Beyond chromatic threshold via (p, q)-theorem, and blow-up phenomenon

Hong Liu* Chong Shangguan† Jozef Skokan‡ Zixiang Xu*

April 17, 2024

Abstract

We establish a novel connection between the well-known chromatic threshold problem in extremal combinatorics and the celebrated (p, q)-theorem in discrete geometry. In particular, for a graph G with bounded clique number and a natural density condition, we prove a (p, q)-theorem for an abstract convexity space associated with G. Our result strengthens those of Thomassen and Nikiforov on the chromatic threshold of cliques. Our (p, q)-theorem can also be viewed as a χ-boundedness result for (what we call) ultra maximal K_r-free graphs.

We further show that the graphs under study are blow-ups of constant size graphs, improving a result of Oberkampf and Schacht on homomorphism threshold of cliques. Our result unravels the cause underpinning such a blow-up phenomenon, differentiating the chromatic and homomorphism threshold problems for cliques. Our result implies that for the homomorphism threshold problem, rather than the minimum degree condition usually considered in the literature, the decisive factor is a clique density condition on co-neighborhoods of vertices. More precisely, we show that if an n-vertex K_r-free graph G satisfies that the common neighborhood of every pair of non-adjacent vertices induces a subgraph with K_{r-2}-density at least $\varepsilon > 0$, then G must be a blow-up of some K_r-free graph F on at most $2^{O(\frac{r}{\varepsilon}\log\frac{1}{\varepsilon})}$ vertices. Furthermore, this single exponential bound is optimal. We construct examples with no K_r-free homomorphic image of size smaller than $2^{\Omega_r(\frac{1}{\varepsilon})}$.

1 Introduction

1.1 Overview

Finding sufficient conditions guaranteeing a graph to have bounded complexity has long been a popular topic in combinatorics and theoretical computer science. There are many natural ways to measure complexity. In this paper, the invariant we are interested in is the chromatic number, and we focus on graphs with bounded clique number. Note that graphs with bounded clique number could have arbitrarily large chromatic number: as one of the famous early applications of probabilistic method in combinatorics, Erdős [18] in 1959 constructed a graph with arbitrarily large girth (hence triangle-free) and chromatic number.

The density condition is a natural one to impose to guarantee constant chromatic number. Turán's theorem [45], a fundamental result in extremal graph theory, states that every n-vertex K_r-free graph

*Extremal Combinatorics and Probability Group (ECOPRO), Institute for Basic Science (IBS), Daejeon, South Korea. Emails: {hongliu, zixiangxu}@ibs.re.kr. Supported by IBS-R029-C4.

†Research Center for Mathematics and Interdisciplinary Sciences, Shandong University, Qingdao 266237, China, and the Frontiers Science Center for Nonlinear Expectations, Ministry of Education, Qingdao 266237, China. Email: theoreming@163.com. Supported by the National Natural Science Foundation of China under Grant Nos. 12101364 and 12231014, and the Natural Science Foundation of Shandong Province under Grant No. ZR2021QA005.

‡Department of Mathematics, London School of Economics, Houghton Street, London WC2A 2AE, UK. Email: j.skokan@lse.ac.uk

contains at most $\frac{r-2}{r-1} \cdot \frac{n^2}{2}$ edges, and, moreover, the balanced complete $(r-1)$-partite graph, also known as the Turán graph $T_{n,r-1}$, is the unique extremal graph. Rephrasing Turán's theorem, we see that every n-vertex K_r-free graph G with $e(G) \geqslant e(T_{n,r-1})$ must be $T_{n,r-1}$ and hence $(r-1)$-colorable. However, the number of edges $e(G)$ is not an ideal density condition. Indeed, if $e(G)$ drops slightly below $e(T_{n,r-1})$, although by a stability theorem of Erdős and Simonovits [21, 41], we know that G is close to $T_{n,r-1}$ (in edit-distance), but G could still have large chromatic number due to small noise, e.g. G could be a disjoint union of $T_{n-o(n),r-1}$ and an Erdős' graph mentioned above on $o(n)$ vertices. Notice that in both this example and the graph from random construction of Erdős, there are vertices with small degree. Motivated by this observation, Andrásfai, Erdős, and Sós [5] initiated the study of the relationship between the chromatic number and the minimum degree of a K_r-free graph. They showed that every n-vertex K_r-free graph with minimum degree larger than $\frac{3r-7}{3r-4} \cdot n$ has chromatic number at most $r-1$. It remained an interesting question to determine the optimal min-degree condition guaranteeing bounded chromatic number. This is the by-now well-known chromatic threshold problem first formulated in the early 1970s by Erdős and Simonovits [20], and since then there has been a large amount of work on this topic. Formally, for a graph H, its chromatic threshold is defined as

$$\delta_\chi(H) := \inf \left\{ \alpha \geqslant 0 : \exists\ C = C(\alpha, H) \text{ s.t. } \forall\ n\text{-vertex } H\text{-free } G,\ \delta(G) \geqslant \alpha n \Rightarrow \chi(G) \leqslant C \right\}.$$

In other words, for $\alpha < \delta_\chi(H)$, there exist H-free graphs with minimum degree αn and arbitrarily large chromatic number, but if $\alpha > \delta_\chi(H)$, then $\chi(G)$ must be bounded.

For the first non-trivial case when H is a triangle, a beautiful construction of Hajnal using the Kneser graph shows that $\delta_\chi(K_3) \geqslant \frac{1}{3}$ (see [20]); this example can be easily extended to obtain $\delta_\chi(K_r) \geqslant \frac{2r-5}{2r-3}$. It was not until 2002 that Thomassen [43] proved that Hajnal's construction is optimal: $\delta_\chi(K_3) = \frac{1}{3}$, as conjectured by Erdős and Simonovits [20]. Later, the chromatic thresholds for all cliques $\delta_\chi(K_r) = \frac{2r-5}{2r-3}$ were determined by Goddard and Lyle [22] and independently by Nikiforov [37]. After a series of results (see, e.g., [10, 11, 12, 24, 28, 34]), the culmination is the remarkable work of Allen, Böttcher, Griffiths, Kohayakawa, and Morris [1], which determined the chromatic thresholds for all graphs H.

One of our main contributions is to study the chromatic threshold problem through a geometric perspective, establishing a surprising connection between this classical problem in extremal combinatorics and the celebrated (p, q)-theorem in discrete geometry. We will elaborate more on this in the next subsection. Below is our first result.

Theorem 1.1. *Let $r \geqslant 3$, $\varepsilon > 0$ and G be an n-vertex K_r-free graph. If for every non-adjacent pair of vertices $u, v \in V(G)$, the induced subgraph $G[N(u) \cap N(v)]$ contains at least εn^{r-2} copies of K_{r-2}, then $\chi(G) = O_{r,\varepsilon}(1)$.*

Our theorem implies and extends the results of Thomassen [43], Goddard and Lyle [22], and Nikiforov [37] as the above clique density condition in the co-neighborhoods is a strictly weaker condition than the minimum degree condition in $\delta_\chi(K_r)$; see the discussion in Section 1.4.

Throughout the rest of this paper, we will refer the graphs in Theorem 1.1 as ε-ultra maximal K_r-free graphs.

1.2 Chromatic threshold meets (p, q)-theorem

In 1921, Radon [39] proved his fundamental lemma in combinatorial convexity which states that any set of $d+2$ points in $\mathbb{R}^d$ can be partitioned into two parts whose convex hulls intersect. Radon's lemma was introduced to prove Helly's theorem [26], yet another fundamental result in convexity. Helly's theorem is a local-global type result, stating that for any finite family of convex sets in $\mathbb{R}^d$, if every $d+1$ of them intersect, then all of them intersect. Since then, many generalizations have been developed. We refer the readers to [8, 9, 13, 17]. One of the important extensions, due to Katchalski and Liu [29], is the fractional Helly theorem. It states that if in a finite family of convex sets in $\mathbb{R}^d$, a positive fraction of $(d+1)$-tuples intersect, then a positive fraction of all sets intersect. For $p \geqslant q$, a family of sets has (p,q)-property if every of its p-tuple contains an intersecting q-tuple. In the early 90s, settling an old conjecture of Hadwiger and Debrunner [23], Alon and Kleitman [3] proved the following famous (p,q)-theorem, a far-reaching generalization of Helly's theorem. For a set system $\mathcal{G}$ defined on a ground set X, a set $T \subseteq X$ is a transversal for $\mathcal{G}$ if for every $A \in \mathcal{G}$, $T \cap A \neq \varnothing$. The transversal number of $\mathcal{G}$, denoted by $\tau(\mathcal{G})$, is the minimum size of a transversal for $\mathcal{G}$.

