arXiv:2610.28574v1 [math.CO] 20 Oct 2026

Vanishing Uniform Turán Densities: Characterization and Jumps

Hao Lin * Jie Ma † Wenling Zhou ‡

Abstract

Uniform Turán density problems ask for the minimum density threshold that guarantees a copy of a fixed 3-uniform hypergraph F in every sufficiently large 3-uniform hypergraph whose edges are uniformly distributed. This line of research was initiated by Erdős and Sós in the 1980s, who introduced the notion of *vertex-set Turán density* $\pi_{\therefore}(F)$. Reiher, Rödl and Schacht later defined two related concepts, namely, the *vertex-pair Turán density* $\pi_{\therefore}(F)$ and the *cherry Turán density* $\pi_{\wedge}(F)$, which satisfy the chain of inequalities $\pi_{\therefore}(F) \geq \pi_{\therefore}(F) \geq \pi_{\wedge}(F)$. In this paper, we investigate the problem for their vanishing densities.

First, we characterize the 3-uniform hypergraphs F for which $\pi_{\wedge}(F) = 0$, thereby answering a question posed by Schacht. We show that these are precisely the so-called *max-layered* 3-uniform hypergraphs recently introduced by Lamaison. Furthermore, we answer an analogous question of Lamaison for the *minimum-codegree palette density* $\pi_{\wedge}^{\mathsf{pal}}$. Together, these results imply that the following are equivalent: $\pi_{\wedge}(F) = 0$, $\pi_{\wedge}^{\mathsf{pal}}(F) = 0$, and F is a max-layered 3-uniform hypergraph.

Let $\star \in \{\therefore, \therefore, \wedge\}$. We say that zero is a *jump* for $\pi_\star$ if there exists $c > 0$ such that no 3-uniform hypergraph F satisfies $\pi_\star(F) \in (0, c)$. A famous result of Reiher, Rödl and Schacht asserts that zero is a jump for $\pi_{\therefore}$. In contrast, we prove that zero is not a jump for either $\pi_{\therefore}$ or $\pi_{\wedge}$. This also yields a negative answer to Schacht's well-ordering question for $\pi_{\therefore}$. Together with existing results, these findings complete the picture for all three notions of uniform Turán density at zero.

1 Introduction

A central problem in Extremal Combinatorics, initiated by Turán [34], asks how dense a large graph or hypergraph can be while avoiding a fixed subgraph. For an integer $k \geq 2$, a *k-uniform hypergraph*, or simply a *k-graph*, is a pair $H = (V, E)$ with $E \subseteq \binom{V}{k}$. The *Turán number* $\mathrm{ex}(n, F)$ of a k-graph F denotes the maximum number of edges in an n-vertex k-graph that is *F-free*, i.e., containing no copy of F as a subgraph. The standard averaging argument shows that $\mathrm{ex}(n, F)/\binom{n}{k}$ is non-increasing in n, so we may legitimately define the *Turán density* of F by

$$\pi(F) = \lim_{n\to\infty} \frac{\mathrm{ex}(n, F)}{\binom{n}{k}}.$$

*School of Mathematical Sciences, University of Science and Technology of China, Hefei, Anhui 230026, China. Email: haolinz6@ustc.edu.cn.

†School of Mathematical Sciences, University of Science and Technology of China, Hefei, Anhui 230026, and Yau Mathematical Sciences Center, Tsinghua University, Beijing 100084, China. Research supported by National Key Research and Development Program of China 2023YFA1010201 and National Natural Science Foundation of China grant 12125106. Email: jiema@ustc.edu.cn.

‡School of Mathematics, Shandong University, Jinan, Shandong 250100, China. Research supported by the National Natural Science Foundation of China (No. 12401457), Shandong Provincial Natural Science Foundation (No. ZR2024QA067), China Postdoctoral Science Foundation (No. 2024M761780), and Young Talent of Lifting engineering for Science and Technology in Shandong, China (No. SDAST2025QTA074). Email: gracezhou@sdu.edu.cn.

Keywords: Uniform Turán density, max-layered 3-graphs, jump points.

The study of Turán densities has been the focus of extensive research. For every $k \geq 2$, let $\Pi^{(k)} = \{\pi(F) : F \text{ is a } k\text{-graph}\}$ consist of all Turán densities of single k-graphs. For graphs, Turán densities are well understood. Mantel's theorem [26] and Turán's theorem [35] yield the exact Turán numbers of complete graphs, while the celebrated Erdős–Stone–Simonovits theorem [10, 12] determines the Turán density $\pi(F) = 1 - \frac{1}{\chi(F)-1}$ for every graph F with chromatic number $\chi(F) \geq 2$. The latter result implies that $\Pi^{(2)} = \{(r-1)/r : r \in \mathbb{N}^+\}$. For $k \geq 3$, however, it is notoriously difficult to understand the Turán density $\pi(F)$ for general k-graphs F; see Keevash [17] for a survey.

Throughout the paper, we focus on 3-graphs. Even the Turán densities of two of the simplest 3-graphs, the complete 3-graph $K_4^{(3)}$ on four vertices and the 3-graph $K_4^{(3)-}$ obtained from it by deleting one edge, remain unknown [1, 8, 13]. Very little has been known about $\Pi^{(3)}$; see, e.g., [6].

It is historically important to understand the behavior of Turán densities at zero. There are two basic questions. The first is to characterize hypergraphs F for which $\pi(F) = 0$. The second is to determine whether the positive Turán densities are bounded away from zero or can be arbitrarily close to zero. More precisely, we say that zero is a *jump* in $\Pi^{(3)}$ if there exists $c > 0$ such that $\Pi^{(3)} \cap (0, c) = \emptyset$. For the classical Turán density of 3-graphs, the first question was settled by Erdős [7], and the second follows from his characterization. A 3-graph is 3-*partite* if its vertex set can be partitioned into three classes such that every edge meets each class in exactly one vertex.

Theorem 1.1 (Erdős [7])**.** *For every 3-graph F, $\pi(F) = 0$ if and only if F is 3-partite.*

Note that if F is not 3-partite, then every complete 3-partite 3-graph is F-free. Taking the balanced complete 3-partite 3-graph gives an F-free 3-graph of asymptotic edge density $2/9$. This implies that $\pi(F) > 0$ in fact forces $\pi(F) \geq \frac{2}{9}$. In particular, this shows that zero is a jump in $\Pi^{(3)}$.

Throughout the paper, for distinct vertices x, y, z, we write xy for $\{x, y\}$ and xyz for $\{x, y, z\}$. For a positive integer ℓ, we denote by $[\ell]$ the set $\{1, \ldots, \ell\}$.

1.1 Uniform Turán densities: background and their behavior at zero

Many extremal constructions for classical hypergraph Turán problems have highly non-uniform edge distributions and often contain large independent sets. This led Erdős and Sós [9, 11] in the 1980s to propose studying Turán problems under the additional requirement that the edges of the host 3-graph be uniformly distributed. In this line of research, three natural notions of uniform Turán density for 3-graphs have been introduced and have attracted considerable attention: the *vertex-set*, *vertex-pair*, and *cherry* Turán densities.

We begin with the vertex-set Turán density. For $d \in [0, 1]$ and $\mu > 0$, a 3-graph $H = (V, E)$ is $(d, \mu, \therefore)$-*dense* if, for all $X, Y, Z \subseteq V$, the number $e_{\therefore}(X, Y, Z)$ of ordered triples $(x, y, z) \in X \times Y \times Z$ with $xyz \in E$ satisfies

$$e_{\therefore}(X, Y, Z) \geq d|X||Y||Z| - \mu|V|^3.$$

The corresponding *vertex-set Turán density* of a fixed 3-graph F is defined by

$$\pi_{\therefore}(F) = \sup\Big\{d \in [0, 1] : \begin{array}{l}\text{for every } \mu > 0 \text{ and } n_0 \in \mathbb{N}, \text{ there exists an } F\text{-free,} \\ (d, \mu, \therefore)\text{-dense 3-graph } H \text{ with } |V(H)| \geq n_0\end{array}\Big\}.$$

With this notation, Erdős and Sós asked to determine $\pi_{\therefore}(K_4^{(3)})$ and $\pi_{\therefore}(K_4^{(3)-})$. The conjecture that $\pi_{\therefore}(K_4^{(3)-}) = 1/4$ was proved by Glebov, Král' and Volec [16], and independently by Reiher, Rödl and Schacht [30]. More recently, the conjecture $\pi_{\therefore}(K_4^{(3)}) = 1/2$ was proved by Kielak, Král', Lamaison, Liu, Shu and Wu [18], and independently by Bucić [3], confirming the optimality of Rödl's construction [32]. Besides, recent work has also determined several new exact values for $\pi_{\therefore}$ including

the minimum positive value [15], tight cycles [4], large stars [22, 25] and families with density 1/4 [5, 23] or 8/27 [14]. Further structural and classification results can be found in [19–21, 24].

The $\therefore$-dense condition provides a natural perspective to avoid extremal constructions with large independent sets. In their treatment of hypergraph Turán problems, Reiher, Rödl and Schacht [28, 31] introduced two stronger notions of uniform density. For $d \in [0,1]$ and $\mu > 0$, a 3-graph $H = (V, E)$ is $(d, \mu, \therefore)$-*dense* if, for any subset $X \subseteq V$ and any set of ordered pairs $P \subseteq V \times V$, the number $e_{\therefore}(X, P)$ of elements $(x, (y, z)) \in X \times P$ with $xyz \in E$ satisfies

$$e_{\therefore}(X, P) \geq d|X||P| - \mu|V|^3.$$

The corresponding *vertex-pair Turán density* of a fixed 3-graph F is defined by

$$\pi_{\therefore}(F) = \sup\Big\{d \in [0,1] : \begin{array}{l}\text{for every } \mu > 0 \text{ and } n_0 \in \mathbb{N}, \text{ there exists an } F\text{-free,} \\ (d, \mu, \therefore)\text{-dense 3-graph } H \text{ with } |V(H)| \geq n_0\end{array}\Big\}.$$

The vertex-pair condition can be strengthened once more by allowing the two ordered pairs (x, y) and (y, z), which share the middle vertex y. More precisely, a 3-graph $H = (V, E)$ is called $(d, \mu, \wedge)$-*dense* if for every $P, Q \subseteq V \times V$, the number $e_{\wedge}(P, Q)$ of members $((x, y), (y, z)) \in P \times Q$ with $xyz \in E$ satisfies

$$e_{\wedge}(P, Q) \geq d|\mathcal{K}_{\wedge}(P, Q)| - \mu|V|^3,$$

where $\mathcal{K}_{\wedge}(P, Q) = \{((x, y), (y, z)) \in P \times Q\}$. The corresponding *cherry Turán density* of a given 3-graph F is defined by

$$\pi_{\wedge}(F) = \sup\Big\{d \in [0,1] : \begin{array}{l}\text{for every } \mu > 0 \text{ and } n_0 \in \mathbb{N}, \text{ there exists an } F\text{-free,} \\ (d, \mu, \wedge)\text{-dense 3-graph } H \text{ with } |V(H)| \geq n_0\end{array}\Big\}.$$

By the above definitions, every $(d, \mu, \wedge)$-dense 3-graph is also $(d, \mu, \therefore)$-dense, and every $(d, \mu, \therefore)$-dense 3-graph is also $(d, \mu, \therefore)$-dense. Therefore, for every 3-graph F,

$$\pi_{\therefore}(F) \geq \pi_{\therefore}(F) \geq \pi_{\wedge}(F). \tag{1.1}$$

For the tetrahedron $K_4^{(3)}$, all three uniform Turán densities considered here are now known. The result [3, 18] mentioned above, together with the main theorem in [28], gives $\pi_{\therefore}(K_4^{(3)}) = \pi_{\therefore}(K_4^{(3)}) = 1/2$, whereas $\pi_{\wedge}(K_4^{(3)}) = 0$ was proved in [31]. More generally, the results in [2, 31] determine $\pi_{\wedge}(K_t^{(3)})$ for every $4 \leq t \leq 16$, except for $t \in \{9, 10\}$.

We now return to the two questions concerning the vanishing densities introduced above. Reiher, Rödl and Schacht [29] gave a characterization of 3-graphs with vertex-set Turán density being zero. To state it, let $\partial F := \{uv \in \binom{V(F)}{2} : uv \subseteq e \text{ for some } e \in E(F)\}$ be the *shadow* of a 3-graph F.

Theorem 1.2 (Reiher, Rödl and Schacht [29])**.** *For every 3-graph F, $\pi_{\therefore}(F) = 0$ if and only if there exist a linear order $\prec$ on $V(F)$, and a 3-coloring $\varphi : \partial F \to \{\text{red}, \text{blue}, \text{green}\}$ such that every edge $uvw \in E(F)$ with $u \prec v \prec w$ satisfies*

$$\varphi(uv) = \text{red}, \qquad \varphi(uw) = \text{blue}, \qquad \varphi(vw) = \text{green}.$$

If a 3-graph F satisfies $\pi_{\therefore}(F) > 0$, then its vertex-set Turán density is at least 1/27 because of the following construction from [29]. Consider a random n-vertex 3-graph H_n as follows: let $V(H_n) = [n]$; color each pair of $[n]$ red, blue, or green independently and uniformly at random, and include ijk with $1 \leq i < j < k \leq n$ as an edge of H_n if ij, ik, jk have colors red, blue, and green, respectively. By Theorem 1.2, it is straightforward to see that H_n is F-free if $\pi_{\therefore}(F) > 0$. On

the other hand, since each triple becomes an edge with probability 1/27, a standard concentration argument shows that, with high probability, H_n is $(1/27, o(1), \therefore)$-dense. Thus, zero is a jump for $\pi_{\therefore}$. In fact, 1/27 is the smallest positive density attained by $\pi_{\therefore}$ (see Garbe, Král' and Lamaison [15]).

