showed that

$$c_1 \log \log n \leq f_3(n) \leq c_2 \sqrt{\log n};$$

and for $k \geq 4$, the best bounds are

$$c_1 \log^{(k-1)}(n) \leq f_k(n) \leq c_2 \sqrt{\log^{(k-2)} n},$$

where $\log^{(k)} n$ denotes the k -times iterated logarithm.

What can be said about hypergraphs with bounded VC-dimension? Let $f_k^d(n)$ be the largest integer m such that every k -uniform hypergraph with n vertices and VC-dimension at most d has a clique or stable set of size at least m . For $k \geq 3$, the best previous bounds are

$$e^{(\log^{(k-1)} n)^{1-o(1)}} \leq f_k^d(n) \leq c_2 \log^{(k-2)}(n),$$

where the upper follows from a construction of Conlon, Fox, Pach, Sudakov, and Suk [17] and the lower bound is due to Fox, Pach, and Suk [22], who observed that it follows from their $2^{(\log n)^{1-o(1)}}$ bound for graphs of bounded VC-dimension and by an adaptation of the classical Erdős–Rado greedy argument [21].

Stronger lower bounds have been obtained in more restrictive settings. Conlon, Fox, Pach, Sudakov, and Suk [17] proved that every semi-algebraic n -vertex k -uniform hypergraph of bounded description complexity admits a clique or stable set of size at least $\exp(c \log^{(k-1)} n)$, where the constant $c > 0$ depends on k and the description complexity of the hypergraph. This was later extended to k -uniform hypergraphs definable by distal structures by Chernikov, Starchenko, and Thomas [12]. Combining the Erdős–Rado argument with our Theorem 1.2 allows us to extend this further to hypergraphs with bounded VC-dimension, improving the lower bound to show that, for fixed $k \geq 3$ and $d \geq 1$,

$$f_k^d(n) = e^{\Theta(\log^{(k-1)} n)}.$$

We give the details for completeness. For integers $d \geq 1$, $k \geq 2$, and $n \geq 1$, let $R_k^d(n)$ be the smallest integer $m \geq 1$ such that every m -vertex k -graph of VC-dimension at most d contains a clique or stable set of size at least n ; then Theorem 1.2 says that $R_2^d(n) \leq n^C$ for all $n \geq 1$, for some $C > 0$ depending on d only.

Theorem 4.1. *For all integers $d \geq 1$, $k \geq 3$, and $n \geq 2$,*

$$R_k^d(n) \leq 2^{\binom{R_{k-1}^d(n-1)}{k-1}} + k - 2.$$

Proof. Let $p := R_{k-1}^d(n-1)$. Let \mathcal{H} be a k -uniform hypergraph with $2^{\binom{p}{k-1}} + k - 2$ vertices and VC-dimension at most d . We claim that:

Claim 4.2. *For every integer q with $k - 2 \leq q \leq p$, there are disjoint $A_q, B_q \subseteq V(\mathcal{H})$ with $|A_q| = q$ and $|B_q| \geq 2^{\binom{p}{k-1} - \binom{q}{k-1}}$ such that for every $S \subseteq \binom{A_q}{k-1}$, either $S \cup \{v\} \in E(\mathcal{H})$ for all $v \in B_q$ or $S \cup \{v\} \notin E(\mathcal{H})$ for all $v \in B_q$.*

Proof. The claim is vacuously true for $q = k - 2$. For $k - 2 \leq q < p$, we shall prove the claim of $q + 1$ assuming that it is true for q . Let $u \in B_q$ be arbitrary, and let $A_{q+1} := A_q \cup \{u\}$. For every $T \in \binom{A_q}{k-2}$ and $v \in B_q \setminus \{u\}$, let $f_T(v)$ be 0 if $T \cup \{u, v\} \in E(\mathcal{H})$ and 1 if $T \cup \{u, v\} \notin E(\mathcal{H})$. Then by the pigeonhole principle, there exists $B_{q+1} \subseteq B_q$ with