Theorem 1.2 ((p,q)-theorem)**.** *Let d, p, q be positive integers with $p \geqslant q \geqslant d+1$. If $\mathcal{F}$ is a finite family of convex sets in $\mathbb{R}^d$ with (p,q)-property, then $\tau(\mathcal{F}) \leqslant O_{d,p,q}(1)$.*

An interesting direction of research in discrete geometry is to prove classical results in Euclidean convexity in an abstract way. We will also work in an axiomatic setting through abstract convexity spaces. By now, abstract versions of Helly's theorem and many of its variants, including the fractional Helly theorem, and (p,q)-theorem have been studied in abstract convexity spaces, see e.g. [4].

We will cast the problem in Theorem 1.1 in the language of convexity spaces, which turns out to be equivalent to the following.

Theorem 1.3. *Let $r \geqslant 3$, $\varepsilon > 0$ and $\mathcal{B}$ be a set system with $(r,2)$-property. If for every intersecting pair $A, B \in \mathcal{B}$, at least ε fraction of the $(r-2)$-tuples $\mathcal{T}$ in $\mathcal{B}$ satisfies that both $\mathcal{T} \cup \{A\}$ and $\mathcal{T} \cup \{B\}$ are matchings of size $(r-1)$, then $\tau(\mathcal{B}) = O_{r,\varepsilon}(1)$.*

1.3 Blow-up phenomenon

Given two graphs G and F, G is homomorphic to F, denoted by $G \xrightarrow{\text{hom}} F$, if there exists a homomorphism $\varphi : V(G) \to V(F)$ preserving adjacencies, i.e. if $uv \in E(G)$, then $\varphi(u)\varphi(v) \in E(F)$. Note that $\chi(G) = t$ is equivalent to G being homomorphic to K_t. A natural extension of the chromatic threshold problem raised in [43] asks for minimum degree condition for an H-free graph guaranteeing a bounded size homomorphic image which is also H-free. Formally, the homomorphism threshold of a graph H is defined as

$$\delta_{\text{hom}}(H) := \inf\left\{\alpha \geqslant 0 : \exists\ H\text{-free } F = F(\alpha, H) \text{ s.t. } \forall\ n\text{-vertex } H\text{-free } G,\ \delta(G) \geqslant \alpha n \Rightarrow G \xrightarrow{\text{hom}} F\right\}.$$

As having a bounded homomorphic image implies bounded chromatic number, $\delta_{\text{hom}}(H) \geqslant \delta_\chi(H)$. The first such result was proved by Łuczak [33], showing that $\delta_{\text{hom}}(K_3) = \frac{1}{3}$. Later, Goddard and Lyle [22] resolved the clique case, proving that $\delta_{\text{hom}}(K_r) = \delta_\chi(K_r) = \frac{2r-5}{2r-3}$ for all positive integers $r \geqslant 3$. The proofs of Łuczak [33], and Goddard and Lyle [22] utilized Szemerédi regularity lemma [42], and therefore gave a tower-type upper bound on the size of the K_r-free homomorphic image F. Recently, Oberkampf and Schacht [38] gave a new proof using a clever probabilistic argument. Their proof yields a double exponential bound $2^{2^{O(1/\varepsilon^2)}}$ on the size of the homomorphic image for K_r-free graphs with minimum degree at least $(\frac{2r-5}{2r-3} + \varepsilon)n$, which was the best-known bound for all $r \geqslant 4$. As for $r = 3$, a beautiful (unpublished) result of Brandt and Thomassé [11] proved an optimal bound that any triangle-free graph G with $\delta(G) \geqslant (\frac{1}{3} + \varepsilon)n$ is homomorphic to a triangle-free graph (Vega graphs) of size $O\big(\frac{1}{\varepsilon}\big)$.

Except for cliques, determining the homomorphism thresholds for any other graph is still widely open. Letzter and Snyder [15] showed that $\delta_{\text{hom}}(C_5) \leqslant \frac{1}{5}$. Later, Ebsen and Schacht [30] proved that $\delta_{\text{hom}}(C_{2r+1}) \leqslant \frac{1}{2r+1}$ for every $r \geqslant 2$. A recent significant advancement, due to Sankar [40], shows that $\delta_{\text{hom}}(C_{2r+1}) > 0$. It was known that $\delta_\chi(C_{2r+1}) = 0$ [44], so her result provides the first example H with $\delta_{\text{hom}}(H) > \delta_\chi(H)$.

An important variation was studied by Łuczak and Thomassé [34] in their ingenious new proof of $\delta_\chi(K_3) = \frac{1}{3}$. They realized that the $n/3$ minimum degree condition can be relaxed to a linear one if we have bounded VC-dimension (see Section 2.1 for definition). That is, any n-vertex triangle-free graph G with bounded VC-dimension and $\delta(G) = \Omega(n)$ has bounded chromatic number (see [32, Theorem 5.1]).

Our next result shows that VC-dimension, while significant in the chromatic threshold problem, is not so influential in the homomorphism threshold problem.

Theorem 1.4. *There exists an n-vertex triangle-free graph G with VC-dimension 3 and $\delta(G) \geqslant \frac{n}{4}$ such that G has no triangle-free homomorphic image of size smaller than $\frac{n}{4}$.*

What is more relevant then? This is the content of our next main result, which strengthens Theorem 1.1. Our theorem shows that the graph under study is in fact a blow-up of a constant size graph, which then necessarily also has to be K_r-free. For a graph H, we write $H[t]$ for the t-blow-up of H obtained by replacing every vertex of H by t independent copies (i.e. every vertex becomes an independent set of size t and every edge becomes a copy of $K_{t,t}$). We simply write $H[\cdot]$ when mentioning a blow-up of H without specifying its size (and the sizes of these independent sets could be different). Obviously, if $G = F[\cdot]$, then $G \xrightarrow{\text{hom}} F$.

Theorem 1.5. *Given $r \geqslant 3, \varepsilon > 0$ and G be an n-vertex K_r-free graph. If for every non-adjacent pair of vertices $u, v \in V(G)$, the induced subgraph $G[N(u) \cap N(v)]$ contains εn^{r-2} copies of K_{r-2}, then $G = F[\cdot]$ for some maximal K_r-free graph F on at most $2^{O((\frac{1}{\varepsilon}+r)\log\frac{1}{\varepsilon})}$ vertices.*

The merit of Theorem 1.5 is that it reveals the cause of the blow-up phenomenon $G = F[\cdot]$. Indeed, since the chromatic and homomorphism thresholds mysteriously coincide for cliques: $\delta_{\text{hom}}(K_r) = \delta_\chi(K_r)$, it was not clear what separates these two problems for cliques. Theorem 1.5 shows that for the blow-up phenomenon, rather than VC-dimension, the main driving factor is the clique density condition in the co-neighborhoods considered above.

Interestingly, in contrast to the $O(\frac{1}{\varepsilon})$ bound on the size of triangle-free homomorphic image in Brandt-Thomassé's work [11] under the stronger minimum degree condition, much to our own surprise, the single exponential bound in $\frac{1}{\varepsilon}$ in Theorem 1.5 is optimal as shown by the following construction.

Theorem 1.6. *For every $r \geqslant 3$, there exists an n-vertex K_r-free graph G such that for every pair of non-adjacent vertices u, v, $G[N(u) \cap N(v)]$ contains at least εn^{r-2} copies of K_{r-2}, but G has no K_r-free homomorphic image of size smaller than $2^{\frac{1}{8r^r\varepsilon}}$.*

1.4 Applications

As a first application, the chromatic and homomorphism thresholds for cliques, $\delta_{\text{hom}}(K_r) = \delta_\chi(K_r) = \frac{2r-5}{2r-3}$, are corollaries of Theorems 1.1 and 1.5, and Theorem 1.5 quantitatively improves the double exponential bound in [38] to a single exponential one. Indeed, the minimum degree condition $\delta(G) \geqslant (\frac{2r-5}{2r-3}+\varepsilon)n$ implies the clique density condition in co-neighborhoods considered in Theorems 1.1 and 1.5, see [32, Propostion 4.9].