Thus, for $\pi_{\therefore}$, both the zero-density class and the behavior immediately above zero are completely characterized. It is natural to ask for analogous results for the two stronger densities $\pi_{\dot{\therefore}}$ and $\pi_{\wedge}$. In his 2022 ICM report, Schacht [33] posed three problems concerning their behavior at zero. One of these asked for a useful characterization of the 3-graphs F with $\pi_{\dot{\therefore}}(F) = 0$; see [33, Problem 4.4]. This question was subsequently answered in the breakthrough work of Lamaison [21].

To state Lamaison's characterization, we first introduce the following notation.

Definition 1.3 (Edge-patterns)**.** Let F be a 3-graph, let $\prec$ be an order on $V(F)$, and let $\lambda : \partial F \to \mathbb{N}$. For an edge $uvw \in E(F)$ with $u \prec v \prec w$, the ordered triple $(\lambda(uv), \lambda(uw), \lambda(vw))$ is called the *edge-pattern* of uvw with respect to $(\prec, \lambda)$. The *edge-pattern set* of F with respect to $(\prec, \lambda)$ is denoted by

$$\mathcal{B}_{\lambda,\prec} := \{(\lambda(uv), \lambda(uw), \lambda(vw)) : uvw \in E(F),\ u \prec v \prec w\}. \tag{1.2}$$

A set $B \subseteq \mathbb{N}^3$ is *min-degenerate* if every member of B has three distinct entries and, whenever two distinct members of B share an entry, that entry is the minimum entry of at least one of them. A 3-graph F is *min-layered* if there exist an order $\prec$ on $V(F)$ and a map $\lambda : \partial F \to \mathbb{N}$ such that $\mathcal{B}_{\lambda,\prec}$ is min-degenerate. Lamaison [21] proved the following characterization.

Theorem 1.4 (Lamaison [21])**.** *For every 3-graph F, $\pi_{\dot{\therefore}}(F) = 0$ if and only if F is min-layered.*

Let $\Pi^{(3)}_{\dot{\therefore}} = \{\pi_{\dot{\therefore}}(F) : F \text{ is a 3-graph}\}$. Although Theorem 1.4 gives a zero-density characterization for $\pi_{\dot{\therefore}}$, it remained open whether zero is a jump in $\Pi^{(3)}_{\dot{\therefore}}$. Another problem of Schacht [33] asked a more general question for the set $\Pi^{(3)}_{\dot{\therefore}}$ as follows. We say a subset $S \subseteq \mathbb{R}$ is *well-ordered* under the usual order if every nonempty subset of S has a least element.

Problem 1.5 ([33], Problem 4.3(iii))**.** Show that $\Pi^{(3)}_{\dot{\therefore}}$ is well-ordered as a subset of the reals.

A positive answer to this problem would, in particular, imply that zero is a jump in $\Pi^{(3)}_{\dot{\therefore}}$.

The last problem of Schacht [33] mentioned above asked for a characterization of the zero-density class of the cherry Turán density.

Problem 1.6 ([33], Problem 5.3(i))**.** Find a useful characterization of the 3-graphs F with $\pi_{\wedge}(F) = 0$.

Lamaison [21] had also asked the analogous zero-density question for the minimum-codegree palette density (see the precise definition of $\pi^{\mathsf{pal}}_{\wedge}(F)$ we give shortly).

These questions stated above provide the driving motivation for our investigation of vanishing uniform Turán densities. We will introduce our main results in the next subsection, which, together with the existing results, give a complete picture for all three notions of uniform Turán density at zero.

1.2 Our results

Our first result characterizes the zero-density class of the cherry Turán density and answers Problem 1.6. It is stated in terms of max-layeredness, a notion introduced by Lamaison [21]. A set $B \subseteq \mathbb{N}^3$ is *max-degenerate* if the maximum entry of every member of B occurs in a unique position and distinct members of B have distinct maximum entries.

Definition 1.7 (Max-layered 3-graphs [21])**.** A 3-graph F is *max-layered* if there exist an order $\prec$ on $V(F)$ and a map $\lambda : \partial F \to \mathbb{N}$ such that the edge-pattern set $\mathcal{B}_{\lambda,\prec}$ is max-degenerate. In this case, $(\prec, \lambda)$ is called a *max-layered witness* for F.

Theorem 1.8. *For every 3-graph F, $\pi_{\Lambda}(F) = 0$ if and only if F is max-layered.*

Our argument also settles a question of Lamaison [21]. A *palette* is a pair $\mathscr{P} = (\mathscr{C}, \mathscr{A})$, where $\mathscr{C}$ is a finite nonempty set of colors and $\mathscr{A} \subseteq \mathscr{C}^3$ is a set of admissible ordered color triples. A 3-graph F *admits* $\mathscr{P}$ if there exist an order $\prec$ on $V(F)$ and a coloring $\varphi : \partial F \to \mathscr{C}$ such that

$$(\varphi(uv), \varphi(uw), \varphi(vw)) \in \mathscr{A}$$

for every edge $uvw \in E(F)$ with $u \prec v \prec w$. We call $(\prec, \varphi)$ an *admission witness* for F and $\mathscr{P}$. In fact, if we view $\mathscr{C}$ as $[|\mathscr{C}|]$, then F *admits* $\mathscr{P}$ if and only if the edge-pattern $\mathcal{B}_{\varphi,\prec} \subseteq \mathscr{A}$.

For $a, b \in \mathscr{C}$, define $d_{12}(a,b) := \big|\{c \in \mathscr{C} : (a,b,c) \in \mathscr{A}\}\big|, d_{13}(a,b) := \big|\{c \in \mathscr{C} : (a,c,b) \in \mathscr{A}\}\big|, d_{23}(a,b) := \big|\{c \in \mathscr{C} : (c,a,b) \in \mathscr{A}\}\big|$. The *minimum codegree* of $\mathscr{P}$ is

$$\delta_{\Lambda}(\mathscr{P}) := \frac{1}{|\mathscr{C}|} \min_{a,b \in \mathscr{C}} \min\{d_{12}(a,b), d_{13}(a,b), d_{23}(a,b)\}. \tag{1.3}$$

The corresponding *minimum-codegree palette density* is

$$\pi_{\Lambda}^{\mathsf{pal}}(F) := \sup\{\delta_{\Lambda}(\mathscr{P}) : \mathscr{P} \text{ is a palette not admitted by } F\}.$$

Lamaison [21] asked whether $\pi_{\Lambda}^{\mathsf{pal}}(F) = 0$ is equivalent to that F is max-layered. We answer this question affirmatively.

Theorem 1.9. *For every 3-graph F, $\pi_{\Lambda}^{\mathsf{pal}}(F) = 0$ if and only if F is max-layered.*

Together, Theorems 1.8 and 1.9 show that, for every 3-graph F, the following are equivalent:

$$\pi_{\Lambda}(F) = 0 \quad \Longleftrightarrow \quad F \text{ is max-layered} \quad \Longleftrightarrow \quad \pi_{\Lambda}^{\mathsf{pal}}(F) = 0. \tag{1.4}$$

We next turn to the densities immediately above zero. In contrast to the classical and vertex-set Turán densities, the positive values of both $\pi_{\therefore}$ and π_{Λ} can be arbitrarily close to zero.

Theorem 1.10. *For every $\varepsilon > 0$, there exists a 3-graph F_ε such that $0 < \pi_{\Lambda}(F_\varepsilon) \le \pi_{\therefore}(F_\varepsilon) \le \varepsilon$.*

Thus, zero is not a jump for either $\pi_{\therefore}$ or π_{Λ}. For $\star \in \{\therefore, \Lambda\}$, let $\Pi_{\star}^{(3)} = \{\pi_{\star}(F) : F \text{ is a 3-graph}\}$. The above theorem also gives the following consequence.

Corollary 1.11. *The sets $\Pi_{\therefore}^{(3)}$ and $\Pi_{\Lambda}^{(3)}$ are not well-ordered subsets of the reals.*

Indeed, Theorem 1.10 gives, for every $\varepsilon > 0$, a positive element of each set below ε. Hence neither set has a least positive element. In particular, Corollary 1.11 gives a negative answer to Problem 1.5.

Together with the existing results, our theorems determine the zero-density classes (i.e., the families of 3-graphs with vanishing density) and the jump behavior at zero for the three uniform Turán densities considered here. The zero-density classes of $\pi_{\therefore}$, $\pi_{\therefore}$, and π_{Λ} are characterized, respectively, by the ordered 3-coloring condition, min-layeredness, and max-layeredness. Moreover, zero is a jump for $\pi_{\therefore}$; in contrast, zero is not a jump for either $\pi_{\therefore}$ or π_{Λ}.

1.3 Proof overview and organization

We first present the proof overview of Theorems 1.8 and 1.9, then present the proof overview for Theorem 1.10, and finally describe the organization of the paper.

The zero-density characterizations. Before we sketch the proofs of Theorems 1.8 and 1.9, we would like to point out that the two zero-density characterizations are closely related. We first prove that every max-layered 3-graph has zero density in both settings. For the minimum-codegree palette density in Theorem 1.9, this follows from a direct greedy coloring of the layers. For the cherry Turán density in Theorem 1.8, the same layer-by-layer idea is carried out in the reduced 3-graph formulation given by Theorem 2.4, which is the content of Lemma 3.1. The converse implications are proved together in Lemma 3.2. If F is not max-layered, we construct a palette of positive minimum codegree that F does not admit. This immediately gives $\pi^{\mathsf{pal}}_{\bigwedge}(F) > 0$, and Lemma 2.5 then yields $\pi_{\bigwedge}(F) > 0$ as well. Together, these arguments establish both zero-density characterizations. We now describe these proof overview in the following.

Suppose first that F is max-layered. A palette has positive minimum codegree if, once colors have been fixed in any two positions of an ordered color triple, there is an admissible choice for the third. We give all pairs in the same layer the same color and consider the layers in increasing order. When a new layer is reached, max-layeredness ensures that there is at most one edge-pattern whose maximum entry is that layer. The other two colors in this pattern have already been fixed, so positive minimum codegree provides a suitable color for the new layer. Continuing in this way shows that F admits every palette of positive minimum codegree, proving the corresponding direction of Theorem 1.9.

For the cherry Turán density, the same layer-by-layer strategy is carried out in reduced 3-graphs (Definition 2.1). Instead of assigning a color to each layer, we choose representatives from the vertex classes of the reduced 3-graph. The reduced density condition provides many possible choices, and a selection lemma is used to coordinate these choices across all relevant triples. Together with Theorem 2.4, this proves the corresponding direction of Theorem 1.8; see Section 4 for more details.

It remains to show that, when F is not max-layered, both densities are positive. We construct a palette of positive minimum codegree that F does not admit. For each of the three positions, we prescribe the missing color as a function of the other two, and declare a triple admissible whenever it satisfies at least one of these three rules. This automatically gives positive minimum codegree.

If F admitted such a palette, an admission witness would give a vertex order and a coloring of the pairs of F. Choosing, for each edge, one of the three rules satisfied by its color triple would then produce a finite system of equations in these pair colors. For each fixed vertex order and choice of rules, the main lemma Lemma 5.4 uses the failure of max-layeredness to choose the functions so that the corresponding system has no solution. Since there are only finitely many such choices, we combine the resulting constructions coordinatewise to obtain a single palette that F does not admit. This construction is given in Section 5. It proves $\pi^{\mathsf{pal}}_{\bigwedge}(F) > 0$, and Lemma 2.5 then gives $\pi_{\bigwedge}(F) > 0$ as well; see Section 5 for more details.

The non-jump theorem. The proof of Theorem 1.10 also uses palettes, but in a different way. Fix $\eta > 0$. We first construct a universal palette $\mathscr{U}_\eta$. A standard probability method shows that every palette whose vertex-pair minimum degree is greater than η receives a homomorphism from a palette on a bounded number of colors, where a palette homomorphism is a color map preserving admissible triples. Taking the product of all such bounded palettes gives $\mathscr{U}_\eta$, which has positive vertex-pair minimum degree and maps homomorphically into every palette above the threshold η; see Lemma 6.4 and (6.3). Consequently, if a 3-graph F admits $\mathscr{U}_\eta$, then it admits every palette with vertex-pair minimum degree greater than η. By Theorem 6.1, this implies $\pi_{\therefore}(F) \le \eta$.