$$\begin{aligned} |B_{q+1}| &\geq \left\lceil 2^{-\binom{|A_q|}{k-2}} (|B_q| - 1) \right\rceil \geq \left\lceil 2^{-\binom{q}{k-2}} \left(2^{\binom{p}{k-1} - \binom{q}{k-1}} - 1 \right) \right\rceil \\ &= \left\lceil 2^{\binom{p}{k-1} - \binom{q+1}{k-1}} - 2^{-\binom{q}{k-2}} \right\rceil = 2^{\binom{p}{k-1} - \binom{q+1}{k-1}}, \end{aligned}$$

such that $f_T(v) = f_T(w)$ for all $T \in \binom{A_q}{k-2}$ and $v, w \in B_{q+1}$; here the last equation holds since $\binom{q+1}{k-1} \leq \binom{p}{k-1}$ and $\binom{q}{k-2} \geq 1$. Thus, by the choice of A_q, B_q , it follows that for every $S \subseteq \binom{A_{q+1}}{k-1}$, either $S \cup \{v\} \in E(\mathcal{H})$ for all $v \in B_{q+1}$ or $S \cup \{v\} \notin E(\mathcal{H})$ for all $v \in B_{q+1}$. This proves Claim 4.2. \square

Now, Claim 4.2 with $q = p$ gives disjoint $A_p, B_p \subseteq V(\mathcal{H})$ with $|A_p| = p$ and $|B_p| \geq 1$ such that for every $S \subseteq \binom{A_p}{k-1}$, either $S \cup \{v\} \in E(\mathcal{H})$ for all $v \in B_p$ or $S \cup \{v\} \notin E(\mathcal{H})$ for all $v \in B_p$. Let $v \in B_p$; and let \mathcal{H}' be the $(k-1)$ -uniform hypergraph with vertex set A_p where for every $S \in \binom{A_p}{k-1}$, $S \in E(\mathcal{H}')$ if and only if $S \cup \{v\} \in E(\mathcal{H})$.

Claim 4.3. *\mathcal{H}' has VC-dimension at most d .*

Proof. Suppose not; then there exists $A \subseteq A_p$ with $|A| > d$ which is shattered by the family $\{N_{\mathcal{H}'}(T) : T \in \binom{A_p}{k-2}\}$ with ground set A_p . However, by the definition of \mathcal{H}' , A is then shattered by the family $\{N_{\mathcal{H}}(T \cup \{v\}) : T \in \binom{A_p}{k-2}\} \subseteq \{N_{\mathcal{H}}(S) : S \in \binom{V(\mathcal{H})}{k-1}\}$ with ground set $V(\mathcal{H})$, contrary to \mathcal{H} having VC-dimension at most d . This proves Claim 4.3. \square

Now, by the definition of p , Claim 4.3 implies that \mathcal{H}' has a clique or stable set S with $|S| \geq n - 1$. Then $S \cup \{v\}$ is a clique or stable set in \mathcal{H} of size at least n . This proves Theorem 4.1. \blacksquare

It is now not hard to iterate Theorem 4.1 and apply the bound $R_2^d(n) \leq n^C$ to get $R_k^d(n) \leq \text{twr}_{k-1}(n^K)$ for some $K > 0$ depending on d, k , where twr_k is defined recursively for $k \geq 1$ by $\text{twr}_1(t) := t$ and $\text{twr}_k(t) := 2^{\text{twr}_{k-1}(t)}$ for all $t > 0$. In other words, $f_k^d(n) \geq (\log^{(k-2)} n)^{1/K}$.

5. THE VIRAL PROPERTY

Let us interpret Theorem 3.1 in the language of forbidden induced subgraphs. A *split* graph is a graph whose vertex set can be partitioned into a clique and a stable set. For graphs H, G , a *copy* of H in G is an injective map $\varphi : V(H) \rightarrow V(G)$ such that for all distinct $u, v \in V(H)$, $uv \in E(H)$ if and only if $\varphi(u)\varphi(v) \in E(G)$. For a finite family \mathcal{F} of graphs, a graph G is \mathcal{F} -free if there is no copy of H in G for all $H \in \mathcal{F}$. The following is a well-known strengthening of Lemma 2.2 (see [7, Theorem 3.3]).