In addition, one can consider the following natural variations on δ_χ and δ_{hom} with increasingly stronger hypotheses. First define a 'higher moment minimum degree': for $a \in \mathbb{N}$, let $\delta^{(a)}(G)$ be the minimum co-degree over all independent sets of size a. So $\delta^{(1)}(G) = \delta(G)$ and $\delta^{(2)}(G) = \min\{|N(u) \cap N(v)| : uv \notin E(G)\}$. Then, we can define a higher moment homomorphism threshold as follows.

$$\delta^{(a)}_{\text{hom}}(H) := \inf\left\{\alpha \geqslant 0 : \exists\ H\text{-free } F = F(\alpha, H) \text{ s.t. } \forall\ n\text{-vertex } H\text{-free } G,\ \delta^{(a)}(G) \geqslant \alpha n \Rightarrow G \xrightarrow{\text{hom}} F\right\}.$$

More generally, for $a, b \in \mathbb{N}$ with $b \geqslant 2$, let $\hat{\delta}^{(a,b)}(G)$ be the minimum relative K_b-density in the subgraph induced by co-neighborhood of I over all independent sets I of size a. In other words,

$\hat{\delta}^{(a,b)}(G) = \min_{I:|I|=a} \frac{k_b(G[N(I)])}{\binom{|N(I)|}{b}}$, where $k_b(\cdot)$ counts the number of copies of K_b. Define

$$\delta^{(a,b)}_{\text{hom}}(H) := \inf\big\{\alpha \geqslant 0 : \exists\ H\text{-free graph } F = F(\alpha, H) \text{ s.t. } \forall\ n\text{-vertex } H\text{-free } G,$$
$$\delta^{(a)}(G) = \Omega(n),\ \hat{\delta}^{(a,b)}(G) \geqslant \alpha \Rightarrow G \xrightarrow{\text{hom}} F\big\}.$$

Thus, $\delta_{\text{hom}} = \delta^{(1)}_{\text{hom}}$ and Theorem 1.5 can be stated as $\delta^{(2,r-2)}_{\text{hom}}(K_r) = 0$.

We shall see that Theorem 1.5 determines all other variations $\delta^{(2,b)}_{\text{hom}}(K_r)$, which take values from generalized Turán densities: $\pi_s(K_t) := \lim_{n\to\infty} \frac{\text{ex}(n,K_s,K_t)}{\binom{n}{s}}$, where $s \leqslant t$. Here $\text{ex}(n, K_s, K_t)$ is the maximum number of K_s in an n-vertex K_t-free graph. Erdős [19] determined back in 1962 all such densities $\pi_s(K_t) = \prod_{i=1}^{s-1} \frac{t-1-i}{t-1}$, which are realized by $T_{n,t-1}$.

Corollary 1.7. *Let $r, a, b \in \mathbb{N}$ with $r \geqslant 3$, $a \geqslant 2$, and $2 \leqslant b \leqslant r-2$. Then we have*

$$\delta^{(a)}_{\text{hom}}(K_r) = \frac{r-3}{r-2} \quad \textit{and} \quad \delta^{(a,b)}_{\text{hom}}(K_r) = \pi_b(K_{r-2}) = \prod_{i=1}^{b-1} \frac{r-3-i}{r-3}.$$

The proofs of Theorem 1.4, Theorem 1.5 and Theorem 1.6 will appear in the full version of this paper [32].

Notations. For a vertex $u \in V(G)$, we will use $N_G(u)$ (or $N(u)$ if the subscript is clear) to denote the set of neighborhood of u. For a pair of vertices $u, v \in V(G)$, we use $N_G(u, v)$ (or $N(u, v)$ if the subscript is clear) to denote the set of common neighbors of u and v, that is, $N_G(u, v) = N_G(u) \cap N_G(v)$. For a subset $T \subseteq V(G)$, we will use $G[T]$ to denote the subgraph induced by T. For the sake of clarity of presentation, we omit floors and ceilings and treat large numbers as integers whenever this does not affect the argument.

Structure of this paper. In Section 2, we will build the connection between graph theory and geometry and show several geometric consequences. We discuss some related results in Section 3.

2 A geometric framework

In this section, first in Section 2.1, we introduce abstract convexity spaces and some useful parameters. Then in Section 2.2, we define a convexity space from a graph, prove some properties of this space, and establish some correspondence between the graph and this space. With this preparation, we are then in a position to lay out our geometric approach in Section 2.3 and list the main lemmas for proving Theorem 1.1. The main lemmas are proved in Section 2.4 and Section 2.5.

2.1 Abstract convexity space

A convexity space is a pair $(X, \mathcal{C})$ where $\mathcal{C} \subseteq 2^X$ is a family of subsets satisfying

- $\varnothing, X \in \mathcal{C}$;
- $\mathcal{C}$ is closed under intersection, i.e. for every $\mathcal{F} \subseteq \mathcal{C}$, $\bigcap \mathcal{F} := \bigcap_{F\in\mathcal{F}} F \in \mathcal{C}$.

We call sets in $\mathcal{C}$ convex sets. The convex hull of a set of points $Y \subseteq X$ is the intersection of all convex sets in $\mathcal{C}$ containing Y, which is also the minimal convex set containing Y. That is, $\text{conv}Y = \bigcap_{Y\subseteq F\in\mathcal{C}} F$. When the set X is clear from the content, we simply refer $\mathcal{C}$ as a convexity space. For more on the theory of convexity space and combinatorial convexity, we refer the interested readers to the books of van de Vel [46] and of Bárány [6]. Two trivial convexity spaces are $\mathcal{C} = \{\varnothing, X\}$ and $\mathcal{C} = 2^X$. Let $\mathcal{C}^d$ be the family of all convex sets in $\mathbb{R}^d$. Then $(\mathbb{R}^d, \mathcal{C}^d)$ is the usual Euclidean convexity space. Here are some non-trivial examples.

- Convex lattice sets: A set of the form $C \cap \mathbb{Z}^d$ for some convex set C in $\mathcal{C}^d$ is called a convex lattice set. Then $(\mathbb{Z}^d, \mathcal{C}^d \cap \mathbb{Z}^d)$ is a convexity space.
- Subgroups: Given a group G with identity e, then $(G\backslash\{e\}, \{H\backslash\{e\} : H \leqslant G\})$ is a convexity space. For a set $S \subseteq G\backslash\{e\}, \text{conv}S = \langle S\rangle\backslash\{e\}$ is a subgroup generated by S (with identity removed).
- Subcubes. A subset $\mathcal{C} \subseteq \{0,1\}^n$ is a subcube if there exists a set of coordinates $I \subseteq [n]$ and a binary vector $\boldsymbol{v} \in \{0,1\}^I$ such that $\mathcal{C}$ consists of all vectors in $\{0,1\}^n$ whose projection on I is $\boldsymbol{v}$. Then $\{0,1\}^n$ together with the family of all subcubes forms a convexity space.
- Subtrees. Given a finite tree T on vertex set V, then V together with all its subtrees forms a convexity space.

Let us now introduce some invariants which are abstractions of properties of Euclidean convexity space.

Definition 2.1. The Radon number of a convexity space $(X, \mathcal{C})$, denoted by $r(\mathcal{C})$, is the minimum integer r such that for any set of points $Y \subseteq X$ with $|Y| \geqslant r$, there is a partition $Y = Y_1 \cup Y_2$ with $\text{conv}Y_1 \cap \text{conv}Y_2 \neq \varnothing$.

Radon's Lemma [39] states, in this notation, that the Radon number of the Euclidean convexity space $(\mathbb{R}^d, \mathcal{C}^d)$ is at most $d+2$.

Definition 2.2. The Helly number of a set system $\mathcal{F}$, denoted by $h(\mathcal{F})$, is the minimum integer h such that in any finite subfamily $\mathcal{G} \subseteq \mathcal{F}$, if every h-tuple of $\mathcal{G}$ is intersecting, then $\mathcal{G}$ is intersecting.