We then seek a positive lower bound for the cherry Turán density of the same 3-graph. For this purpose, we construct the affine palette $\mathscr{Q}_p$ from Definition 6.5, whose minimum codegree is $1/p$. By Lemma 6.6, the prime p can be chosen so that $\mathscr{U}_\eta$ has no homomorphism to either $\mathscr{Q}_p$ or its

reverse. The palette separation theorem Theorem 6.2 then gives a 3-graph F which admits $\mathscr{U}_\eta$ but does not admit $\mathscr{Q}_p$. The first property gives $\pi_{\lambda}(F) \le \eta$, whereas the second and Lemma 2.5 give $\pi_{\lambda}(F) \ge 1/p > 0$. Since η can be chosen arbitrarily small, this proves the non-jump theorem.

Organization. The remainder of the paper is organized as follows. In Section 2, we recall reduced 3-graphs for the cherry Turán density. In Section 3, we state the two key lemmas (Lemmas 3.1 and 3.2) and deduce Theorems 1.8 and 1.9. The two key lemmas are proved in Sections 4 and 5, respectively. In Section 6, we prove Theorem 1.10 and obtain Corollary 1.11. Finally, in Section 7, we give some new explicit examples of 3-graphs F with $\pi_{\lambda}(F) \ne \pi_{\lambda}(F)$ and discuss further questions.

2 Preliminaries

In this section, we recall reduced 3-graphs and palette constructions for the cherry Turán density.

2.1 Reduced 3-graphs

In this subsection, we introduce the concept of reduced 3-graphs, which encapsulates hypergraph regularity arguments about the uniform Turán density of 3-graphs. We follow some notions introduced in Reiher's survey [27, Section 3]. It should be noted that the concept of reduced 3-graphs has also been mentioned in other related works, e.g., [4, 5, 15, 23, 30].

Definition 2.1 (Reduced 3-graphs)**.** Let $I \subseteq \mathbb{N}$ be a finite index set. For each $ij \in \binom{I}{2}$, let $\mathcal{P}^{ij}$ be a nonempty vertex set, and assume that these sets are pairwise disjoint. For every $ijk \in \binom{I}{3}$, let $\mathcal{A}^{ijk}$ be a 3-partite 3-graph with vertex partition $\{\mathcal{P}^{ij}, \mathcal{P}^{ik}, \mathcal{P}^{jk}\}$. Then the $\binom{|I|}{2}$-partite 3-graph $\mathcal{A}$ with

$$V(\mathcal{A}) = \bigcup_{ij \in \binom{I}{2}} \mathcal{P}^{ij} \qquad \text{and} \qquad E(\mathcal{A}) = \bigcup_{ijk \in \binom{I}{3}} E(\mathcal{A}^{ijk})$$

is called an *I-reduced 3-graph*. The sets $\mathcal{P}^{ij}$ are its *vertex classes*, and the 3-graphs $\mathcal{A}^{ijk}$ are its *constituents*.

Throughout, the superscripts of vertex classes and constituents are unordered: $\mathcal{P}^{ij} = \mathcal{P}^{ji}$, and $\mathcal{A}^{ijk}$ depends only on the set ijk. Whenever an order is needed, the indices are written increasingly.

Definition 2.2 (Reduced λ-density)**.** Let $d \in [0,1]$ and let $\mathcal{A}$ be an I-reduced 3-graph. We say that $\mathcal{A}$ is *(d, λ)-dense* if, for every $i < j < k$ in I, every ordering (X, Y, Z) of the three classes $\mathcal{P}^{ij}, \mathcal{P}^{ik}, \mathcal{P}^{jk}$, and all $x \in X$ and $y \in Y$,

$$\left|N_{\mathcal{A}^{ijk}}(x, y; Z)\right| \ge d|Z|,$$

where

$$N_{\mathcal{A}^{ijk}}(x, y; Z) := \{z \in Z : \{x, y, z\} \in E(\mathcal{A}^{ijk})\}.$$

Definition 2.3 (Reduced images)**.** Let F be a 3-graph and let $\mathcal{A}$ be an I-reduced 3-graph. A *reduced map* from F to $\mathcal{A}$ is a pair (ι, ψ) such that

(i) $\iota : V(F) \to I$ and $\psi : \partial F \to V(\mathcal{A})$;

(ii) if $xy \in \partial F$, then $\iota(x) \ne \iota(y)$ and $\psi(xy) \in \mathcal{P}^{\iota(x)\iota(y)}$;

(iii) if $xyz \in E(F)$, then $\{\psi(xy), \psi(xz), \psi(yz)\} \in E(\mathcal{A}^{\iota(x)\iota(y)\iota(z)})$.

If such a reduced map exists, then $\mathcal{A}$ contains a *reduced image* of F; otherwise, $\mathcal{A}$ is *F-free*.

The following theorem provides the bridge between this reduced formalism and the cherry Turán density by characterizing $\pi_{\curlywedge}(F)$ in terms of F-free reduced 3-graphs with arbitrarily many indices.

Theorem 2.4 (Reiher [27])**.** *Let F be a 3-graph. Then*

$$\pi_{\curlywedge}(F) = \sup\Big\{d \in [0,1] : \begin{array}{l} \textit{for every } m \in \mathbb{N} \textit{ there exists an } F\textit{-free,} \\ (d,\curlywedge)\textit{-dense reduced 3-graph with an index set of size } m \end{array}\Big\}.$$

2.2 Palette lower bounds

Palette constructions go back to Rödl's random pair-coloring idea [32], while the minimum-codegree palette density for the cherry setting was introduced by Lamaison [21, Section 7]. In this subsection, we use the reduced formalism above to relate this palette density to the cherry Turán density.

The following lemma shows that palette obstructions give cherry-density lower bounds. We prove it directly by converting any palette not admitted by F into an F-free reduced 3-graph whose reduced $\curlywedge$-density is at least the minimum codegree of the palette.

Lemma 2.5. *For every 3-graph F, $\pi_{\curlywedge}^{\mathsf{pal}}(F) \le \pi_{\curlywedge}(F)$.*

Proof. Let $\mathscr{P} = (\mathscr{C}, \mathscr{A})$ be a palette not admitted by F, and put $d := \delta_{\curlywedge}(\mathscr{P})$. We show that $d \le \pi_{\curlywedge}(F)$. By Theorem 2.4, it suffices to construct, for every $m \in \mathbb{N}$, an F-free $(d,\curlywedge)$-dense reduced 3-graph with m indices.

Fix $m \in \mathbb{N}$ and take the index set to be $[m]$. For every $ij \in \binom{[m]}{2}$, let $\mathcal{P}^{ij}$ be a disjoint copy of $\mathscr{C}$, and write (ij, a) for the copy of $a \in \mathscr{C}$. For $i < j < k$, define the constituent $\mathcal{A}^{ijk}$ by

$$\{(ij,a),(ik,b),(jk,c)\} \in E(\mathcal{A}^{ijk}) \quad \Longleftrightarrow \quad (a,b,c) \in \mathscr{A}.$$

Denote the resulting reduced 3-graph by $\mathcal{A}_m(\mathscr{P})$.

The definition of $\delta_{\curlywedge}(\mathscr{P})$ gives the required reduced density. Indeed, fixing vertices in the classes $\mathcal{P}^{ij}$ and $\mathcal{P}^{ik}$ leaves at least $d|\mathscr{C}|$ choices in $\mathcal{P}^{jk}$ by d_{12}; the pairs of classes $(\mathcal{P}^{ij}, \mathcal{P}^{jk})$ and $(\mathcal{P}^{ik}, \mathcal{P}^{jk})$ are handled by d_{13} and d_{23}, respectively. Thus $\mathcal{A}_m(\mathscr{P})$ is $(d,\curlywedge)$-dense.

It remains to prove that $\mathcal{A}_m(\mathscr{P})$ is F-free. Suppose to the contrary that (ι, ψ) is a reduced map from F to $\mathcal{A}_m(\mathscr{P})$. For each $xy \in \partial F$, let $\varphi(xy) \in \mathscr{C}$ be the unique color whose copy in $\mathcal{P}^{\iota(x)\iota(y)}$ is $\psi(xy)$. Choose a linear order $\prec$ on $V(F)$ by ordering the vertices increasingly according to ι and breaking ties arbitrarily.

Let $uvw \in E(F)$ with $u \prec v \prec w$. The reduced-map condition implies that $\iota(u), \iota(v), \iota(w)$ are pairwise distinct, and hence $\iota(u) < \iota(v) < \iota(w)$. Moreover,

$$\{\psi(uv), \psi(uw), \psi(vw)\} \in E\big(\mathcal{A}^{\iota(u)\iota(v)\iota(w)}\big).$$

By the definition of the constituent,

$$\big(\varphi(uv), \varphi(uw), \varphi(vw)\big) \in \mathscr{A}.$$

Thus $(\prec, \varphi)$ is an admission witness for F and $\mathscr{P}$, contrary to the choice of $\mathscr{P}$. Therefore $\mathcal{A}_m(\mathscr{P})$ is F-free.

Since this construction works for every $m \in \mathbb{N}$, Theorem 2.4 gives $\delta_{\curlywedge}(\mathscr{P}) \le \pi_{\curlywedge}(F)$. Taking the supremum over all palettes not admitted by F proves the lemma. $\square$

3 Proofs of the zero characterization theorems

This section states the two key lemmas underlying our zero characterizations and then proves Theorems 1.8 and 1.9. Their proofs are deferred to Sections 4 and 5, respectively.

The first key lemma gives the implication from max-layeredness to the vanishing of $\pi_{\boldsymbol{\Lambda}}$.

Lemma 3.1. *If F is a max-layered 3-graph, then $\pi_{\boldsymbol{\Lambda}}(F) = 0$.*

The second key lemma provides a positive lower bound for every non-max-layered 3-graph.

Lemma 3.2. *If a 3-graph F is not max-layered, then we have $\pi^{\mathsf{pal}}_{\boldsymbol{\Lambda}}(F) \geq 3^{-|V(F)|!\,3^{2|E(F)|}} > 0$.*

Together with Lemma 2.5, this immediately implies that $\pi_{\boldsymbol{\Lambda}}(F) > 0$ whenever F is not max-layered.

Proof of Theorem 1.8. If F is max-layered, then Lemma 3.1 gives $\pi_{\boldsymbol{\Lambda}}(F) = 0$.

Conversely, if F is not max-layered, then Lemma 3.2 gives $\pi^{\mathsf{pal}}_{\boldsymbol{\Lambda}}(F) > 0$. By Lemma 2.5, this implies $\pi_{\boldsymbol{\Lambda}}(F) > 0$. Hence $\pi_{\boldsymbol{\Lambda}}(F) = 0$ is impossible. This proves the equivalence. □

The preceding proof combines the two key lemmas to characterize the vanishing of $\pi_{\boldsymbol{\Lambda}}$. We now turn to the palette analogue. The non-max-layered direction again follows from Lemma 3.2. For the max-layered direction, we show directly that F admits every palette with positive minimum codegree by constructing the required coloring.

Proof of Theorem 1.9. If F is not max-layered, then Lemma 3.2 gives $\pi^{\mathsf{pal}}_{\boldsymbol{\Lambda}}(F) > 0$.

Conversely, suppose F is max-layered, witnessed by an order $\prec$ on $V(F)$ and a map $\lambda : \partial F \to \mathbb{N}$ such that $\mathcal{B}_{\lambda,\prec}$ is max-degenerate. Let $\mathscr{P} = (\mathscr{C}, \mathscr{A})$ be a palette satisfying $\delta_{\boldsymbol{\Lambda}}(\mathscr{P}) > 0$. By the definition, for every $a, b \in \mathscr{C}$ we have

$$d_{12}(a,b),\ d_{13}(a,b),\ d_{23}(a,b) \geq 1. \tag{3.1}$$

We may assume $\lambda(\partial F) = [L]$ without loss of generality. Then, we recursively construct a map $\kappa : [L] \to \mathscr{C}$ by assigning colors to $[L]$ in increasing order. First, we choose $\kappa(1) \in \mathscr{C}$ arbitrarily. Now let $2 \leq \ell \leq L$, and suppose that $\kappa(1), \ldots, \kappa(\ell - 1)$ have already been chosen so that

$$\big(\kappa(b_1), \kappa(b_2), \kappa(b_3)\big) \in \mathscr{A} \tag{3.2}$$

for every edge-pattern $(b_1, b_2, b_3) \in \mathcal{B}_{\lambda,\prec}$ whose maximum entry is smaller than ℓ. Since $\mathcal{B}_{\lambda,\prec}$ is max-degenerate, there is at most one edge-pattern in $\mathcal{B}_{\lambda,\prec}$ whose maximum entry is ℓ.

If no such edge-pattern exists, choose $\kappa(\ell) \in \mathscr{C}$ arbitrarily. Otherwise, let $\boldsymbol{b}_\ell = (b_{\ell,1}, b_{\ell,2}, b_{\ell,3}) \in \mathcal{B}_{\lambda,\prec}$ be the unique edge-pattern whose maximum entry is ℓ. The entry of $\boldsymbol{b}_\ell$ equal to ℓ is unique, and its other two entries are smaller than ℓ and hence have already been colored. We distinguish three cases.