Lemma 5.1. *For every two bipartite graphs H_1, H_2 and every split graph J , there exists $d \geq 1$ such that every $\{H_1, \overline{H_2}, J\}$ -free graph has VC-dimension at most d . Conversely, for every $d \geq 1$, there are bipartite graphs H_1, H_2 and a split graph J such that every graph of VC-dimension at most d is $\{H_1, \overline{H_2}, J\}$ -free.*

Thus, Theorem 1.5 can be rewritten as:

Theorem 5.2. *For every two bipartite graphs H_1, H_2 and every split graph J , the class of $\{H_1, \overline{H_2}, J\}$ -free graphs has the polynomial Rödl property.*

We say that a finite family \mathcal{F} of graphs is *viral* if there exists $C > 0$ such that for every $\varepsilon \in (0, \frac{1}{2})$ and for every graph G with at most $(\varepsilon^C |G|)^{|H|}$ copies of H for all $H \in \mathcal{F}$, there is an ε -restricted induced subgraph of G with at least $\varepsilon^C |G|$ vertices. Thus, if \mathcal{F} is viral then the class of \mathcal{F} -free graphs has the polynomial Rödl property. It is conjectured in [40] that $\{H\}$ is viral for all graphs H , which would mean that Nikiforov's theorem [42] holds with polynomial dependence. Recently, Gishboliner and Shapira [28] showed that for every two bipartite graphs H_1, H_2 and every split graph J , $\{H_1, \overline{H_2}, J\}$ is viral if and

only if the class of $\{H_1, \overline{H_2}, J\}$ -free graphs has the Erdős-Hajnal property, by using a strengthening of Theorem 2.1 for graphs with few copies of $H_1, \overline{H_2}, J$. Later, Bucić, Fox, and Pham [8] showed that for any finite family \mathcal{F} , the class of \mathcal{F} -free graphs has the Erdős-Hajnal property if and only if \mathcal{F} is viral. Thus one can bootstrap Theorem 5.2 into the following.

Theorem 5.3. *For every two bipartite graphs H_1, H_2 and every split graph J , $\{H_1, \overline{H_2}, J\}$ is viral.*

6. TOURNAMENTS

If Q is a tournament, a tournament is Q -free if it has no subtournament isomorphic to Q . Say that Q has the *Erdős-Hajnal property* if there exists $c > 0$ such that every n -vertex Q -free tournament admits a transitive subtournament with at least n^c vertices. (A warning: earlier we were using this name applied to hereditary class, and now we are using it for a tournament, not for the class that does not contain it.) Alon, Pach, and Solymosi [3] proved that Conjecture 1.1 is equivalent to the following statement.

Conjecture 6.1. *Every tournament has the Erdős-Hajnal property.*

For a tournament Q and an ordering $\phi = (v_1, \dots, v_n)$ of $V(Q)$, the *backedge graph* of Q with respect to ϕ is the graph G with vertex set $V(Q)$, where for all $i, j \in [n]$ with $i < j$, $v_i v_j \in E(G)$ if and only if $(v_j, v_i) \in E(Q)$. For two tournaments Q_1, Q_2 , the tournament Q obtained by *substituting* Q_2 for $v \in V(Q_1)$ has vertex set $V(Q_2) \cup (V(Q_1) \setminus \{v\})$, such that $Q[V(Q_1) \setminus \{v\}] = Q_1 \setminus v$, $Q[V(Q_2)] = Q_2$, and for all $u \in V(Q_1) \setminus \{v\}$ and $w \in V(Q_2)$, $(u, w) \in E(Q)$ if and only if $(u, v) \in E(Q_1)$. A tournament is *prime* if it cannot be obtained by substitution from tournaments with fewer vertices. By an adaptation of another argument of Alon, Pach, and Solymosi in [3], Conjecture 1.1 reduces to proving that all prime tournaments have the Erdős-Hajnal property. Partial results in this direction include [5, 13, 39, 51–53]; these show the Erdős-Hajnal property of several types of prime tournaments, that all admit a forest backedge graph except those from [39]. In another direction, Berger, Choromanski, and Chudnovsky [6] proved Conjecture 6.1 for every six-vertex tournament except for the seven-vertex Paley tournament with one vertex removed.