Thus, Helly's theorem [26] implies that $h(\mathcal{C}^d) \leqslant d+1$, which is less than its Radon number. This relation between Radon and Helly numbers holds for all convexity space $(X, \mathcal{C})$: $h(\mathcal{C}) < r(\mathcal{C})$, as shown by Levi [31]. The gap between Radon number and Helly number, however, could be arbitrarily large. One such example is the space of subcubes in $\{0,1\}^n$, which has Helly number 2, but Radon number $\lfloor \log_2(n+1) \rfloor + 1$.

We also need the fractional extension of Helly number defined as follows.

Definition 2.3. The fractional Helly number of a set system $\mathcal{F}$, denoted by $h^*(\mathcal{F})$, is the smallest natural number k such that for every $\alpha > 0$, there exists $\beta = \beta(\alpha) > 0$ such that the following holds. For every $\{F_1, F_2, \ldots, F_m\} \subseteq \mathcal{F}$, if the number of intersecting k-tuples $I \in \binom{[m]}{k}$ with $\bigcap_{i \in I} F_i \neq \varnothing$ is at least $\alpha\binom{m}{k}$, then $\{F_1, F_2, \ldots, F_m\}$ contains an intersecting subfamily of size at least βm.

We remark that there is no direct relation between the Helly number and its fractional counterpart. For example, for the Euclidean convexity space, $h(\mathcal{C}^d) = h^*(\mathcal{C}^d) = d+1$; for the convex lattice set, $h(\mathcal{C}^d \cap \mathbb{Z}^d) = 2^d$ and $h^*(\mathcal{C}^d \cap \mathbb{Z}^d) = d+1$ [7, 14]; for the family $\mathcal{B}^d$ of all axis-aligned boxes in $\mathbb{R}^d$, $h(\mathcal{B}^d) = 2$ and $h^*(\mathcal{B}^d) = d+1$ [16].

The concept of (weak) ε-net is well-studied in computational geometry, combinatorics, and machine learning; it is particularly useful in many algorithmic applications, including range searching and geometric optimization.

Definition 2.4. Let X be a set and $\mathcal{C} \subseteq 2^X$ be a family of subsets of X. Given a finitely supported probability measure μ on X, a set $N \subseteq X$ is a weak ε-net for $\mathcal{C}$ with respect to μ, if $N \cap C \neq \varnothing$ for every $C \in \mathcal{C}$ with $\mu(C) \geqslant \varepsilon$. We say that $\mathcal{C}$ has weak ε-nets of size $m = m(\mathcal{C}, \varepsilon)$ if there is a weak ε-net for $\mathcal{C}$ with respect to μ of size at most m for every finitely supported probability measure μ on X.

Note that $m(\mathcal{C}, \varepsilon)$, the size of weak ε-nets for $\mathcal{C}$, does not depend on the choice of the probability measure μ.

We need the following fractional version of transversal number.

Definition 2.5. A fractional transversal for a set system $\mathcal{G}$ on ground set V is a function f from V to $[0,1]$ such that for every set $A \in \mathcal{G}$, $\sum_{v\in A} f(v) \geqslant 1$. The size of f is $\sum_{v\in V} f(v)$ and the fractional transversal number of $\mathcal{G}$ is the minimum size of a fractional transversal for $\mathcal{G}$, denoted by $\tau^*(\mathcal{G})$.

In particular, the characteristic function of a transversal is a fractional transversal and so for any $\mathcal{G}$, we have $\tau^*(\mathcal{G}) \leqslant \tau(\mathcal{G})$.

Definition 2.6. A matching in a set system $\mathcal{H}$ is a collection pairwise disjoint sets in $\mathcal{H}$. The matching number of $\mathcal{H}$ is size of a largest matching in $\mathcal{H}$, denoted by $\nu(\mathcal{H})$.

Definition 2.7. A fractional matching in a set system $\mathcal{H}$ on ground set V is a function g from $\mathcal{H}$ to $[0,1]$ such that for every $v \in V$, $\sum_{v\in A} g(A) \leqslant 1$. The size of g is $\sum_{A\in\mathcal{H}} g(A)$ and the fractional matching number of $\mathcal{H}$ is the maximum size of a fractional matching in $\mathcal{H}$, denoted by $\nu^*(\mathcal{H})$.

Similarly, $\nu(\mathcal{H}) \leqslant \nu^*(\mathcal{H})$ holds for every set system $\mathcal{H}$. Moreover, LP duality infers that $\nu^*(\mathcal{H}) = \tau^*(\mathcal{H})$. Therefore, for any $\mathcal{H}$, we have $\nu(\mathcal{H}) \leqslant \nu^*(\mathcal{H}) = \tau^*(\mathcal{H}) \leqslant \tau(\mathcal{H})$.

The Vapnik-Chervonenkis dimension (VC-dimension for short) is a parameter that measures the complexity of various combinatorial objects, and plays an important role in statistics, algebraic geometry, learning theory, and model theory. For a set system $\mathcal{F} \subseteq 2^X$, the Vapnik-Chervonenkis dimension of $\mathcal{F}$, denoted by $\dim_{\mathsf{VC}}(\mathcal{F})$, is the largest integer d for which there exists a subset $S \subseteq X$ with $|S| = d$ such that for every subset $B \subseteq S$, one can find a member $A \in \mathcal{F}$ with $A \cap S = B$. In such case, we say that S is shattered by $\mathcal{F}$.

The well-known ε-net theorem of Haussler and Welzl [25] (also see the book of Matoušek [35]) provides an inverse inequality for transversal number and its fractional version in set systems with bounded VC-dimension.

Theorem 2.8 (ε-net theorem)**.** *There exists an absolute constant C such that the following holds. Let $\mathcal{F}$ be a set system with VC-dimension d. Then $\tau(\mathcal{F}) \leqslant Cd\tau^*(\mathcal{F})\log\tau^*(\mathcal{F})$.*

2.2 The unusual suspect

In this subsection, we construct an abstract convexity space from a graph. Given a set system $\mathcal{F}$, the intersection closure of $\mathcal{F}$, denoted by $\mathcal{F}^\cap$, is the set system obtained from $\mathcal{F}$ by taking all possible intersections of subfamilies of $\mathcal{F}$. Given a graph G, let $\mathrm{MIS}(G)$ be the family of all maximal independent sets of G.

We can now define our convexity space $\mathcal{C}(G)$ from a given graph G as follows:

- Let $\mathcal{B}(G) := \{K_v : v \in V(G)\}$ be the set system on the ground set $\mathrm{MIS}(G)$ indexed by $V(G)$, where $K_v = \{I \in \mathrm{MIS}(G) : v \in I\}$ consists of all maximal independent sets of G containing v.
- Let $\mathcal{C}(G) = \mathcal{B}(G)^\cap$ be the intersection closure of $\mathcal{B}(G)$. Then $(X, \mathcal{C}(G))$ is a convexity space, where $X = \mathrm{MIS}(G)$.

Note that $\mathcal{B}(G)$ and $\mathcal{C}(G)$ could be a family of multi-sets. For example, if two vertices u, v have the same neighborhood in G, then $K_u = K_v$. Equivalently, $\mathcal{C}(G)$ can be defined as $\mathcal{C}(G) = \{K_S : S \subseteq V(G)\}$, where for each subset $S \subseteq V(G)$, $K_S = \{I \in X : S \subseteq I\}$ is the set of all maximal independent sets containing S. The dual of a set system $\mathcal{H}$ is a set system obtained from $\mathcal{H}$ by swapping the roles of ground elements and sets in $\mathcal{H}$. Note that the dual of $\mathcal{B}(G)$ is the set system induced by $\mathrm{MIS}(G)$, which we denote by $\mathcal{M}(G) = \{I : I \in \mathrm{MIS}(G)\}$.

The convex hull operator of this convexity space $(X, \mathcal{C}(G))$ can be described as follows. Given a subset $Y = \{I_1, \ldots, I_m\} \subseteq X$, $\mathrm{conv}Y$ is the intersection of all convex sets in $\mathcal{C}(G)$ containing Y, so $\mathrm{conv}Y = K_{I_1\cap\cdots\cap I_m}$.