If $\boldsymbol{b}_\ell = (\ell, b_{\ell,2}, b_{\ell,3})$, then by (3.1), we may therefore choose $c \in \mathscr{C}$ such that

$$\big(c, \kappa(b_{\ell,2}), \kappa(b_{\ell,3})\big) \in \mathscr{A}$$

and set $\kappa(\ell) := c$. Similarly, if $\boldsymbol{b}_\ell = (b_{\ell,1}, \ell, b_{\ell,3})$, we choose $c \in \mathscr{C}$ such that $\big(\kappa(b_{\ell,1}), c, \kappa(b_{\ell,3})\big) \in \mathscr{A}$ and set $\kappa(\ell) := c$. If $\boldsymbol{b}_\ell = (b_{\ell,1}, b_{\ell,2}, \ell)$, we choose $c \in \mathscr{C}$ such that $\big(\kappa(b_{\ell,1}), \kappa(b_{\ell,2}), c\big) \in \mathscr{A}$ and set $\kappa(\ell) := c$. Thus (3.2) holds for every edge-pattern whose maximum entry is at most ℓ, so the induction continues.

After step L, the map κ is defined on $[L]$, and (3.2) holds for every edge-pattern in $\mathcal{B}_{\lambda,\prec}$.

Define

$$\varphi(xy) := \kappa(\lambda(xy)) \qquad \text{for every } xy \in \partial F.$$

Let $uvw \in E(F)$ with $u \prec v \prec w$. Since its edge-pattern belongs to $\mathcal{B}_{\lambda,\prec}$, by (3.2) we have

$$(\varphi(uv), \varphi(uw), \varphi(vw)) = (\kappa(\lambda(uv)), \kappa(\lambda(uw)), \kappa(\lambda(vw))) \in \mathscr{A}.$$

Thus the order $\prec$ and the coloring φ show that F admits $\mathscr{P}$.

Thus, we have proved that every palette of positive minimum codegree is admitted by F. Consequently, we have $\pi^{\mathsf{pal}}_{\curlywedge}(F) = 0$. This proves the equivalence. □

4 Proof of Lemma 3.1

To prove Lemma 3.1, we choose, for each layer, one representative in every pair-class of the reduced 3-graph. These representatives are chosen in increasing order of the layers. Once the representatives of the earlier layers are fixed, each relevant triple of indices gives a set of possible choices for the new representative. The reduced density condition makes each of these sets large. A pair-class is involved in many triples, however, so its representative must belong to all the corresponding sets on the indices we retain.

Before starting the proof, we recall the following selection lemma, which follows by combining [15, Lemmas 8–10]; see also [29, Corollaries 4.2, 4.3 and 4.6]. Roughly speaking, the lemma states that if each triple has a positive proportion of available choices, then after restricting to a smaller set of indices, we can choose one vertex from each vertex class in a way that works for every remaining triple.

For $i < j < k$, define $p_1(i,j,k) := ij$, $p_2(i,j,k) := ik$, and $p_3(i,j,k) := jk$.

Lemma 4.1. *For every $n \in \mathbb{N}$ and $d > 0$, there exists an integer $N = N(n,d)$ with the following property. Let $\mathcal{A}$ be an I-reduced 3-graph with $|I| \geq N$, and fix $r \in [3]$. Suppose that for every $i < j < k$ in I we are given a set*

$$C_{ijk} \subseteq \mathcal{P}^{p_r(i,j,k)} \qquad \text{with} \qquad |C_{ijk}| \geq d\,|\mathcal{P}^{p_r(i,j,k)}|.$$

Then there exist a set $I' \subseteq I$ with $|I'| = n$ and vertices $c_{ij} \in \mathcal{P}^{ij}$ for all $ij \in \binom{I'}{2}$ such that

$$c_{p_r(i,j,k)} \in C_{ijk} \qquad \text{for every } i < j < k \text{ in } I'.$$

Proof of Lemma 3.1. Fix $d \in (0,1]$, and choose a max-layered witness $(\prec, \lambda)$ for F. Write $V(F) = \{v_1, \ldots, v_n\}$, where $v_1 \prec v_2 \prec \cdots \prec v_n$. We may assume $\lambda(\partial F) = [L]$ without loss of generality. For each $\ell \in [L]$, there is at most one edge-pattern in $\mathcal{B}_{\lambda,\prec}$ whose maximum entry is ℓ. If such an edge-pattern exists, denote it by $\boldsymbol{b}_\ell = (b_{\ell,1}, b_{\ell,2}, b_{\ell,3})$, and let $r_\ell \in [3]$ be the unique position such that $b_{\ell,r_\ell} = \ell$. Thus $b_{\ell,q} < \ell$ for every $q \in [3] \setminus \{r_\ell\}$.

Let $N(n,d)$ be a threshold function supplied by Lemma 4.1. Define integers $m_{L+1}, m_L, \ldots, m_1$ backwards by setting $m_{L+1} = n$ and

$$m_\ell = \begin{cases} N(m_{\ell+1}, d), & \text{if an edge-pattern } \boldsymbol{b}_\ell \text{ exists,} \\ m_{\ell+1}, & \text{otherwise.} \end{cases} \tag{4.1}$$

Let $\mathcal{A}$ be a $(d, \curlywedge)$-dense I-reduced 3-graph with $|I| \geq m_1$. Choose an arbitrary set $I_1 \subseteq I$ of size m_1. Now, we recursively construct nested sets

$$I_1 \supseteq I_2 \supseteq \cdots \supseteq I_{L+1}, \qquad \text{with } |I_\ell| = m_\ell,$$

together with vertices

$$c_{ij}^{b} \in \mathcal{P}^{ij} \qquad \text{for every } b < \ell \text{ and } ij \in \binom{I_\ell}{2}.$$

The construction maintains the following induction condition: for every edge-pattern $\boldsymbol{b} = (b_1, b_2, b_3) \in \mathcal{B}_{\lambda,\prec}$ with $\max\{b_1, b_2, b_3\} < \ell$ and every $i < j < k$ in I_ℓ,

$$\left\{c_{ij}^{b_1}, c_{ik}^{b_2}, c_{jk}^{b_3}\right\} \in E(\mathcal{A}^{ijk}). \tag{4.2}$$

For $\ell = 1$, condition (4.2) is vacuous since no edge-pattern in $\mathcal{B}_{\lambda,\prec}$ has maximum entry 1. Thus we set $I_2 = I_1$ and choose an arbitrary vertex $c_{ij}^{1} \in \mathcal{P}^{ij}$ for every $ij \in \binom{I_2}{2}$. Consequently, (4.2) holds at the beginning of step 2.

For each $\ell \in \{2, \dots, L\}$, the construction divides into two cases. Suppose first that no edge-pattern in $\mathcal{B}_{\lambda,\prec}$ has maximum entry ℓ. In this case, set $I_{\ell+1} = I_\ell$ and choose an arbitrary vertex $c_{ij}^{\ell} \in \mathcal{P}^{ij}$ for every $ij \in \binom{I_{\ell+1}}{2}$. Since no new edge-pattern has to be realized at this step, condition (4.2) holds with $\ell + 1$.

Now suppose that $\boldsymbol{b}_\ell$ exists. For each $q \in [3] \setminus \{r_\ell\}$, the entry $b_{\ell,q}$ is smaller than ℓ, so for every $ij \in \binom{I_\ell}{2}$, the vertex $c_{ij}^{b_{\ell,q}}$ has already been chosen. For every $i < j < k$ in I_ℓ, define

$$C_{ijk} := \left\{x \in \mathcal{P}^{p_{r_\ell}(i,j,k)} : \{x\} \cup \left\{c_{p_q(i,j,k)}^{b_{\ell,q}} : q \in [3] \setminus \{r_\ell\}\right\} \in E(\mathcal{A}^{ijk})\right\}. \tag{4.3}$$

Since $\mathcal{A}$ is $(d, \boldsymbol{\lambda})$-dense, by Definition 2.2 we have

$$|C_{ijk}| \geq d\left|\mathcal{P}^{p_{r_\ell}(i,j,k)}\right| \qquad \text{for every } i < j < k \text{ in } I_\ell.$$

As $|I_\ell| = N(m_{\ell+1}, d)$, we may apply Lemma 4.1 to the restriction of $\mathcal{A}$ to I_ℓ, with $m_{\ell+1}$ and r_ℓ. We obtain a set $I_{\ell+1} \subseteq I_\ell$ of size $m_{\ell+1}$ and vertices $c_{ij}^{\ell} \in \mathcal{P}^{ij}$ for all $ij \in \binom{I_{\ell+1}}{2}$ such that

$$c_{p_{r_\ell}(i,j,k)}^{\ell} \in C_{ijk} \qquad \text{for every } i < j < k \text{ in } I_{\ell+1}.$$

Then by (4.3), the edge-pattern $\boldsymbol{b}_\ell$ is now realized on every triple of $I_{\ell+1}$, thus the condition (4.2) remains valid with $\ell + 1$.

After the final step, $|I_{L+1}| = |V(F)|$, and condition (4.2) applies to every edge-pattern in $\mathcal{B}_{\lambda,\prec}$. Write $I_{L+1} = \{\alpha_1 < \cdots < \alpha_n\}$, define $\iota : V(F) \to I$ by

$$\iota(v_i) = \alpha_i \qquad (i \in [n]),$$

and, for every $v_i v_j \in \partial F$ with $i < j$, define

$$\psi(v_i v_j) := c_{\alpha_i \alpha_j}^{\lambda(v_i v_j)} \in \mathcal{P}^{\alpha_i \alpha_j}.$$

Thus, for every shadow pair $v_i v_j$ with $i < j$, we have

$$\iota(v_i) \neq \iota(v_j) \qquad \text{and} \qquad \psi(v_i v_j) \in \mathcal{P}^{\iota(v_i)\iota(v_j)}.$$

Let $v_i v_j v_k \in E(F)$ with $i < j < k$. Its edge-pattern is $(\lambda(v_i v_j), \lambda(v_i v_k), \lambda(v_j v_k)) \in \mathcal{B}_{\lambda,\prec}$. Thus, applying the condition (4.2) to $\alpha_i < \alpha_j < \alpha_k$ gives

$$\{\psi(v_i v_j), \psi(v_i v_k), \psi(v_j v_k)\} \in E(\mathcal{A}^{\alpha_i \alpha_j \alpha_k}).$$

Hence (ι, ψ) is a reduced map from F to $\mathcal{A}$.

We have therefore shown that every $(d, \boldsymbol{\lambda})$-dense reduced 3-graph with at least m_1 indices contains a reduced image of F. Since $d > 0$ is arbitrary, Theorem 2.4 yields $\pi_{\boldsymbol{\lambda}}(F) = 0$. □

5 Proof of Lemma 3.2

In this section, we prove Lemma 3.2. We will define the palette using three functions, one for each position of the missing color. A triple is allowed whenever one of its colors is the prescribed function of the other two. On a finite color set, this guarantees positive minimum codegree.

To exclude F, we first fix a vertex order and choose, for each edge, which of the three prescriptions it is to satisfy. These choices give a system of equations in the colors assigned to the pairs of F. The type families introduced below record which variable is determined by which two others in each equation. We first construct functions for which each such system has no solution, and then combine the constructions to handle all vertex orders and choices simultaneously.

In the following, we begin by introducing the definition of type families and explaining their connection with max-layeredness, thereby preparing the ground for the palette construction that follows. We then construct the desired palette, and conclude the section by proving the auxiliary lemma on which the construction relies.

5.1 Type families and max-layeredness

Let X be a finite set, and put $\mathcal{T}(X) := X^3 \times [3]$. An ordered 4-tuple $\boldsymbol{\tau} = (h, s, t, i) \in \mathcal{T}(X)$ is called a *type* on X; we call h its *head*, (s, t) its *ordered pair of tails*, and i its *position*. A *type family* on X is a subset $\mathcal{T} \subseteq \mathcal{T}(X)$.

Definition 5.1. A type family $\mathcal{T} \subseteq \mathcal{T}(X)$ is *layer-realizable* if there exists a map $f : X \to \mathbb{N}$ such that the following conditions hold.

(P1) For every type $\boldsymbol{\tau} = (h, s, t, i) \in \mathcal{T}$, we have $f(h) > f(s)$ and $f(h) > f(t)$.

(P2) For any two types $\boldsymbol{\tau} = (h, s, t, i)$ and $\boldsymbol{\tau}' = (h', s', t', i')$ in $\mathcal{T}$, if $f(h) = f(h')$, then $f(s) = f(s')$, $f(t) = f(t')$, and $i = i'$.

Such a map f is called a *layer realization* of $\mathcal{T}$.

We now associate the type family with a 3-graph F. Let $\prec$ be an order on $V(F)$. For every edge $e = uvw \in E(F)$ with $u \prec v \prec w$, write $p_1(e) := uv$, $p_2(e) := uw$, and $p_3(e) := vw$. A *branch choice function* for $(F, \prec)$ is a map $\theta : E(F) \to [3]$.