For an integer $k \geq 0$, a tournament is k -colourable if its vertex set can be partitioned into k subsets each inducing a transitive subtournament. It is known (see [15, Lemma 3.2]) that a class of tournaments closed under taking subtournaments does not contain all two-colourable tournaments if and only if for some d , the VC-dimension of the set of in-neighbourhoods of each of its members is at most d . Moreover, a standard argument together with Lemma 2.2 proves the following (see [38, Section 4.6] for a proof).

Lemma 6.2. *For every $d \geq 1$, there exists $m \geq 1$ such that for every tournament T , if the set of its in-neighbourhoods has VC-dimension at most d then every backedge graph of T has VC-dimension at most m . Conversely, for every $m \geq 1$, there exists $d \geq 1$ such that if a backedge graph of a tournament T has VC-dimension at most m , then the set of in-neighbourhoods of T has VC-dimension at most d .*

Therefore, since every backedge graph of every n -vertex transitive tournament is a comparability graph (and so has a clique or stable set of size at least \sqrt{n}), Theorem 1.2 can be restated as follows.

Theorem 6.3. *Every two-colourable tournament has the Erdős-Hajnal property.*

Tournaments that admit forest backedge graphs are two-colourable, so this result contains all of the aforementioned results on prime tournaments with the Erdős-Hajnal property, except those from [39]; and since all six-vertex tournaments are two-colourable (see [37]), Theorem 6.3 implies that they all satisfy Conjecture 6.1 as well. One can also formulate a viral version for tournaments similar to the one for graphs and prove that every two-colourable tournament is viral; we omit the details.

ACKNOWLEDGEMENTS

We would like to thank Artem Chernikov and Jacob Fox for helpful discussions. We are also grateful to the anonymous referees for helpful remarks.

REFERENCES

- [1] N. Alon, E. Fischer, and I. Newman. Efficient testing of bipartite graphs for forbidden induced subgraphs. *SIAM J. Comput.*, 37(3):959–976, 2007. 3
- [2] N. Alon, J. Pach, R. Pinchasi, R. Radoičić, and M. Sharir. Crossing patterns of semi-algebraic sets. *J. Combin. Theory Ser. A*, 111(2):310–326, 2005. 1, 2
- [3] N. Alon, J. Pach, and J. Solymosi. Ramsey-type theorems with forbidden subgraphs. *Combinatorica*, 21(2):155–170, 2001. Paul Erdős and his mathematics (Budapest, 1999). 10