The graph G	**Set syetem $\mathcal{B} = \mathcal{B}(G)$**
$uv \in E(G)$	$K_u \cap K_v = \varnothing$
(maximal) independent set I	(maximal) intersecting subfamily $\{K_v : v \in I\}$
Chromatic number $\chi(G)$	Transversal number $\tau(\mathcal{B})$
Clique number $w(G)$	Matching number $\nu(\mathcal{B})$
K_r-freeness	$(r, 2)$-property

Table 1: Graph terminology vs. geometric terminology

In the rest of this subsection, we establish some correspondence between a graph G and the set system $\mathcal{B}(G)$, which is summarized in Table 1.

We first observe that G is isomorphic to the disjointness graph of $\mathcal{B}(G)$. Here, the disjointness graph of a set system $\mathcal{F}$, denoted by $D(\mathcal{F})$, is the graph with vertex set $\mathcal{F}$ and a pair of sets in $\mathcal{F}$ are adjacent in $D(\mathcal{F})$ if and only if they are disjoint. On the other hand, if we start with any set system $\mathcal{F}$ and consider $\mathcal{B}(D(\mathcal{F}))$, we will get a set system whose nerves have the same 1-skeleton. The nerve of a set system is an abstract simplicial complex that records all intersecting subfamilies. In other words, $\mathcal{F}$ and $\mathcal{B}(D(\mathcal{F}))$ have the same pairwise intersections. One can think the two operations $D(\cdot)$ and $\mathcal{B}(\cdot)$ as dual of each other.

Proposition 2.9. *Given any graph G, we have $G \cong D(\mathcal{B}(G))$. On the other hand, given any set system $\mathcal{F}$, $\mathcal{F}$ and $\mathcal{B}(D(\mathcal{F}))$ have the same pairwise intersections.*

Proof. By the definition of $\mathcal{B}(G)$, $K_u \cap K_v = \varnothing$ if and only if there is no maximal independent set containing both u and v, in other words $uv \in E(G)$. Thus, mapping K_v to v for each $v \in V(G)$ is a graph isomorphism between $D(\mathcal{B}(G))$ and G.

For the second part, we need to show that $\mathcal{F}$ and $\mathcal{B}(D(\mathcal{F}))$ have the same pairwise intersections. This amounts to proving that $D(\mathcal{F}) \cong D(\mathcal{B}(D(\mathcal{F})))$, which follows from the first part. □

A simple but useful fact is that the operation: $G \to \mathcal{C}(G)$ always produces a convexity space with Helly number 2.

Proposition 2.10. *For any graph G with at least one edge, $\mathcal{C}(G)$ has Helly number 2.*

Proof. We need to show that for any subfamily $\mathcal{S} = \{K_{S_i} : S_i \subseteq V(G)\} \subseteq \mathcal{C}(G)$, if $\mathcal{S}$ is pairwise intersecting, then it is in fact intersecting. Similarly to Proposition 2.9, it is easy to see that if K_S and K_T intersect, then $S \cup T$ is an independent set. Thus, $\mathcal{S}$ being pairwise intersecting infers that $S = \bigcup\limits_{S_i \in \mathcal{S}} S_i$ is an independent set in G. Let I be an arbitrary maximal independent set containing S, then $I \in \bigcap \mathcal{S}$ as desired. □

Proposition 2.9 in particular implies that there is a one-to-one correspondence between (maximal) independent sets I in G and subfamilies of $\mathcal{B}(G)$ that are (maximally) pairwise intersecting; and as $\mathcal{B}(G)$ has Helly number 2, all these subfamilies are intersecting.

A key correspondence is as follows.

Proposition 2.11. *For any graph G, we have $\chi(G) = \tau(\mathcal{B}(G))$.*

Proof. Observe that a collection of maximal independent sets pierces $\mathcal{B}(G) = \{K_u : u \in V(G)\}$ if and only if their union covers $V(G)$. We first show that $\chi(G) \leqslant \tau(\mathcal{B}(G))$. Let $k = \tau(\mathcal{B}(G))$, then there are $I_1, I_2, \ldots, I_k \in \mathrm{MIS}(G)$ such that they pierce $\mathcal{B}(G)$. As $\bigcup_{j \in [k]} I_j = V(G)$, for every vertex $v \in V(G)$, we can color it with the smallest index $j \in [k]$ such that $v \in I_j$. This provides a proper k-coloring of G, and so $\chi(G) \leqslant k$.

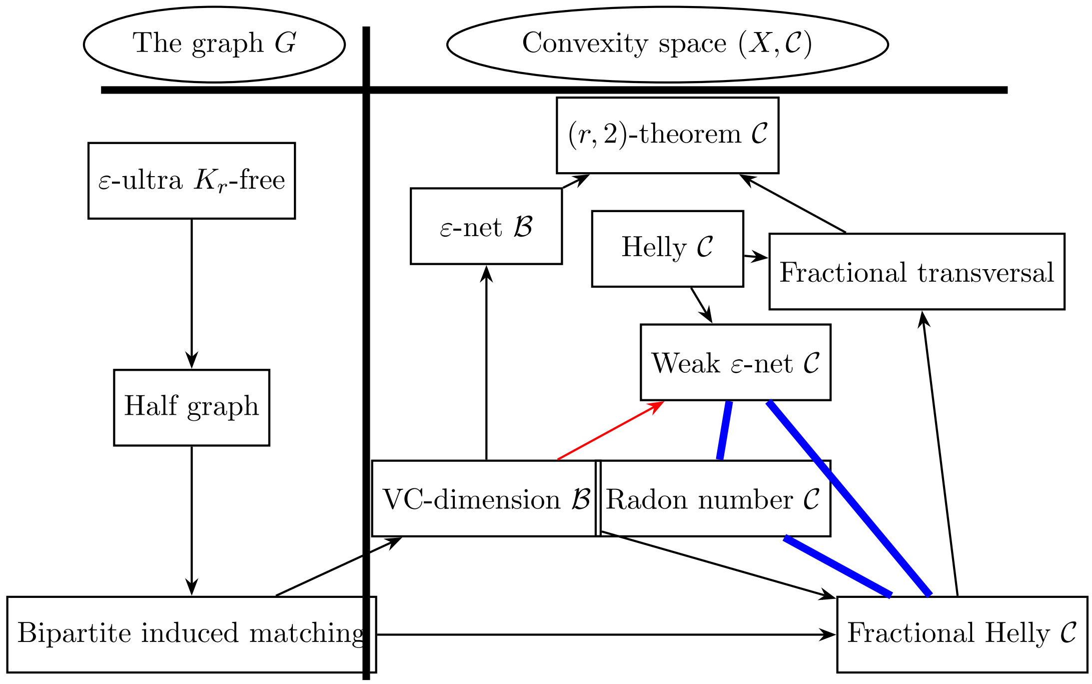

Figure 2.1: The relationship between graphs and convexity spaces

To show $\tau(\mathcal{B}(G)) \leqslant \chi(G)$, suppose that $\chi(G) = r$. Then we can partition $V(G)$ into r independent sets $V_1 \cup V_2 \cup \cdots \cup V_r$. For each V_j, $j \in [r]$, let $I_j \supseteq V_j$ be an arbitrary maximal independent set containing it. Then we have $\bigcup_{j\in[r]} I_j = V(G)$ and therefore I_j, $j \in [r]$, pierce $\mathcal{B}(G)$, which implies that $\tau(\mathcal{B}(G)) \leqslant r$. $\square$

Remark 2.12. Using the correspondence in Table 1, it is not hard to see that Theorem 1.1 is equivalent to Theorem 1.3, which is an $(r,2)$-theorem for $\mathcal{B}(G)$. A class of graphs are χ-bounded if its chromatic number can be bounded by a function of its clique number. The correspondence in Table 1 shows that Theorem 1.1, the $(r,2)$-theorem for $\mathcal{B}(G)$, can be rephrased as the statement that ultra maximal K_r-free graphs are χ-bounded. Finally note that [32, Theorem 1.4] is a strengthening as it is an $(r,2)$-theorem for the convexity space $\mathcal{C}(G)$, because the $(r,2)$-theorem holds for any family of convex sets in the convexity space, not just the ones of the family $\mathcal{B}(G)$.

2.3 Main lemmas and proof of Theorem 1.1

In this subsection, we present our geometric framework, depicted in Fig. 2.1. By Proposition 2.11, our goal is to bound the transversal number of $\mathcal{B}(G)$. Our approach consists of two sides.