For each $e \in E(F)$, let $h_e := p_{\theta(e)}(e)$ be the shadow pair in position $\theta(e)$, and let (s_e, t_e) be the ordered pair obtained by deleting that shadow pair from the ordered triple $(p_1(e), p_2(e), p_3(e))$. Thus

$$(s_e, t_e) := \begin{cases} (p_2(e), p_3(e)), & \text{if } \theta(e) = 1, \\ (p_1(e), p_3(e)), & \text{if } \theta(e) = 2, \\ (p_1(e), p_2(e)), & \text{if } \theta(e) = 3. \end{cases} \tag{5.1}$$

We define the type family of F associated with $\prec$ and θ by

$$\mathcal{T}_{\prec,\theta}(F) := \{(h_e, s_e, t_e, \theta(e)) : e \in E(F)\} \subseteq \mathcal{T}(\partial F). \tag{5.2}$$

Lemma 5.2. *A 3-graph F is max-layered if and only if there exist an order $\prec$ on $V(F)$ and a branch choice function $\theta : E(F) \to [3]$ such that the associated type family $\mathcal{T}_{\prec,\theta}(F)$ is layer-realizable.*

Proof. Suppose first that F is max-layered, witnessed by an order $\prec$ on $V(F)$ and a map $\lambda : \partial F \to \mathbb{N}$ such that $\mathcal{B}_{\lambda,\prec}$ is max-degenerate.

For every edge $e \in E(F)$, let $\theta(e)$ be the position of the unique maximum entry of its edge-pattern. We claim that the map $\lambda : \partial F \to \mathbb{N}$ is a layer realization of $\mathcal{T}_{\prec,\theta}(F)$.

For each $e \in E(F)$, the shadow pair h_e occupies the position of the unique maximum entry of the edge-pattern of e. Hence $\lambda(h_e) > \lambda(s_e)$ and $\lambda(h_e) > \lambda(t_e)$, so condition (P1) holds. Now let $e, e' \in E(F)$ satisfy $\lambda(h_e) = \lambda(h_{e'})$. The corresponding edge-patterns have the same maximum entry. Since distinct elements of $\mathcal{B}_{\lambda,\prec}$ have distinct maximum entries, these two edge-patterns must be equal. It follows that $\theta(e) = \theta(e')$, $\lambda(s_e) = \lambda(s_{e'})$, and $\lambda(t_e) = \lambda(t_{e'})$. Thus condition (P2) holds, and $\mathcal{T}_{\prec,\theta}(F)$ is layer-realizable.

Conversely, suppose that there exist an order $\prec$ on $V(F)$, a branch choice function $\theta : E(F) \to [3]$, and a layer realization $f : \partial F \to \mathbb{N}$ of $\mathcal{T}_{\prec,\theta}(F)$. We show that $(\prec, f)$ is a max-layered witness for F. By condition (P1), for every edge $e \in E(F)$, the entry in position $\theta(e)$ is the unique maximum entry of its edge-pattern.

Now let $e, e' \in E(F)$, and suppose that their edge-patterns with respect to $(\prec, f)$ have the same maximum entry. Since these maxima are $f(h_e)$ and $f(h_{e'})$, respectively, we have $f(h_e) = f(h_{e'})$. Condition (P2) gives $\theta(e) = \theta(e')$, $f(s_e) = f(s_{e'})$, and $f(t_e) = f(t_{e'})$. Hence the two edge-patterns are equal. Therefore $\mathcal{B}_{f,\prec}$ is max-degenerate, and F is max-layered. □

The above lemma shows that a 3-graph F is not max-layered if and only if, for every order $\prec$ and every branch choice function θ, the associated type family $\mathcal{T}_{\prec,\theta}(F)$ is not layer-realizable. We next introduce color systems and define compatibility between a type family and a color system. A *color-system* is a tuple $\mathcal{S} = (\mathcal{C}, g_1, g_2, g_3)$, where $\mathcal{C}$ is a finite nonempty color set and $g_i : \mathcal{C}^2 \to \mathcal{C}$ for each $i \in [3]$.

Definition 5.3. Let $\mathcal{T}$ be a type family on X, and let $\mathcal{S}$ be a color-system. We say that $\mathcal{T}$ is *compatible with* $\mathcal{S}$ if there exists a map $\varphi : X \to \mathcal{C}$ such that

$$\varphi(h) = g_i(\varphi(s), \varphi(t)) \qquad \text{for every } (h, s, t, i) \in \mathcal{T}. \tag{5.3}$$

Such a map φ is called a *compatible color function* for $\mathcal{T}$ and $\mathcal{S}$.

For fixed $\mathcal{T}$ and $\mathcal{S}$, the equations in (5.3) form a system of compatibility equations whose unknown is the map $\varphi : X \to \mathcal{C}$. The next lemma connects the solvability of these compatibility equations with layer realizability and plays a key role in the proof of Lemma 3.2. We defer its proof to Section 5.3.

Lemma 5.4. *Let X be a finite set and let $\mathcal{T} \subseteq \mathcal{T}(X)$ be a type family. If $\mathcal{T}$ is not layer-realizable, then there exists a color-system $\mathcal{S} = (\mathcal{C}, g_1, g_2, g_3)$ such that $\mathcal{T}$ is not compatible with $\mathcal{S}$. Moreover, $\mathcal{C}$ may be chosen so that $|\mathcal{C}| \le 3^{3^{|\mathcal{T}|}}$.*

5.2 The palette construction

With Lemma 5.4 in hand, we now combine the resulting non-solvability into the palette construction.

Proof of Lemma 3.2. Suppose that F is not max-layered. Let

$$\Omega := \{(\prec, \theta) : \prec \text{ is an order on } V(F) \text{ and } \theta \text{ is a branch choice function for } (F, \prec)\}.$$

There are $|V(F)|!$ choices for $\prec$ and, for each such order, $3^{|E(F)|}$ choices for θ. Hence

$$|\Omega| = |V(F)|!\, 3^{|E(F)|}.$$

For every $\omega = (\prec, \theta) \in \Omega$, put $\mathcal{T}_\omega := \mathcal{T}_{\prec,\theta}(F)$, then $|\mathcal{T}_\omega| = |E(F)|$. By Lemma 5.2, the type family $\mathcal{T}_\omega$ is not layer-realizable. Hence Lemma 5.4 provides a color-system $\mathcal{S}_\omega = (\mathcal{C}_\omega, g_1^\omega, g_2^\omega, g_3^\omega)$ such that $\mathcal{T}_\omega$ is not compatible with $\mathcal{S}_\omega$ and $|\mathcal{C}_\omega| \le 3^{3^{|E(F)|}}$.

Let

$$\mathcal{C} := \prod_{\omega \in \Omega} \mathcal{C}_\omega.$$

Then the set $\mathcal{C}$ is finite and nonempty, and

$$|\mathcal{C}| = \prod_{\omega \in \Omega} |\mathcal{C}_\omega| \le \left(3^{3^{|E(F)|}}\right)^{|V(F)|!\,3^{|E(F)|}} = 3^{|V(F)|!\,3^{2|E(F)|}}.$$

For $\boldsymbol{a} = (a_\omega)_{\omega\in\Omega}$ and $\boldsymbol{b} = (b_\omega)_{\omega\in\Omega}$ in $\mathcal{C}$, define $g_i : \mathcal{C}^2 \to \mathcal{C}$ coordinatewise by

$$g_i(\boldsymbol{a}, \boldsymbol{b}) := \left(g_i^\omega(a_\omega, b_\omega)\right)_{\omega\in\Omega} \qquad (i \in [3]),$$

and put $\mathcal{S} := (\mathcal{C}, g_1, g_2, g_3)$. Then we define $\mathscr{A}^{(1)} := \{(g_1(\boldsymbol{b}, \boldsymbol{c}), \boldsymbol{b}, \boldsymbol{c}) : \boldsymbol{b}, \boldsymbol{c} \in \mathcal{C}\}$, $\mathscr{A}^{(2)} := \{(\boldsymbol{a}, g_2(\boldsymbol{a}, \boldsymbol{c}), \boldsymbol{c}) : \boldsymbol{a}, \boldsymbol{c} \in \mathcal{C}\}$ and $\mathscr{A}^{(3)} := \{(\boldsymbol{a}, \boldsymbol{b}, g_3(\boldsymbol{a}, \boldsymbol{b})) : \boldsymbol{a}, \boldsymbol{b} \in \mathcal{C}\}$, and let

$$\mathscr{A} := \mathscr{A}^{(1)} \cup \mathscr{A}^{(2)} \cup \mathscr{A}^{(3)}, \qquad \mathscr{P} := (\mathcal{C}, \mathscr{A}).$$

Thus for all $\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{c} \in \mathcal{C}$, the triples

$$(\boldsymbol{a}, \boldsymbol{b}, g_3(\boldsymbol{a}, \boldsymbol{b})), \qquad (\boldsymbol{a}, g_2(\boldsymbol{a}, \boldsymbol{c}), \boldsymbol{c}), \qquad (g_1(\boldsymbol{b}, \boldsymbol{c}), \boldsymbol{b}, \boldsymbol{c})$$

belong to $\mathscr{A}$. Hence $d_{12}(\boldsymbol{a}, \boldsymbol{b}) \ge 1$, $d_{13}(\boldsymbol{a}, \boldsymbol{c}) \ge 1$, and $d_{23}(\boldsymbol{b}, \boldsymbol{c}) \ge 1$, and therefore

$$\delta_{\boldsymbol{\lambda}}(\mathscr{P}) \ge \frac{1}{|\mathcal{C}|} \ge 3^{-\,|V(F)|!\,3^{2|E(F)|}}.$$

We now show that F does not admit $\mathscr{P}$. Suppose otherwise, and let $\prec$ be an order on $V(F)$ and $\varphi : \partial F \to \mathcal{C}$ be a map such that $(\prec, \varphi)$ is an admission witness for F and $\mathscr{P}$. For every $e = uvw \in E(F)$ with $u \prec v \prec w$, choose $\theta(e) \in [3]$ such that

$$\left(\varphi(uv), \varphi(uw), \varphi(vw)\right) \in \mathscr{A}^{(\theta(e))}.$$

If this triple belongs to more than one of the three sets, choose one such index arbitrarily.

By the definitions of h_e, s_e, t_e and $\mathscr{A}^{(1)}, \mathscr{A}^{(2)}, \mathscr{A}^{(3)}$, we have

$$\varphi(h_e) = g_{\theta(e)}\left(\varphi(s_e), \varphi(t_e)\right) \qquad (e \in E(F)).$$

Thus $\mathcal{T}_{\prec,\theta}(F)$ is compatible with $\mathcal{S}$, witnessed by the compatible color function φ.

Let $\omega_0 = (\prec, \theta) \in \Omega$, and define $\varphi_{\omega_0} : \partial F \to \mathcal{C}_{\omega_0}$ by

$$\varphi_{\omega_0}(x) := \left(\varphi(x)\right)_{\omega_0} \qquad (x \in \partial F).$$

Taking the component indexed by ω_0 in the preceding compatibility equations gives

$$\varphi_{\omega_0}(h_e) = g_{\theta(e)}^{\omega_0}\left(\varphi_{\omega_0}(s_e), \varphi_{\omega_0}(t_e)\right) \qquad (e \in E(F)).$$

Hence $\mathcal{T}_{\omega_0}$ is compatible with $\mathcal{S}_{\omega_0}$, contradicting the choice of $\mathcal{S}_{\omega_0}$.

Therefore F does not admit $\mathscr{P}$, and we obtain

$$\pi_{\boldsymbol{\lambda}}^{\mathsf{pal}}(F) \ge \delta_{\boldsymbol{\lambda}}(\mathscr{P}) \ge 3^{-\,|V(F)|!\,3^{2|E(F)|}} > 0,$$

as required. □

5.3 Proof of Lemma 5.4

It remains to prove Lemma 5.4. The proof is by induction on the number of types. The inductive reduction arises when two types have the same head and position: condition (P2) then forces their corresponding tails to receive the same layer values. We now formalize this reduction in the following definition.

Definition 5.5 (Merge map)**.** Let $\mathcal{T}$ be a type family on X, and suppose that $\mathcal{T}$ contains two types

$$\boldsymbol{\tau}_1 = (h', s_1, t_1, j) \qquad \text{and} \qquad \boldsymbol{\tau}_2 = (h', s_2, t_2, j)$$

with $(s_1, t_1) \neq (s_2, t_2)$. Let $\sim$ be the equivalence relation on X generated by $s_1 \sim s_2$ and $t_1 \sim t_2$.

Choose one representative from each equivalence class of $\sim$, and let $X' \subseteq X$ be the set of these representatives. Define $\rho : X \to X'$ by mapping each element of X to the representative of its equivalence class. We call ρ a *merge map* associated with $\boldsymbol{\tau}_1$ and $\boldsymbol{\tau}_2$.

Define the type family $\mathcal{T}_\rho$ on X' by

$$\mathcal{T}_\rho := \{(\rho(h), \rho(s), \rho(t), i) : (h, s, t, i) \in \mathcal{T}\}.$$

We call $\mathcal{T}_\rho$ the *ρ-merge type family* of $\mathcal{T}$.

Since $\boldsymbol{\tau}_1$ and $\boldsymbol{\tau}_2$ are distinct types in $\mathcal{T}$ but induce the same type in $\mathcal{T}_\rho$, we have $|\mathcal{T}_\rho| \leq |\mathcal{T}| - 1$. The next lemma records the two properties of this merge operation needed in the induction.

Lemma 5.6. *Let $\mathcal{T}$ be a type family on X, and let $\rho : X \to X'$ be a merge map associated with two types*

$$\boldsymbol{\tau}_1 = (h', s_1, t_1, j) \qquad \textit{and} \qquad \boldsymbol{\tau}_2 = (h', s_2, t_2, j)$$

in $\mathcal{T}$ with $(s_1, t_1) \neq (s_2, t_2)$. Then the following hold.