- [4] S. Basu. Combinatorial complexity in o-minimal geometry. *Proc. Lond. Math. Soc. (3)*, 100(2):405–428, 2010. [2](#)
- [5] E. Berger, K. Choromanski, and M. Chudnovsky. Forcing large transitive subtournaments. *J. Combin. Theory Ser. B*, 112:1–17, 2015. [10](#)
- [6] E. Berger, K. Choromanski, and M. Chudnovsky. On the Erdős–Hajnal conjecture for six-vertex tournaments. *European J. Combin.*, 75:113–122, 2019. [10](#)
- [7] N. Bousquet, A. Lagoutte, Z. Li, A. Parreau, and S. Thomassé. Identifying codes in hereditary classes of graphs and VC-dimension. *SIAM J. Discrete Math.*, 29(4):2047–2064, 2015. [3](#), [9](#)
- [8] M. Bucić, J. Fox, and H. Pham. Equivalence between Erdős–Hajnal and polynomial Rödl and Nikiforov conjectures, [arXiv:2403.08303](#), 2024. [3](#), [10](#)
- [9] M. Bucić, T. Nguyen, A. Scott, and P. Seymour. Induced subgraph density. I. A loglog step towards Erdős–Hajnal, *Int. Math. Res. Not. IMRN*, (12):9991–1004, 2024. [1](#), [2](#)
- [10] A. Chernikov and S. Starchenko. A note on the Erdős–Hajnal property for stable graphs. *Proc. Amer. Math. Soc.*, 146(2):785–790, 2018. [2](#)
- [11] A. Chernikov and S. Starchenko. Regularity lemma for distal structures. *J. Eur. Math. Soc.*, 20(10):2437–2466, 2018. [2](#)
- [12] A. Chernikov, S. Starchenko, and M. E. M. Thomas. Ramsey growth in some NIP structures. *J. Inst. Math. Jussieu*, 20(1):1–29, 2021. [2](#), [8](#)
- [13] K. Choromanski. EH-suprema of tournaments with no nontrivial homogeneous sets. *J. Combin. Theory Ser. B*, 114:97–123, 2015. [10](#)
- [14] M. Chudnovsky. The Erdős–Hajnal conjecture—a survey. *J. Graph Theory*, 75(2):178–190, 2014. [1](#)
- [15] M. Chudnovsky, R. Kim, C.-H. Liu, P. Seymour, and S. Thomassé. Domination in tournaments. *J. Combin. Theory Ser. B*, 130:98–113, 2018. [10](#)
- [16] M. Chudnovsky, A. Scott, P. Seymour, and S. Spirkl. Erdős–Hajnal for graphs with no 5-hole. *Proc. Lond. Math. Soc. (3)*, 126(3):997–1014, 2023. [1](#), [2](#)
- [17] D. Conlon, J. Fox, J. Pach, B. Sudakov, and A. Suk. Ramsey-type results for semi-algebraic relations. *Trans. Amer. Math. Soc.*, 366(9):5043–5065, 2014. [8](#)
- [18] P. Erdős and A. Hajnal. On spanned subgraphs of graphs. In *Contributions to Graph Theory and Its Applications (Internat. Colloq., Oberhof, 1977) (German)*, pages 80–96. Tech. Hochschule Ilmenau, Ilmenau, 1977. [1](#)
- [19] P. Erdős and A. Hajnal. Ramsey-type theorems. *Discrete Applied Math.*, 25(1-2):37–52, 1989. [1](#)
- [20] P. Erdős, A. Hajnal, A and R. Rado, Partition relations for cardinal numbers, *Acta Math. Acad. Sci. Hungar.* 16, 93–196, 1965 [8](#)
- [21] P. Erdős and R. Rado. Combinatorial theorems on classifications of subsets of a given set. *Proc. London Math. Soc. (3)*, 2:417–439, 1952. [8](#)
- [22] J. Fox, J. Pach, and A. Suk. Erdős–Hajnal conjecture for graphs with bounded VC-dimension. *Discrete Comput. Geom.*, 61(4):809–829, 2019. [1](#), [3](#), [8](#)
- [23] J. Fox, J. Pach, and A. Suk. Bounded VC-dimension implies the Schur–Erdős conjecture. *Combinatorica*, 41(6):803–813, 2021. [1](#)
- [24] J. Fox, J. Pach, and A. Suk. Sunflowers in set systems of bounded dimension. *Combinatorica*, 43(1):187–202, 2023. [1](#)
- [25] J. Fox, J. Pach, and C. D. Tóth. Intersection patterns of curves. *J. Lond. Math. Soc. (2)*, 83(2):389–406, 2011. [1](#)
- [26] J. Fox and B. Sudakov. Induced Ramsey-type theorems. *Adv. Math.*, 219(6):1771–1800, 2008. [2](#)
- [27] Y. Fu. A note on Erdős–Hajnal property for graphs with VC dimension ≤ 2 , [arXiv:2310.16970](#), 2023. [2](#)
- [28] L. Gishboliner and A. Shapira. On Rödl’s theorem for cographs. *Electron. J. Combin.*, 30(4):Paper No. 4.13, 2023. [9](#)
- [29] A. Gyárfás. Reflections on a problem of Erdős and Hajnal. In *The Mathematics of Paul Erdős, II*, volume 14 of *Algorithms Combin.*, pages 93–98. Springer, Berlin, 1997. [1](#)
- [30] O. Janzer and C. Pohoata. On the Zarankiewicz problem for graphs with bounded VC-dimension. *Combinatorica* 44(4):839–848, 2024. [1](#)
- [31] D. Larman, J. Matoušek, J. Pach, and J. Törőcsik. A Ramsey-type result for convex sets. *Bull. London Math. Soc.*, 26(2):132–136, 1994. [1](#)
- [32] L. Lovász and B. Szegedy. Regularity partitions and the topology of graphons. In *An Irregular Mind*, volume 21 of *Bolyai Soc. Math. Stud.*, pages 415–446. János Bolyai Math. Soc., Budapest, 2010. [3](#)
- [33] M. Malliaris and S. Shelah. Regularity lemmas for stable graphs. *Trans. Amer. Math. Soc.*, 366(3):1551–1585, 2014. [2](#)
- [34] M. Malliaris and S. Shelah. Notes on the stable regularity lemma. *Bull. Symb. Log.*, 27(4):415–425, 2021. [2](#)
- [35] M. Malliaris and C. Terry. On unavoidable-induced subgraphs in large prime graphs. *J. Graph Theory*, 88(2):255–270, 2018. [2](#)
- [36] J. Matoušek. *Lectures on Discrete Geometry*, volume 212 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 2002. [2](#)
- [37] V. Neumann Lara. The 3- and 4-dichromatic tournaments of minimum order. *Discrete Math.*, 135(1-3):233–243, 1994. [10](#)
- [38] H. T. Nguyen. *Induced Subgraph Density*. PhD thesis, Princeton University, May 2025. [10](#)
- [39] T. Nguyen, A. Scott, and P. Seymour. Induced subgraph density. III. Cycles and subdivisions, [arXiv:2307.06379](#), 2024. [3](#), [4](#), [9](#)
- [40] T. Nguyen, A. Scott, and P. Seymour. Induced subgraph density. IV. New graphs with the Erdős–Hajnal property, [arXiv:2307.06455](#), 2023. [1](#), [3](#), [4](#), [10](#)
- [41] T. Nguyen, A. Scott, and P. Seymour. Induced subgraph density. VII. The five-vertex path, [arXiv:2312.15333](#), 2025. [1](#), [4](#)
- [42] V. Nikiforov. Edge distribution of graphs with few copies of a given graph. *Combin. Probab. Comput.*, 15(6):895–902, 2006. [9](#)