On the side of convexity space, to bound $\tau(\mathcal{B}(G))$, by the ε-net theorem, Theorem 2.8, it suffices to bound the VC-dimension of $\mathcal{B}(G)$ and the fractional transversal number $\tau^*(\mathcal{B}(G))$. The latter can in turn be bounded by the fractional Helly number of $\mathcal{B}(G)$. To this end, we prove in Lemma 2.13 and Lemma 2.14 that both the VC-dimension and the fractional Helly number of $\mathcal{B}(G)$ can be controlled by certain induced matching defined as follows.

A matching $\{u_iv_i\}_{i\in[t]}$ in a graph G is a <u>bipartite induced matching</u> of size t if $u_iv_j \in E(G)$ if and only if $i = j$. We call the size of a largest such matching in G its <u>bipartite induced matching number</u>, denoted by $\nu_{\mathrm{bi}}(G)$. If we further require that both $\{u_i\}_{i\in[t]}$ and $\{v_i\}_{i\in[t]}$ are independent sets, then we call it an <u>induced matching</u>.

Lemma 2.13. *For any graph G, if* $\dim_{\mathsf{VC}}(\mathcal{B}(G)) \geqslant 3$, *then* $\dim_{\mathsf{VC}}(\mathcal{B}(G)) \leqslant \nu_{\mathrm{bi}}(G)$.

We remark that the bound on the VC-dimension above is optimal. Indeed, consider the graph G with four vertices x, y, z, w, where x, y, z form a copy of triangle, and w is an isolated vertex. Then it is easy to check that $\mathcal{B}(G)$ has VC-dimension 2, however $\nu_{\text{bi}}(G) = 1$.

Lemma 2.14. *For any graph G, $h^*(\mathcal{B}(G)) \leqslant \nu_{\text{bi}}(G) + 1$.*

We have then to control the bipartite induced matching number on the graph side. Recall that an n-vertex K_r-free graph G is ε-ultra maximal K_r-free if for every non-adjacent pair of vertices $u, v \in V(G)$, the induced subgraph $G[N(u) \cap N(v)]$ contains εn^{r-2} copies of K_{r-2}.

We show that all ε-ultra maximal K_r-free graphs have bounded bipartite induced matching number.

Theorem 2.15. *For $\varepsilon > 0$ and $r \geqslant 3$, every ε-ultra maximal K_r-free graph G satisfies $\nu_{\text{bi}}(G) \leqslant (2/\varepsilon)^{2/\varepsilon}$.*

We will prove Theorem 2.15 in two steps. First, we show that in an ε-ultra K_r-free graph G, there does not exist a large copy of certain half graph, see Lemma 2.22. From not containing a large half graph, we can then derive that G cannot have a large bipartite induced matching, see Lemma 2.23.

Combining Lemmas 2.13 and 2.14 and Theorem 2.15, we get the following.

Corollary 2.16. *Let $\varepsilon > 0$, $r \geqslant 3$ and G be an ε-ultra maximal K_r-free graph. Then both the VC-dimension and fractional Helly number of $\mathcal{B}(G)$ are at most $(2/\varepsilon)^{2/\varepsilon} + 1$.*

Remark 2.17. We would like to mention a related (p, q)-theorem for set systems with bounded VC-dimension due to Matoušek [36]. It states that if the VC-dimension of the dual of a set system $\mathcal{F}$ is at most $k - 1$, then any finite $\mathcal{G} \subseteq \mathcal{F}$ with (r, k)-property, $r \geqslant k$, satisfies $\tau(\mathcal{G}) = O_{r,k}(1)$. Notice that even though we do have a bound on the VC-dimension of dual of $\mathcal{B}(G)$, that is, $\dim_{\mathsf{VC}}(\mathcal{M}(G))$ (from Corollary 2.16 or Lemma 2.20). However, we cannot apply Matoušek's result to bound $\tau(\mathcal{B}(G))$ directly. Indeed, by Table 1, the K_r-freeness of G corresponds to $(r, 2)$-property of $\mathcal{B}(G)$. To utilize Matoušek's result, we would need to (i) either show the (r, k)-property for $\mathcal{B}(G)$, (ii) or prove that $\mathcal{M}(G)$ has VC-dimension 1. Having (i) for any $k > 2$ would require a stronger hypothesis that in G, there is an independent set of size k in a subgraph induced by any set of r vertices. We cannot hope to have (ii) either: for ε-ultra maximal K_r-free graph G, $\mathcal{M}(G)$ could have VC-dimension $2^{\Omega(1/\varepsilon)}$, see e.g. the construction in Theorem 1.5 and [32, Theorem 1.7].

To prove Theorem 1.1, we are left to bound the fractional transversal number of $\mathcal{B}(G)$.

Lemma 2.18. *Let $\varepsilon > 0$, $r \geqslant 3$ and G be an ε-ultra maximal K_r-free graph. Then $\tau^*(\mathcal{B}(G)) = O_{\varepsilon,r}(1)$.*

Before we prove this result, we need the following.

Theorem 2.19 ([2])**.** *Let $r, \ell \geqslant k \geqslant 2$ and $\mathcal{F}$ be a set system with (r, k)-property. If the Helly number of $\mathcal{F}$ is k and the fractional Helly number of $\mathcal{F}$ is ℓ, then $\tau^*(\mathcal{F}) = O_{r,\ell,k}(1)$.*

Proof of Lemma 2.18. By the correspondence in Table 1, $\mathcal{B}(G)$ has $(r, 2)$-property as G is K_r-free. The Helly number of $\mathcal{B}(G)$ is 2 by Proposition 2.10 and the fractional Helly number of $\mathcal{B}(G)$ is also bounded by Corollary 2.16. The conclusion then follows from Theorem 2.19 with $k = 2$. □

Finally, combining Theorem 2.8, Corollary 2.16 and Lemma 2.18, we obtain Theorem 1.1.

2.4 VC-dimension and fractional Helly number

In this subsection, we prove Lemmas 2.13 and 2.14.

Proof of Lemma 2.13. Suppose the VC-dimension of $\mathcal{B}(G)$ is $d \geqslant 3$, then there is a set $A := \{I_1, \dots, I_d\}$ of d maximal independent sets shattered by $\mathcal{B}(G)$. We shall find a bipartite induced matching of size d in G. Note that for each $i \in [d]$, there exists some $K_{v_i} \in \mathcal{B}(G)$ such that $K_{v_i} \cap A = A \backslash \{I_i\}$. Then $v_i \notin I_i$, which together with the maximality of I_i implies that there is a vertex $u_i \in I_i$ such that $v_i u_i \in E(G)$. Moreover, for all distinct $i, j \in [d]$, $v_i \in I_j$, implying that $v_i u_j \notin E(G)$. That is, v_i is adjacent to u_j if and only if $i = j$. In particular, all v_i are distinct vertices and all u_i are distinct vertices.

We now show that for any $i, j \in [d]$, $v_i \neq u_j$. Suppose $v_i = u_j$. Then $i \neq j$, for otherwise $v_j = v_i = u_j$ contradicting $u_j v_j \in E(G)$. Note also that for any distinct $i, j \in [d]$, $K_{v_i} \cap K_{v_j} \cap A = A \backslash \{I_i, I_j\} \neq \varnothing$ as $d \geqslant 3$, implying that $u_j v_j = v_i v_j \notin E(G)$, a contradiction. Thus, $\{u_i v_i\}_{i \in [d]}$ is a bipartite induced matching of size d as desired. □

We can also bound the VC-dimension of the dual of $\mathcal{B}(G)$ by its bipartite induced matching number.