(i) *The type family $\mathcal{T}$ is layer-realizable if and only if $\mathcal{T}_\rho$ is layer-realizable.*

(ii) *If $\overline{\mathcal{S}} = (\overline{\mathcal{C}}, \overline{g}_1, \overline{g}_2, \overline{g}_3)$ is a color-system such that $\mathcal{T}_\rho$ is not compatible with $\overline{\mathcal{S}}$, then there exists a color-system $\mathcal{S} = (\mathcal{C}, g_1, g_2, g_3)$ such that $\mathcal{T}$ is not compatible with $\mathcal{S}$ and $|\mathcal{C}| = |\overline{\mathcal{C}}|^3$.*

Proof. For (i), first suppose that $f' : X' \to \mathbb{N}$ is a layer realization of $\mathcal{T}_\rho$, and put $f := f' \circ \rho$. For every $(h, s, t, i) \in \mathcal{T}$, the type $(\rho(h), \rho(s), \rho(t), i)$ belongs to $\mathcal{T}_\rho$. Condition (P1) for f' therefore gives $f(h) > f(s)$ and $f(h) > f(t)$.

Now let $(h, s, t, i), (\widetilde{h}, \widetilde{s}, \widetilde{t}, \widetilde{i}) \in \mathcal{T}$ satisfy $f(h) = f(\widetilde{h})$. Then $f'(\rho(h)) = f'(\rho(\widetilde{h}))$, so condition (P2) applied to the corresponding types in $\mathcal{T}_\rho$ gives

$$i = \widetilde{i}, \qquad f(s) = f(\widetilde{s}), \qquad \text{and} \qquad f(t) = f(\widetilde{t}).$$

Thus f is a layer realization of $\mathcal{T}$.

Conversely, suppose that $f : X \to \mathbb{N}$ is a layer realization of $\mathcal{T}$. Applying condition (P2) to $\boldsymbol{\tau}_1$ and $\boldsymbol{\tau}_2$ gives

$$f(s_1) = f(s_2) \qquad \text{and} \qquad f(t_1) = f(t_2).$$

Since $\sim$ is generated by $s_1 \sim s_2$ and $t_1 \sim t_2$, the map f is constant on every equivalence class of $\sim$. Hence the map $f' : X' \to \mathbb{N}$ defined by

$$f'(\rho(x)) := f(x) \qquad (x \in X)$$

is well-defined. Choosing preimages in $\mathcal{T}$ for types in $\mathcal{T}_\rho$, conditions (P1) and (P2) for f' follow directly from the corresponding conditions for f. Thus f' is a layer realization of $\mathcal{T}_\rho$.

For (ii), choose $c_* \in \overline{\mathcal{C}}$ and put $\mathcal{C} := \overline{\mathcal{C}}^3$. Write $\boldsymbol{a} = (a_0, a_1, a_2)$ and $\boldsymbol{b} = (b_0, b_1, b_2)$ for elements of $\mathcal{C}$. For each $i \in [3]$, define

$$g_i(\boldsymbol{a}, \boldsymbol{b}) := \begin{cases} (\overline{g}_i(a_0, b_0), a_0, b_0), & \text{if } i = j, \\ (\overline{g}_i(a_0, b_0), c_*, c_*), & \text{if } i \neq j. \end{cases}$$

Let $\mathcal{S} := (\mathcal{C}, g_1, g_2, g_3)$. Clearly $|\mathcal{C}| = |\overline{\mathcal{C}}|^3$.

Suppose for a contradiction that $\mathcal{T}$ is compatible with $\mathcal{S}$, and let $\varphi : X \to \mathcal{C}$ be a compatible color function. Let $\varphi_0 : X \to \overline{\mathcal{C}}$ be its first-coordinate projection. Taking first coordinates in the compatibility equation for each $(h, s, t, i) \in \mathcal{T}$ gives $\varphi_0(h) = \overline{g}_i(\varphi_0(s), \varphi_0(t))$. For the two selected types, compatibility also gives

$$g_j(\varphi(s_1), \varphi(t_1)) = \varphi(h') = g_j(\varphi(s_2), \varphi(t_2)).$$

Comparing the second and third coordinates yields $\varphi_0(s_1) = \varphi_0(s_2)$ and $\varphi_0(t_1) = \varphi_0(t_2)$. Hence φ_0 is constant on every equivalence class of $\sim$, and therefore induces a well-defined map $\overline{\varphi} : X' \to \overline{\mathcal{C}}$ given by

$$\overline{\varphi}(\rho(x)) := \varphi_0(x) \qquad (x \in X).$$

For every $(h, s, t, i) \in \mathcal{T}$, we have

$$\overline{\varphi}(\rho(h)) = \overline{g}_i(\overline{\varphi}(\rho(s)), \overline{\varphi}(\rho(t))).$$

By the definition of $\mathcal{T}_\rho$, this means that $\overline{\varphi}$ is a compatible color function for $\mathcal{T}_\rho$ and $\overline{\mathcal{S}}$, a contradiction. □

Now, we give the proof of Lemma 5.4.

Proof of Lemma 5.4. We argue by induction on $n := |\mathcal{T}|$. If $n = 0$, then $\mathcal{T}$ is empty and hence layer-realizable, so there is nothing to prove. Let $n \geq 1$, and assume that the assertion holds for every type family with fewer than n types.

We now consider the following three cases.

Case 1: a position conflict. Suppose that $\mathcal{T}$ contains two types $\boldsymbol{\tau}_1 = (h', s_1, t_1, j_1)$ and $\boldsymbol{\tau}_2 = (h', s_2, t_2, j_2)$ with $j_1 \neq j_2$. Let $\mathcal{C} := \{0, 1\}$, let g_{j_2} be constantly 1, and let g_q be constantly 0 for every $q \in [3] \setminus \{j_2\}$. Put $\mathcal{S} := (\mathcal{C}, g_1, g_2, g_3)$. If $\varphi : X \to \mathcal{C}$ were a compatible color function, then the compatibility equations for $\boldsymbol{\tau}_1$ and $\boldsymbol{\tau}_2$ would give $\varphi(h') = 0$ and $\varphi(h') = 1$, respectively, a contradiction. Thus $\mathcal{T}$ is not compatible with $\mathcal{S}$, and $|\mathcal{C}| = 2 \leq 3^{3^n}$.

Case 2: a merge reduction. Suppose that Case 1 does not occur, but $\mathcal{T}$ contains two types $\boldsymbol{\tau}_1 = (h', s_1, t_1, j)$ and $\boldsymbol{\tau}_2 = (h', s_2, t_2, j)$ with $(s_1, t_1) \neq (s_2, t_2)$. Let $\rho : X \to X'$ be a merge map associated with these two types. Since they induce the same type in $\mathcal{T}_\rho$, we have $|\mathcal{T}_\rho| \leq n - 1$. Moreover, part (i) of Lemma 5.6 implies that $\mathcal{T}_\rho$ is not layer-realizable. By the induction hypothesis, there is a color-system $\overline{\mathcal{S}}$ with color set $\overline{\mathcal{C}}$ such that $\mathcal{T}_\rho$ is not compatible with $\overline{\mathcal{S}}$ and $|\overline{\mathcal{C}}| \leq 3^{3^{|\mathcal{T}_\rho|}}$. Part (ii) of Lemma 5.6 then yields a color-system $\mathcal{S}$ with color set $\mathcal{C}$ such that $\mathcal{T}$ is not compatible with $\mathcal{S}$ and

$$|\mathcal{C}| = |\overline{\mathcal{C}}|^3 \leq \left(3^{3^{|\mathcal{T}_\rho|}}\right)^3 = 3^{3^{|\mathcal{T}_\rho|+1}} \leq 3^{3^n}.$$

Case 3: a dependency cycle. Suppose that neither Case 1 nor Case 2 occurs. Thus any two types with the same head have the same position and the same ordered pair of tails. Define the *dependency*

digraph D on vertex set X by adding the arcs $h \to s$ and $h \to t$ for every $(h, s, t, i) \in \mathcal{T}$; a loop is regarded as a directed cycle.

If D were acyclic, a reverse ordering of D would give an injective map $f : X \to [|X|]$ such that $f(x) > f(y)$ for every arc $x \to y$. Thus condition (P1) would hold. Moreover, if $(h, s, t, i), (\widetilde{h}, \widetilde{s}, \widetilde{t}, \widetilde{i}) \in \mathcal{T}$ satisfy $f(h) = f(\widetilde{h})$, then injectivity gives $h = \widetilde{h}$. Since neither Case 1 nor Case 2 occurs, we have $i = \widetilde{i}$ and $(s, t) = (\widetilde{s}, \widetilde{t})$, so condition (P2) also holds. Hence f would be a layer realization of $\mathcal{T}$, a contradiction. Therefore D contains a directed cycle

$$h_0 \longrightarrow h_1 \longrightarrow \cdots \longrightarrow h_{\ell-1} \longrightarrow h_0,$$

where $h_0, \ldots, h_{\ell-1}$ are distinct, $\ell \geq 1$, and all subscripts are read modulo ℓ.

For each $r \in \{0, \ldots, \ell - 1\}$, choose a type $\boldsymbol{\tau}_r = (h_r, s_r, t_r, i_r) \in \mathcal{T}$ witnessing the arc $h_r \to h_{r+1}$. Choose $\sigma_r \in \{1, 2\}$ so that $h_{r+1} = s_r$ if $\sigma_r = 1$ and $h_{r+1} = t_r$ if $\sigma_r = 2$; if both equalities hold, choose either value. Since the heads $h_0, \ldots, h_{\ell-1}$ are distinct, these selected types are distinct, and hence $\ell \leq n$.

Set $\mathcal{C} := (\mathbb{Z}/(\ell + 1)\mathbb{Z})^\ell$, with coordinates indexed by $0, \ldots, \ell - 1$. For $\boldsymbol{a}, \boldsymbol{b} \in \mathcal{C}$ and $q \in [3]$, define $g_q(\boldsymbol{a}, \boldsymbol{b})$ coordinatewise by

$$\big(g_q(\boldsymbol{a}, \boldsymbol{b})\big)_r = \begin{cases} a_{r+1} + 1, & \text{if } q = i_r \text{ and } \sigma_r = 1, \\ b_{r+1} + 1, & \text{if } q = i_r \text{ and } \sigma_r = 2, \\ 0, & \text{if } q \neq i_r, \end{cases} \tag{5.4}$$

where the coordinate subscripts are read modulo ℓ and all arithmetic is performed in $\mathbb{Z}/(\ell + 1)\mathbb{Z}$. Put $\mathcal{S} := (\mathcal{C}, g_1, g_2, g_3)$.

Suppose for a contradiction that $\varphi : X \to \mathcal{C}$ is a compatible color function. Considering the compatibility equation for $\boldsymbol{\tau}_r$ in coordinate r gives

$$\big(\varphi(h_r)\big)_r = \big(\varphi(h_{r+1})\big)_{r+1} + 1 \qquad (0 \leq r < \ell).$$

Chaining these equations around the cycle yields

$$\big(\varphi(h_0)\big)_0 = \big(\varphi(h_0)\big)_0 + \ell \qquad \text{in } \mathbb{Z}/(\ell + 1)\mathbb{Z},$$

which is impossible since $\ell \not\equiv 0 \pmod{\ell + 1}$. Therefore $\mathcal{T}$ is not compatible with $\mathcal{S}$. Finally, since $1 \leq \ell \leq n$,

$$|\mathcal{C}| = (\ell + 1)^\ell \leq (n + 1)^n \leq 3^{n^2} \leq 3^{3^n}.$$

In all three cases, we obtain a color-system $\mathcal{S}$ such that $\mathcal{T}$ is not compatible with $\mathcal{S}$ and $|\mathcal{C}| \leq 3^{3^{|\mathcal{T}|}}$. This completes the proof. □

6 Non-jump at zero

In this section we prove Theorem 1.10. We first define the palette density corresponding to the vertex-pair density notion. For a palette $\mathscr{P} = (\mathscr{C}, \mathscr{A})$, a position $r \in [3]$, and a color $a \in \mathscr{C}$, define

$$D_r^{\mathscr{P}}(a) := \big|\{(x_1, x_2, x_3) \in \mathscr{A} : x_r = a\}\big|.$$

The *vertex-pair minimum degree* of $\mathscr{P}$ is

$$\delta_{\dot{\cdot}}(\mathscr{P}) := \frac{1}{|\mathscr{C}|^2} \min_{a \in \mathscr{C}} \min_{r \in [3]} D_r^{\mathscr{P}}(a). \tag{6.1}$$

For a 3-graph F, set

$$\pi_{\therefore}^{\mathsf{pal}}(F) = \sup\{\delta_{\therefore}(\mathscr{P}) : F \text{ does not admit } \mathscr{P}\}.$$

We shall use Lamaison's palette theorem for the vertex-pair parameter.