- [43] J. Pach and G. Tóth. How many ways can one draw a graph? *Combinatorica*, 26(5):559–576, 2006. [2](#)
- [44] V. Rödl. On universality of graphs with uniformly distributed edges. *Discrete Math.*, 59(1-2):125–134, 1986. [2](#)
- [45] L. Rónyai, L. Babai, and M. K. Ganapathy. On the number of zero-patterns of a sequence of polynomials. *J. Amer. Math. Soc.*, 14(3):717–735, 2001. [2](#)
- [46] P. Simon. *A guide to NIP theories*, volume 44 of *Lecture Notes in Logic*. Association for Symbolic Logic, Chicago, IL; Cambridge Scientific Publishers, Cambridge, 2015. [2](#)
- [47] P. Simon. A note on “Regularity lemma for distal structures”. *Proc. Amer. Math. Soc.*, 144(8):3573–3578, 2016. [2](#)
- [48] B. Sudakov and I. Tomon. Ramsey properties of algebraic graphs and hypergraphs. *Forum Math. Sigma*, 10:Paper No. e95, 20, 2022. [2](#)
- [49] I. Tomon. String graphs have the Erdős–Hajnal property. *J. Eur. Math. Soc.*, 26(1):275–287, 2024. [2](#)
- [50] V. N. Vapnik and A. Y. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. In *Measures of Complexity*, pages 11–30. Springer, Cham, 2015. Reprint of Theor. Probability Appl. 16 (1971), 264–280. [1](#)
- [51] S. Zayat and S. Ghazal. Erdős-Hajnal conjecture for new infinite families of tournaments. *J. Graph Theory*, 102(2):388–417, 2023. [10](#)
- [52] S. Zayat and S. Ghazal. Tournaments and the Erdős-Hajnal conjecture. *Australas. J. Combin.*, 86:351–372, 2023. [10](#)
- [53] S. Zayat and S. Ghazal. About the Erdős-Hajnal conjecture for seven-vertex tournaments. *Discrete Math.*, 347(1):Paper No. 113733, 23, 2024. [10](#)

PRINCETON UNIVERSITY, PRINCETON, NJ 08544, USA
Email address: tunghn@math.princeton.edu

MATHEMATICAL INSTITUTE, UNIVERSITY OF OXFORD, OXFORD OX2 6GG, UK
Email address: scott@maths.ox.ac.uk

PRINCETON UNIVERSITY, PRINCETON, NJ 08544, USA
Email address: pds@math.princeton.edu