Lemma 2.20. *For any graph G, if the VC-dimension of the dual of $\mathcal{B}(G)$ is $d \geqslant 1$, then $d \leqslant \nu_{bi}(G)$.*

Proof. Recall that the dual of $\mathcal{B}(G)$ is the maximal independent set hypergraph $\mathcal{M} = \mathcal{M}(G)$. By assumption, there is a set D of d vertices $v_1, \dots, v_d$ such that for any subset $D' \subseteq D$, there is a maximal independent set $I_{D'}$ such that $I_{D'} \cap D = D'$. In particular, D is an independent set by considering I_D. For each $i \in [d]$, as $v_i \notin I_{D \backslash \{v_i\}}$, the maximality of $I_{D \backslash \{v_i\}}$ implies that there exists some vertex $u_i \in I_{D \backslash \{v_i\}}$ such that $v_i u_i \in E(G)$ and $v_j u_i \notin E(G)$ for any $j \neq i$. Since D is an independent set, we must have $u_i \notin D$. Moreover, $u_1, u_2, \dots, u_d$ are all distinct. Indeed, if $u_j = u_i$ for $j \neq i$, then $v_j u_j = v_j u_i \notin E(G)$, a contradiction. Thus, $\{v_i u_i\}_{i \in [d]}$ is a bipartite induced matching of size d. □

We now bound the fractional Helly number of $\mathcal{B}(G)$. Matoušek [36] proved that if the VC-dimension of the dual of a set system $\mathcal{F}$ is at most $k-1$, then the fractional Helly number of $\mathcal{F}$ is at most k. This, together with Lemma 2.20, implies that $h^*(\mathcal{B}(G)) \leqslant \dim_{\mathsf{VC}}(\mathcal{M}(G)) + 1 \leqslant \nu_{\mathrm{bi}}(G) + 1$. Here, we give a different proof using the following result of Holmsen [27].

Theorem 2.21 ([27])**.** *For any $k \geqslant 2$ and $\alpha > 0$, there exists $\beta > 0$ such that the following holds. Let G be an n-vertex graph with at least $\alpha\binom{n}{k}$ independent sets of size k. If G does not contain an induced matching of size k, then $\alpha(G) \geqslant \beta n$.*

Proof of Lemma 2.14. Let $\nu_{\mathrm{bi}}(G) = k - 1$. By the definition of fractional Helly number, we need to show that for every $\alpha > 0$, there exists $\beta > 0$ such that every family $\mathcal{F} = \{K_{u_1}, \dots, K_{u_m}\} \subseteq \mathcal{B}(G)$ with at least $\alpha\binom{m}{k}$ intersecting k-subsets must contain an intersecting subfamily of size βm. By the correspondence in Table 1, this amounts to proving that for every $U = \{u_1, \dots, u_m\} \subseteq V(G)$, if $G[U]$ has at least $\alpha\binom{m}{k}$ independent sets of size k, then $\alpha(G[U]) \geqslant \beta m$. This follows immediately from Theorem 2.21 as G does not contain an induced matching of size $\nu_{\mathrm{bi}}(G) + 1 = k$. □

2.5 Bipartite induced matching in ε-ultra maximal K_r-free

In this subsection, we prove Theorem 2.15. We need a notion of half graphs. Let $\mathcal{H}_k$ be the family of graphs which consists of all of the graphs H with vertices $\{x_i, y_i : i \in [k]\}$ such that

- $x_i y_i \notin E(H)$ for all $i \in [k]$; and
- $x_i y_j \in E(H)$ for all $1 \leqslant j < i \leqslant k$.

Note that we put no restriction on pairs $\{x_j, y_i\}$ for all $1 \leqslant j < i \leqslant k$ and no restriction on $H[\{x_1, \ldots, x_k\}]$ and $H[\{y_1, \ldots, y_k\}]$. We call graphs in $\mathcal{H}_k$ half graphs.

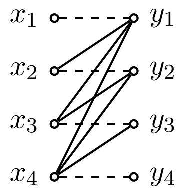

Figure 2.2: Half graph

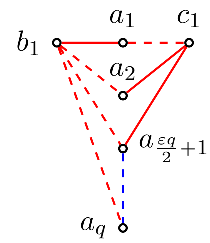

Figure 2.3: Claim 2.24

Lemma 2.22. *Let $\varepsilon > 0$, $r \geqslant 3$ and G be an ε-ultra maximal K_r-free graph. Then G does not contain any half graph in $\mathcal{H}_k$ as a subgraph, where $k = \frac{1}{\varepsilon} + 1$.*

Proof. Suppose to the contrary that G contains a copy of some half graph $H \in \mathcal{H}_k$ with vertex set $\{x_i, y_i : i \in [k]\}$. As G is ε-ultra maximal K_r-free and $x_i y_i \notin E(G)$ for all $i \in [k]$, the induced subgraph $G[N(x_i, y_i)]$ contains at least εn^{r-2} many $(r-2)$-cliques. By the pigeonhole principle, there must exist a copy of K_{r-2} lying in at least $\lceil \frac{k\varepsilon n^{r-2}}{\binom{n}{r-2}} \rceil \geqslant \lceil k\varepsilon \rceil \geqslant 2$ common neighborhood of distinct pairs, say $G[N(x_i, y_i)]$ and $G[N(x_j, y_j)]$ with $1 \leqslant j < i \leqslant k$. Then, as $x_i y_j \in E(G)$, this copy of K_{r-2} together with x_i and y_j forms a copy of K_r, a contradiction. □

Lemma 2.23. *Let $\varepsilon > 0$, $r \geqslant 3$ and G be an ε-ultra maximal K_r-free graph. If G contains a bipartite induced matching of size t, then G contains a copy of some graph $H \in \mathcal{H}_k$ as a subgraph, where $k = \log t / \log \frac{2}{\varepsilon}$.*

Proof. We first establish the following claim, which will be used iteratively to build a large half graph from a large bipartite induced matching.

Claim 2.24. *Let $a_1, \ldots, a_q$ and b_1 be vertices with $q \geqslant 2/\varepsilon$, $a_1 b_1 \in E(G)$ and $a_i b_1 \notin E(G)$ for all $2 \leqslant i \leqslant q$. Then there exists a vertex c_1 such that $a_1 c_1 \notin E(G)$ and $a_i c_1 \in E(G)$ for $2 \leqslant i \leqslant \varepsilon q/2 + 1$.*

Proof of claim. Consider the set of non-adjacent pairs $\{a_i, b_1\}$, $2 \leqslant i \leqslant q$. By assumption, for each such non-adjacent pair there are at least εn^{r-2} many $(r-2)$-cliques in $G[N(a_i, b_1)]$. Then, by the pigeonhole principle, there exists a copy K of $(r-2)$-clique lying in at least $\frac{(q-1)\varepsilon n^{r-2}}{\binom{n}{r-2}} \geqslant \varepsilon q/2$ many common neighborhood of distinct pairs $\{a_i, b_1\}$. By relabelling if necessary, we may assume that K lies in $G[N(a_i, b_1)]$ for $2 \leqslant i \leqslant \varepsilon q/2 + 1$. Note that, the vertex a_1 cannot be adjacent to all vertices in K for otherwise K together with a_1 and b_1 would form a copy of K_r. Therefore, we can pick a vertex c_1 in K such that $a_1 c_1 \notin E(G)$ and $a_i c_1 \in E(G)$ for $2 \leqslant i \leqslant \varepsilon q/2 + 1$ as desired. ■

We now build a large half graph. Let $\{x_i z_i\}_{i \in [t]}$ be a bipartite induced matching in G. In the first round, applying Claim 2.24 with $(a_i, q, b_1) = (x_i, t, z_1)$, we obtain a vertex y_1 such that $x_1 y_1 \notin E(G)$ and $x_i y_1 \in E(G)$ for $2 \leqslant i \leqslant \varepsilon t/2 + 1$. In general, for $2 \leqslant j \leqslant k$, in the j-th round, we apply Claim 2.24 with $(a_i, q, b_1) = (x_{i+j-1}, (\frac{\varepsilon}{2})^{j-1} t, z_j)$, we obtain a vertex y_j such that $x_j y_j \notin E(G)$ and $x_i y_j \in E(G)$ for $j + 1 \leqslant i \leqslant (\frac{\varepsilon}{2})^j t + j$. Since $t \geqslant (\frac{2}{\varepsilon})^k$, we can indeed invoke Claim 2.24 for k rounds to obtain vertices $y_1, \ldots, y_k$. Note that all y_i vertices are distinct as for any $1 \leqslant j < i \leqslant k$, $x_i y_j \in E(G)$ but $x_i y_i \notin E(G)$. Then, $\{x_i, y_i : i \in [k]\}$ induces a half graph in $\mathcal{H}_k$ with $k = \log t / \log \frac{2}{\varepsilon}$ as desired. □

Theorem 2.15 now follows immediately from Lemmas 2.22 and 2.23.