Theorem 6.1 ([21, Theorem 5.1])**.** *For every 3-graph F, we have $\pi_{\therefore}(F) = \pi_{\therefore}^{\mathsf{pal}}(F)$.*

For palettes $\mathscr{P} = (\mathscr{C}_P, \mathscr{A}_P)$ and $\mathscr{Q} = (\mathscr{C}_Q, \mathscr{A}_Q)$, a map $f : \mathscr{C}_P \to \mathscr{C}_Q$ is a *palette homomorphism* if

$$(f(x), f(y), f(z)) \in \mathscr{A}_Q \qquad \text{for every } (x, y, z) \in \mathscr{A}_P.$$

We write $\mathscr{P} \preceq \mathscr{Q}$ when such a map exists. Palette admission is monotone under this relation: if a 3-graph admits $\mathscr{P}$ and $\mathscr{P} \preceq \mathscr{Q}$, then it also admits $\mathscr{Q}$. For a palette $\mathscr{P} = (\mathscr{C}_P, \mathscr{A}_P)$, put

$$\mathrm{rev}(\mathscr{P}) := (\mathscr{C}_P, \{(z, y, x) : (x, y, z) \in \mathscr{A}_P\}).$$

Reversing the witnessing vertex order shows that a 3-graph admits $\mathscr{P}$ if and only if it admits $\mathrm{rev}(\mathscr{P})$. Consequently, if some 3-graph admits $\mathscr{P}$ but does not admit $\mathscr{Q}$, then necessarily $\mathscr{P} \not\preceq \mathscr{Q}$ and $\mathscr{P} \not\preceq \mathrm{rev}(\mathscr{Q})$. The following palette separation theorem states that these two necessary conditions are also sufficient.

Theorem 6.2 ([20, Theorem 13])**.** *Let $\mathscr{P}$ and $\mathscr{Q}$ be palettes. There exists a 3-graph F which admits $\mathscr{P}$ but does not admit $\mathscr{Q}$ if and only if $\mathscr{P} \not\preceq \mathscr{Q}$ and $\mathscr{P} \not\preceq \mathrm{rev}(\mathscr{Q})$.*

We now define the palette product construction used below.

Definition 6.3 (Palette product)**.** Let I be a nonempty finite set, and for each $i \in I$ let $\mathscr{P}_i = (\mathscr{C}_i, \mathscr{A}_i)$ be a palette. Their *palette product* is the palette

$$\mathscr{P}_\times := \prod_{i \in I} \mathscr{P}_i = (\mathscr{C}_\times, \mathscr{A}_\times), \qquad \mathscr{C}_\times := \prod_{i \in I} \mathscr{C}_i,$$

where for product colors $\boldsymbol{x} = (x_i)_{i \in I}$, $\boldsymbol{y} = (y_i)_{i \in I}$ and $\boldsymbol{z} = (z_i)_{i \in I}$,

$$(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z}) \in \mathscr{A}_\times \quad \Longleftrightarrow \quad (x_i, y_i, z_i) \in \mathscr{A}_i \text{ for every } i \in I.$$

For every $i \in I$, the coordinate projection $\mathscr{C}_\times \to \mathscr{C}_i$ is a palette homomorphism, and hence $\mathscr{P}_\times \preceq \mathscr{P}_i$. Moreover, for every $r \in [3]$ and every product color $\boldsymbol{a} = (a_i)_{i \in I}$, $D_r^{\mathscr{P}_\times}(\boldsymbol{a}) = \prod_{i \in I} D_r^{\mathscr{P}_i}(a_i)$. It follows that

$$\delta_{\therefore}(\mathscr{P}_\times) \geq \prod_{i \in I} \delta_{\therefore}(\mathscr{P}_i). \tag{6.2}$$

The next lemma reduces all palettes above a fixed vertex-pair minimum-degree threshold to palettes on a bounded color set.

Lemma 6.4. *For every $0 < \eta < 1$, there exists an integer $n = n(\eta)$ with the following property. If $\mathscr{R} = (\mathscr{C}_R, \mathscr{A}_R)$ is a palette satisfying $\delta_{\therefore}(\mathscr{R}) > \eta$, then there exists a palette $\mathscr{S} = ([n], \mathscr{A}_S)$ such that $\delta_{\therefore}(\mathscr{S}) \geq \eta/2$ and $\mathscr{S} \preceq \mathscr{R}$.*

Proof. Choose $n = n(\eta) \geq 12$ so large that $3n \exp(-\eta^2 n/32) < 1$. Select a random map $f : [n] \to \mathscr{C}_R$ by choosing $f(1), \ldots, f(n)$ independently and uniformly from $\mathscr{C}_R$. Define the palette $\mathscr{S}_f = ([n], \mathscr{A}_f)$ by

$$(i, j, k) \in \mathscr{A}_f \quad \Longleftrightarrow \quad (f(i), f(j), f(k)) \in \mathscr{A}_R.$$

Then f is a palette homomorphism from $\mathscr{S}_f$ to $\mathscr{R}$, and hence $\mathscr{S}_f \preceq \mathscr{R}$. It remains to show that $\delta_{\therefore}(\mathscr{S}_f) \geq \eta/2$ with positive probability.

Fix $i \in [n]$. Let $Z_i^{(1)}$ count the ordered pairs (j,k) of distinct elements of $[n] \setminus \{i\}$ for which $(i,j,k) \in \mathscr{A}_f$. Conditional on $f(i) = a$, every such pair is counted with probability $\frac{D_1^{\mathscr{R}}(a)}{|\mathscr{C}_R|^2} \geq \eta$. Hence

$$\mathbb{E}[Z_i^{(1)} \mid f(i) = a] \geq \eta(n-1)(n-2) \geq \frac{3\eta n^2}{4},$$

where the last inequality uses $n \geq 12$.

Still conditioning on $f(i)$, the variable $Z_i^{(1)}$ depends on the independent random variables $f(\ell)$, $\ell \neq i$. Changing one of them affects at most $2(n-2) \leq 2n$ counted pairs. Therefore, McDiarmid's bounded-differences inequality gives

$$\Pr(Z_i^{(1)} < \eta n^2/2 \mid f(i)) \leq \exp\left(-\frac{2(\eta n^2/4)^2}{(n-1)(2n)^2}\right) \leq \exp(-\eta^2 n/32).$$

The same estimate holds for the analogous variables $Z_i^{(2)}$ and $Z_i^{(3)}$, in which i occupies the second and third positions, respectively. By the union bound and the choice of n, with positive probability

$$Z_i^{(r)} \geq \eta n^2/2 \qquad \text{for every } i \in [n] \text{ and } r \in [3].$$

Since $D_r^{\mathscr{S}_f}(i) \geq Z_i^{(r)}$, this event implies $\delta_{\therefore}(\mathscr{S}_f) \geq \eta/2$. Fixing such an outcome of f and taking $\mathscr{S} := \mathscr{S}_f$ proves the lemma. □

Fix $0 < \eta < 1$, and let $n = n(\eta)$ be supplied by Lemma 6.4. Let $\mathfrak{P}_\eta$ be the collection of all palettes $\mathscr{S} = ([n], \mathscr{A}_S)$ satisfying $\delta_{\therefore}(\mathscr{S}) \geq \eta/2$. This collection is finite and nonempty, since it contains the full palette $([n], [n]^3)$. We define the *universal palette*

$$\mathscr{U}_\eta := \prod_{\mathscr{S} \in \mathfrak{P}_\eta} \mathscr{S}.$$

By (6.2), we have $\delta_{\therefore}(\mathscr{U}_\eta) \geq (\eta/2)^{|\mathfrak{P}_\eta|} > 0$. Moreover, if $\mathscr{R}$ is any palette with $\delta_{\therefore}(\mathscr{R}) > \eta$, then Lemma 6.4 gives some $\mathscr{S} \in \mathfrak{P}_\eta$ with $\mathscr{S} \preceq \mathscr{R}$, while the corresponding coordinate projection gives $\mathscr{U}_\eta \preceq \mathscr{S}$. Thus

$$\mathscr{U}_\eta \preceq \mathscr{R} \qquad \text{whenever } \delta_{\therefore}(\mathscr{R}) > \eta. \tag{6.3}$$

We now introduce a family of special *affine palettes* with positive minimum codegree. For every fixed palette with positive vertex-pair minimum degree, a suitable member of this family can be separated from it. This is the key ingredient in the proof of Theorem 1.10.

Definition 6.5 (Affine palette)**.** For a prime $p > 2$, the *affine palette over* $\mathbb{F}_p$ is

$$\mathscr{Q}_p = (\mathscr{C}_p, \mathscr{A}_p), \qquad \mathscr{C}_p := \mathbb{F}_p, \qquad \mathscr{A}_p := \{(x,y,z) \in \mathbb{F}_p^3 : x + y - 2z = 1\}.$$

Since the coefficients $1, 1, -2$ are non-zero in $\mathbb{F}_p$, fixing any two coordinates determines the remaining coordinate uniquely, whereas fixing any one coordinate leaves exactly p choices for the other two. Hence

$$\delta_{\wedge}(\mathscr{Q}_p) = \delta_{\therefore}(\mathscr{Q}_p) = \frac{1}{p}. \tag{6.4}$$

Lemma 6.6. *Let $0 < \eta < 1$, and let $\mathscr{P} = (\mathscr{C}, \mathscr{A})$ be a palette with $\delta_{\therefore}(\mathscr{P}) > 0$. Then there exists a prime $p > \max\{2, 1/\eta\}$ such that*

$$\mathscr{P} \not\preceq \mathscr{Q}_p \qquad \text{and} \qquad \mathscr{P} \not\preceq \mathrm{rev}(\mathscr{Q}_p).$$

Proof. First consider the system of affine equations over $\mathbb{Q}$

$$X_a + X_b - 2X_c = 1 \qquad ((a, b, c) \in \mathscr{A}), \tag{6.5}$$

with one unknown X_x for each $x \in \mathscr{C}$. We claim that this system has no solution over $\mathbb{Q}$. Indeed, suppose that $X : \mathscr{C} \to \mathbb{Q}$ is a solution, and choose $c_0 \in \mathscr{C}$ with $X_{c_0} = \max_{x \in \mathscr{C}} X_x$. Since $\delta_{\therefore}(\mathscr{P}) > 0$, the color c_0 appears as the third coordinate of some admissible triple. Thus there are $a, b \in \mathscr{C}$ such that $(a, b, c_0) \in \mathscr{A}$. The corresponding equation gives

$$X_a + X_b - 2X_{c_0} = 1.$$

But $X_a \le X_{c_0}$ and $X_b \le X_{c_0}$, so the left-hand side is at most 0, a contradiction.

We now pass from $\mathbb{Q}$ to $\mathbb{F}_p$. We use the following standard consequence of Gaussian elimination. If a finite linear system with integer coefficients has no solution over $\mathbb{Q}$, then there is an integer linear combination of its equations whose left-hand side is identically 0 and whose right-hand side is a non-zero integer. To see this, perform Gaussian elimination over $\mathbb{Q}$ on the augmented matrix; since the system is inconsistent, one obtains a row $0 \ \cdots \ 0 \mid r$ with $r \in \mathbb{Q}^\times$. This row is a rational linear combination of the original equations, and clearing denominators gives the desired integer linear combination.

Applying this observation to (6.5), we obtain integers γ_{abc}, indexed by $(a, b, c) \in \mathscr{A}$, such that

$$\sum_{(a,b,c) \in \mathscr{A}} \gamma_{abc}(X_a + X_b - 2X_c) = 0$$

as a formal linear expression, while $N_1 := \sum_{(a,b,c) \in \mathscr{A}} \gamma_{abc}$ is a non-zero integer. Therefore, if $p \nmid N_1$, then (6.5) has no solution over $\mathbb{F}_p$: otherwise, reducing the above integer linear combination modulo p would give $0 = N_1 \pmod{p}$, a contradiction. A map $f : \mathscr{C} \to \mathbb{F}_p$ is a palette homomorphism from $\mathscr{P}$ to $\mathscr{Q}_p$ if and only if the assignment $X_x = f(x)$ gives a solution of (6.5) over $\mathbb{F}_p$. Hence $\mathscr{P} \not\preceq \mathscr{Q}_p$ whenever $p \nmid N_1$.

It remains to exclude homomorphisms from $\mathscr{P}$ to $\mathrm{rev}(\mathscr{Q}_p)$. A map $f : \mathscr{C} \to \mathbb{F}_p$ is such a homomorphism if and only if the assignment $X_x = f(x)$ solves the reverse system over $\mathbb{F}_p$:

$$X_c + X_b - 2X_a = 1 \qquad ((a, b, c) \in \mathscr{A}). \tag{6.6}$$

The same maximum argument shows that this system has no solution over $\mathbb{Q}$. Indeed, if X were a rational solution, choose $a_0 \in \mathscr{C}$ maximizing X. Since $\delta_{\therefore}(\mathscr{P}) > 0$, there are $b, c \in \mathscr{C}$ such that $(a_0, b, c) \in \mathscr{A}$. Then $X_c + X_b - 2X_{a_0} = 1$, while the left-hand side is at most 0, a contradiction.

Applying the same Gaussian-elimination observation to (6.6), we obtain a non-zero integer N_2 such that the reverse system has no solution over $\mathbb{F}_p$ whenever $p \nmid N_2$. Hence $\mathscr{P} \not\preceq \mathrm{rev}(\mathscr{Q}_p)$ for all such primes p.