3 Concluding remarks

In the full version of paper [32], we shall prove an $(r, 2)$-theorem for $\mathcal{C}$ instead of just $\mathcal{B}(G)$. We do so by utilizing the equivalence of (i) having bounded Radon number, (ii) having a weak ε-net theorem, and (iii) having bounded fractionally Helly number for a convexity space. Notice that an $(r, 2)$-theorem in the convexity space $\mathcal{C}$ is a stronger statement as the $(r, 2)$-theorem holds for any family of convex sets in the convexity space, not just the ones of the family $\mathcal{B}(G)$.

References

[1] P. Allen, J. Böttcher, S. Griffiths, Y. Kohayakawa, and R. Morris. The chromatic thresholds of graphs. Adv. Math., 235:261–295, 2013.

[2] N. Alon, G. Kalai, J. Matoušek, and R. Meshulam. Transversal numbers for hypergraphs arising in geometry. Adv. in Appl. Math., 29(1):79–101, 2002.

[3] N. Alon and D. J. Kleitman. Piercing convex sets and the Hadwiger-Debrunner (p, q)-problem. Adv. Math., 96(1):103–112, 1992.

[4] N. Amenta, J. A. De Loera, and P. Soberón. Helly's theorem: new variations and applications. In Algebraic and geometric methods in discrete mathematics. AMS special session on algebraic and geometric methods in applied discrete mathematics, San Antonio, TX, USA, January 11, 2015. Proceedings, pages 55–95. Providence, RI: American Mathematical Society (AMS), 2017.

[5] B. Andrásfai, P. Erdős, and V. T. Sós. On the connection between chromatic number, maximal clique and minimal degree of a graph. Discrete Math., 8:205–218, 1974.

[6] I. Bárány. Combinatorial convexity, volume 77 of University Lecture Series. American Mathematical Society, Providence, RI, [2021] ©2021.

[7] I. Bárány and J. Matoušek. A fractional Helly theorem for convex lattice sets. Adv. Math., 174(2):227–235, 2003.

[8] I. Bárány and P. Soberón. Tverberg's theorem is 50 years old: a survey. Bull. Amer. Math. Soc. (N.S.), 55(4):459–492, 2018.

[9] P. V. M. Blagojević and G. M. Ziegler. Beyond the Borsuk-Ulam theorem: the topological Tverberg story. In A journey through discrete mathematics, pages 273–341. Springer, Cham, 2017.

[10] S. Brandt. On the structure of dense triangle-free graphs. Combin. Probab. Comput., 8(3):237–245, 1999.

[11] S. Brandt and S. Thomassé. Dense triangle-free graphs are four-colorable: A solution to the Erdős-Simonovits problem. preprint, 2011.

[12] C. C. Chen, G. P. Jin, and K. M. Koh. Triangle-free graphs with large degree. Combin. Probab. Comput., 6(4):381–396, 1997.

[13] J. A. De Loera, X. Goaoc, F. Meunier, and N. H. Mustafa. The discrete yet ubiquitous theorems of Carathéodory, Helly, Sperner, Tucker, and Tverberg. Bull. Amer. Math. Soc. (N.S.), 56(3):415–511, 2019.

[14] J.-P. Doignon. Convexity in cristallographical lattices. J. Geom., 3:71–85, 1973.

[15] O. Ebsen and M. Schacht. Homomorphism thresholds for odd cycles. Combinatorica, 40(1):39–62, 2020.

[16] J. Eckhoff. Intersection properties of boxes. I. An upper-bound theorem. Israel J. Math., 62(3):283–301, 1988.

[17] J. Eckhoff. Helly, Radon, and Carathéodory type theorems. In Handbook of convex geometry, Vol. A, B, pages 389–448. North-Holland, Amsterdam, 1993.

[18] P. Erdős. Graph theory and probability. Canadian J. Math., 11:34–38, 1959.

[19] P. Erdős. On the number of complete subgraphs contained in certain graphs. Magyar Tud. Akad. Mat. Kutató Int. Közl., 7:459–464, 1962.

[20] P. Erdős and M. Simonovits. On a valence problem in extremal graph theory. Discrete Math., 5:323–334, 1973.

[21] P. Erdős. Some recent results on extremal problems in graph theory. (Results). Theory Graphs, Int. Symp. Rome 1966, 117-123 (English), 124-130 (French) (1967)., 1967.

[22] W. Goddard and J. Lyle. Dense graphs with small clique number. J. Graph Theory, 66(4):319–331, 2011.

[23] H. Hadwiger and H. Debrunner. Über eine Variante zum Hellyschen Satz. Arch. Math. (Basel), 8:309–313, 1957.

[24] R. Häggkvist. Odd cycles of specified length in nonbipartite graphs. In Graph theory (Cambridge, 1981), North-Holland Math. Stud., 62, pages 89–99. 1982.

[25] D. Haussler and E. Welzl. ϵ-nets and simplex range queries. Discrete Comput. Geom., 2(2):127–151, 1987.

[26] E. Helly. Über Mengen konvexer Körper mit gemeinschaftlichen Punkten. Jahresber. Dtsch. Math.-Ver., 32:175–176, 1923.

[27] A. F. Holmsen. Large cliques in hypergraphs with forbidden substructures. Combinatorica, 40(4):527–537, 2020.

[28] G. P. Jin. Triangle-free four-chromatic graphs. Discrete Math., 145(1-3):151–170, 1995.

[29] M. Katchalski and A. Liu. A problem of geometry in $\mathbb{R}^n$. Proc. Amer. Math. Soc., 75(2):284–288, 1979.

[30] S. Letzter and R. Snyder. The homomorphism threshold of $\{C_3, C_5\}$-free graphs. J. Graph Theory, 90(1):83–106, 2019.

[31] F. W. Levi. On Helly's theorem and the axioms of convexity. J. Indian Math. Soc. (N.S.), 15(Part):65–76, 1951.

[32] H. Liu, C. Shangguan, J. Skokan, and Z. Xu. Beyond chromatic threshold via (p, q)-theorem, and sharp blow-up phenomenon. arXiv preprint, arXiv: 2403.17910, 2024.

[33] T. Łuczak. On the structure of triangle-free graphs of large minimum degree. Combinatorica, 26(4):489–493, 2006.

[34] T. Łuczak and S. Thomassé. Coloring dense graphs via VC-dimension. arXiv preprint, arXiv: 1007.1670, 2010.

[35] J. Matoušek. Lectures on discrete geometry, volume 212 of Graduate Texts in Mathematics. Springer-Verlag, New York, 2002.

[36] J. Matoušek. Bounded VC-dimension implies a fractional Helly theorem. Discrete Comput. Geom., 31(2):251–255, 2004.

[37] V. Nikiforov. Chromatic number and minimum degree of K_r-free graphs. arXiv preprint, arXiv: 1001.2070, 2010.

[38] H. Oberkampf and M. Schacht. On the structure of dense graphs with bounded clique number. Comb. Probab. Comput., 29(5):641–649, 2020.

[39] J. Radon. Mengen konvexer Körper, die einen gemeinsamen Punkt enthalten. Math. Ann., 83(1-2):113–115, 1921.

[40] M. Sankar. Homotopy and the homomorphism threshold of odd cycles. arXiv preprint, arXiv: 2206.07525, 2022.

[41] M. Simonovits. A method for solving extremal problems in graph theory, stability problems. Theory of Graphs, Proc. Colloq. Tihany, Hungary 1966, 279-319 (1968)., 1968.

[42] E. Szemerédi. Regular partitions of graphs. Problèmes combinatoires et théorie des graphes, Orsay 1976, Colloq. int. CNRS No. 260, 399-401 (1978)., 1978.

[43] C. Thomassen. On the chromatic number of triangle-free graphs of large minimum degree. Combinatorica, 22(4):591–596, 2002.

[44] C. Thomassen. On the chromatic number of pentagon-free graphs of large minimum degree. Combinatorica, 27(2):241–243, 2007.

[45] P. Turán. Eine extremalaufgabe aus der graphentheorie. Fiz Lapok, pages 436–452, 1941.

[46] M. L. J. van de Vel. Theory of convex structures, volume 50 of North-Holland Mathematical Library. North-Holland Publishing Co., Amsterdam, 1993.