Finally, choose a prime $p > \max\{2, 1/\eta\}$ such that $p \nmid N_1 N_2$. This is possible because only finitely many primes divide the non-zero integer $N_1 N_2$. This proves the lemma. □

Proof of Theorem 1.10. Put $\eta := \min\{\varepsilon, 1/2\}$ and construct the universal palette $\mathscr{U}_\eta$ above. Since $\delta_{\therefore}(\mathscr{U}_\eta) > 0$, Lemma 6.6 gives a prime $p > \max\{2, 1/\eta\}$ such that

$$\mathscr{U}_\eta \not\preceq \mathscr{Q}_p \qquad \text{and} \qquad \mathscr{U}_\eta \not\preceq \mathrm{rev}(\mathscr{Q}_p).$$

By (6.4), $0 < \delta_{\wedge}(\mathscr{Q}_p) = \delta_{\therefore}(\mathscr{Q}_p) = 1/p < \eta$. Therefore, Theorem 6.2 gives a 3-graph F_ε which admits $\mathscr{U}_\eta$ but does not admit $\mathscr{Q}_p$.

We first establish the upper bound. Let $\mathscr{R}$ be any palette with $\delta_{\therefore}(\mathscr{R}) > \eta$. By (6.3), $\mathscr{U}_\eta \preceq \mathscr{R}$. Since F_ε admits $\mathscr{U}_\eta$, it also admits $\mathscr{R}$. Thus every palette not admitted by F_ε has vertex-pair minimum degree at most η, and so $\pi^{\mathsf{pal}}_{\therefore}(F_\varepsilon) \le \eta$. By Theorem 6.1, we have $\pi_{\therefore}(F_\varepsilon) \le \eta \le \varepsilon$.

On the other hand, F_ε does not admit $\mathscr{Q}_p$. Hence, by (6.4), we have $\pi^{\mathsf{pal}}_{\wedge}(F_\varepsilon) \ge \delta_{\wedge}(\mathscr{Q}_p) = \frac{1}{p}$. It follows from Lemma 2.5 that $\pi_{\wedge}(F_\varepsilon) \ge 1/p$. Combining the two bounds with the general inequality $\pi_{\wedge}(F_\varepsilon) \le \pi_{\therefore}(F_\varepsilon)$ gives

$$0 < \frac{1}{p} \le \pi_{\wedge}(F_\varepsilon) \le \pi_{\therefore}(F_\varepsilon) \le \eta \le \varepsilon,$$

as required. □

7 Concluding remarks

To conclude this paper, we first provide an explicit family of 3-graphs that separates the cherry and vertex-pair Turán densities as an application of Theorem 1.8, and then discuss further directions concerning related positive densities.

7.1 An explicit family

The following family gives an explicit separation between the cherry and vertex-pair Turán densities.

Lemma 7.1. *For every $r \ge 4$, let*

$$S_r^+ = \Big([r], \{1ij : 2 \le i < j \le r\} \cup \{234\}\Big).$$

Then we have $\pi_{\wedge}(S_r^+) = 0$ and $\pi_{\therefore}(S_r^+) \ge \frac{1}{2}$.

Proof. The 3-graph S_r^+ is the full 3-uniform star with center 1, together with the additional edge 234; in particular, $S_4^+ = K_4^{(3)}$.

We first construct a max-layered witness for S_r^+. Let $\prec$ be the natural order $1 \prec 2 \prec \cdots \prec r$. Let

$$\xi : \{ij : 2 \le i < j \le r,\ j \ge 5\} \longrightarrow \{4, 5, 6, \ldots\}$$

be injective, and define $\lambda : \partial S_r^+ \to \mathbb{N}$ by

$$\lambda(12) = 2, \qquad \lambda(34) = 3, \qquad \lambda(1i) = 1 \ \ (3 \le i \le r), \qquad \lambda(23) = \lambda(24) = 1,$$

and

$$\lambda(ij) = \xi(ij) \qquad (2 \le i < j \le r,\ j \ge 5).$$

With respect to $(\prec, \lambda)$, the edges 123 and 124 have the same edge-pattern $(2, 1, 1)$, while the edges 134 and 234 have the same edge-pattern $(1, 1, 3)$. Every remaining edge has the form $1ij$ with $j \ge 5$, and its edge-pattern is $(\lambda(1i), \lambda(1j), \xi(ij))$. Since $\xi(ij) > 3$, its third entry is the unique maximum entry. The two edge-patterns $(2, 1, 1)$ and $(1, 1, 3)$ have maximum entries 2 and 3, respectively, while the injectivity of ξ ensures that the remaining edge-patterns have pairwise distinct maximum entries. Hence the maximum entry of every edge-pattern in $\mathcal{B}_{\lambda,\prec}$ occurs in a unique position, and distinct edge-patterns in $\mathcal{B}_{\lambda,\prec}$ have distinct maximum entries. Thus $\mathcal{B}_{\lambda,\prec}$ is max-degenerate, so $(\prec, \lambda)$ is a max-layered witness for S_r^+. By Theorem 1.8, $\pi_{\wedge}(S_r^+) = 0$.

For the vertex-pair Turán density, observe that $K_4^{(3)} \subseteq S_r^+$ on the vertex set $\{1,2,3,4\}$. Since $F \subseteq G$ implies $\pi_{\therefore}(F) \leq \pi_{\therefore}(G)$, the theorem of Reiher, Rödl and Schacht [28] gives

$$\pi_{\therefore}(S_r^+) \geq \pi_{\therefore}(K_4^{(3)}) = \frac{1}{2},$$

completing the proof. □

For $r \geq 5$, this family does not arise from the known example $K_4^{(3)}$ through the blow-up construction. Indeed, S_r^+ is not contained in any blow-up of $K_4^{(3)}$: the shadow of every 3-graph contained in such a blow-up is 4-colorable, whereas $\partial S_r^+ = K_r$ has chromatic number $r > 4$. Thus the family $(S_r^+)_{r\geq 4}$ gives a separation between the cherry and vertex-pair parameters with unbounded shadow chromatic number. Moreover, the equality $\pi_{\wedge}(S_r^+) = 0$ cannot be deduced from $\pi_{\wedge}(K_4^{(3)}) = 0$ by monotonicity, since $K_4^{(3)} \subseteq S_r^+$ only yields $\pi_{\wedge}(K_4^{(3)}) \leq \pi_{\wedge}(S_r^+)$.

7.2 Further directions

By Theorem 1.10, for every $\varepsilon > 0$ there exists a finite 3-graph F_ε such that

$$0 < \pi_{\wedge}(F_\varepsilon) \leq \pi_{\therefore}(F_\varepsilon) \leq \varepsilon.$$

The proof constructs a universal palette as the product of finitely many palettes and then separates it from a suitable affine palette of positive minimum codegree. Theorem 6.2 consequently produces a finite 3-graph that admits the universal palette but does not admit the affine palette. Although this proves that positive cherry Turán densities occur arbitrarily close to zero, it does not yield an explicit family. It would therefore be interesting to find a natural sequence $(G_m)_{m\geq 1}$ of small or simply described 3-graphs such that

$$0 < \pi_{\wedge}(G_m) \longrightarrow 0.$$

Such a sequence would give a more explicit description of the positive part of the cherry Turán density near zero.

AI Declaration

All mathematical ideas, arguments, results, and proofs were developed by the authors. AI tools were only used for language polishing and proofreading in this paper.

References

[1] R. Baber and J. Talbot. Hypergraphs do jump. *Combin. Probab. Comput.*, 20(2):161–171, 2011.

[2] S. Berger, S. Piga, C. Reiher, V. Rödl, and M. Schacht. Turán density of cliques of order five in 3-uniform hypergraphs with quasirandom links. *Trans. Amer. Math. Soc.*, 378(12):8283–8318, 2025.

[3] M. Bucić. The uniform Turán density of the tetrahedron. *arXiv preprint arXiv:2609.11802*, 2026.

[4] M. Bucić, J. W. Cooper, D. Král', S. Mohr, and D. Munhá Correia. Uniform Turán density of cycles. *Trans. Amer. Math. Soc.*, 376(7):4765–4809, 2023.

[5] A. Chen and B. Schülke. Beyond the broken tetrahedron. *Combin. Probab. Comput.*, 35(1):59–70, 2026.

[6] D. Conlon and B. Schülke. Hypergraphs accumulate. *International Mathematics Research Notices*, 2025(2):rnae289, 2025.

[7] P. Erdős. On extremal problems of graphs and generalized graphs. *Israel J. Math.*, 2:183–190, 1964.

[8] P. Erdős. Paul Turán, 1910–1976: his work in graph theory. *J. Graph Theory*, 1(2):97–101, 1977.

[9] P. Erdős. Problems and results on graphs and hypergraphs: similarities and differences. In *Mathematics of Ramsey theory*, volume 5 of *Algorithms Combin.*, pages 12–28. Springer, Berlin, 1990.

[10] P. Erdős and M. Simonovits. A limit theorem in graph theory. *Studia Sci. Math. Hungar.*, 1:51–57, 1966.

[11] P. Erdős and V. T. Sós. On Ramsey-Turán type theorems for hypergraphs. *Combinatorica*, 2(3):289–295, 1982.

[12] P. Erdős and A. H. Stone. On the structure of linear graphs. *Bull. Amer. Math. Soc.*, 52:1087–1091, 1946.

[13] P. Frankl and Z. Füredi. An exact result for 3-graphs. *Discrete Math.*, 50(2-3):323–328, 1984.

[14] F. Garbe, D. Il'kovič, D. Král', F. Kučerák, and A. Lamaison. Hypergraphs with uniform Turán density equal to 8/27. *Combin. Probab. Comput.*, 35(2):280–289, 2026.

[15] F. Garbe, D. Král', and A. Lamaison. Hypergraphs with minimum positive uniform Turán density. *Israel J. Math.*, 259(2):701–726, 2024.

[16] R. Glebov, D. Král', and J. Volec. A problem of Erdős and Sós on 3-graphs. *Israel J. Math.*, 211(1):349–366, 2016.

[17] P. Keevash. Hypergraph Turán problems. In *Surveys in combinatorics 2011*, volume 392 of *London Math. Soc. Lecture Note Ser.*, pages 83–139. Cambridge Univ. Press, Cambridge, 2011.

[18] B. Kielak, D. Král, A. Lamaison, H. Liu, X. Shu, and Z. Wu. Solution of uniform Turán's tetrahedron problem. *arXiv preprint arXiv:2609.08336*, 2026.

[19] D. King, S. Piga, M. Sales, and B. Schülke. On possible uniform Turán densities. *arXiv preprint arXiv: 2504.21220*, 2025.

[20] D. Král, F. Kučerák, A. Lamaison, and G. Tardos. Uniform Turán density–palette classification. *arXiv preprint arXiv:2505.17325*, 2025.

[21] A. Lamaison. Palettes determine uniform Turán density. *arXiv preprint arXiv:2408.09643*, 2024.

[22] A. Lamaison and Z. Wu. The uniform Turán density of large stars. *arXiv preprint arXiv:2409.03699*, 2024.

[23] H. Li, H. Lin, G. Wang, and W. Zhou. Hypergraphs with a quarter uniform Turán density. *J. Oper. Res. Soc. China*, 2025.

[24] H. Lin, G. Wang, W. Zhou, and Y. Zhou. Extremal problems in uniformly dense hypergraphs and digraphs. *arXiv preprint arXiv:2603.10766*, 2026.

[25] H. Lin and W. Zhou. Turán density of stars in uniformly dense hypergraphs. *arXiv preprint arXiv:2510.12576*, 2025.

[26] W. Mantel. Vraagstuk XXVIII. *Wiskundige Opgaven met de Oplossingen*, 10:60–61, 1907.

[27] C. Reiher. Extremal problems in uniformly dense hypergraphs. *European J. Combin.*, 88:103117, 22, 2020.

[28] C. Reiher, V. Rödl, and M. Schacht. Embedding tetrahedra into quasirandom hypergraphs. *J. Combin. Theory Ser. B*, 121:229–247, 2016.

[29] C. Reiher, V. Rödl, and M. Schacht. Hypergraphs with vanishing Turán density in uniformly dense hypergraphs. *J. Lond. Math. Soc. (2)*, 97(1):77–97, 2018.

[30] C. Reiher, V. Rödl, and M. Schacht. On a Turán problem in weakly quasirandom 3-uniform hypergraphs. *J. Eur. Math. Soc.*, 20(5):1139–1159, 2018.

[31] C. Reiher, V. Rödl, and M. Schacht. Some remarks on $\pi_{\wedge}$. In *Connections in discrete mathematics*, pages 214–239. Cambridge Univ. Press, Cambridge, 2018.

[32] V. Rödl. On universality of graphs with uniformly distributed edges. *Discrete Math.*, 59(1-2):125–134, 1986.

[33] M. Schacht. Restricted problems in extremal combinatorics. In *Proceedings of the International Congress of Mathematicians 2022*, volume 6, pages 4646–4658. EMS Press, 2022.

[34] P. Turán. Eine Extremalaufgabe aus der Graphentheorie. *Mat. Fiz. Lapok*, 48:436–452, 1941.

[35] P. Turán. On the theory of graphs. *Colloq. Math.*, 3:19–30, 1954.