Three-Chromatic Geometric Hypergraphs

Gábor Damásdi[1]
MTA-ELTE Lendület Combinatorial Geometry Research Group, Dept. of Computer Science, ELTE Eötvös Loránd University, Budapest, Hungary

Dömötör Pálvölgyi
MTA-ELTE Lendület Combinatorial Geometry Research Group, Dept. of Computer Science, ELTE Eötvös Loránd University, Budapest, Hungary

Abstract

We prove that for any planar convex body C there is a positive integer m with the property that any finite point set P in the plane can be three-colored such that there is no translate of C containing at least m points of P, all of the same color. As a part of the proof, we show a strengthening of the Erdős-Sands-Sauer-Woodrow conjecture. Surprisingly, the proof also relies on the two dimensional case of the Illumination conjecture.

2012 ACM Subject Classification Mathematics of computing → Hypergraphs

Keywords and phrases Discrete geometry, Geometric hypergraph coloring, Decomposition of multiple coverings

Digital Object Identifier 10.4230/LIPIcs.SoCG.2022.32

Related Version *Full Version*: https://arxiv.org/abs/2112.01820

Funding *Gábor Damásdi*: Supported by the ÚNKP-21-3 New National Excellence Program of the Ministry for Innovation and Technology from the source of the National Research, Development and Innovation Fund, and both researchers supported by the Lendület program of the Hungarian Academy of Sciences (MTA), under grant number LP2017-19/2017.

1 Introduction

Our main result is the following.

▶ **Theorem 1.** *For any planar convex body C there is a positive integer $m = m(C)$ such that any finite point set P in the plane can be three-colored in a way that there is no translate of C containing at least m points of P, all of the same color.*

This result closes a long line of research about coloring points with respect to planar range spaces that consist of translates of a fixed set, a problem that was initiated by Pach over forty years ago [21]. In general, a pair $(P, \mathcal{S})$, where P is a set of points in the plane and $\mathcal{S}$ is a family of subsets of the plane, called the *range space*, defines a *primal* hypergraph $\mathcal{H}(P, \mathcal{S})$ whose vertex set is P, and for each $S \in \mathcal{S}$ we add the edge $S \cap P$ to the hypergraph. Given any hypergraph $\mathcal{H}$, a planar realization of $\mathcal{H}$ is defined as a pair $(P, \mathcal{S})$ for which $\mathcal{H}(P, \mathcal{S})$ is isomorphic to $\mathcal{H}$. If $\mathcal{H}$ can be realized with some pair $(P, \mathcal{S})$ where $\mathcal{S}$ is from some family $\mathcal{F}$, then we say that $\mathcal{H}$ is realizable with $\mathcal{F}$. The dual of the hypergraph $\mathcal{H}(P, \mathcal{S})$, where the elements of the range space $\mathcal{S}$ are the vertices and the points P define the edges such that $\{S \in \mathcal{S} \mid p \in S\}$ is an edge for every $p \in P$, is known as the *dual* hypergraph and is denoted by $\mathcal{H}(\mathcal{S}, P)$. If $\mathcal{H} = \mathcal{H}(\mathcal{S}, P)$ where $\mathcal{S}$ is from some family $\mathcal{F}$, then we say that $\mathcal{H}$

[1] Corresponding author

has a dual realization with $\mathcal{F}$. Pach observed [21, 24] that if $\mathcal{F}$ is the family of translates of some set, then $\mathcal{H}$ has a dual realization with $\mathcal{F}$ if and only if $\mathcal{H}$ has a (primal) realization with $\mathcal{F}$.

Pach proposed to study the chromatic number of hypergraphs realizable with different geometric families $\mathcal{F}$. It is important to distinguish between two types of hypergraph colorings that we will use, the *proper* coloring and the *polychromatic* coloring.

▶ **Definition 2.** *A hypergraph is* properly k-colorable *if its vertices can be colored with k colors so that each edge contains points from at least two color classes. Such a coloring is called a* proper k-coloring*. If a hypergraph has a proper k-coloring but not a proper $(k-1)$-coloring, then it is called* k-chromatic*.*

A hypergraph is polychromatic k-colorable *if its vertices can be colored with k colors so that each edge contains points from each color class. Such a coloring is called a* polychromatic k-coloring*.*

Note that for a polychromatic k-coloring to exist, it is necessary that each edge of the underlying hypergraph has at least k vertices. More generally, we say that a hypergraph is *m-heavy* if each of its edges has at least m vertices.

The main question that Pach raised can be rephrased as follows.

▶ **Question 3.** *For which planar families $\mathcal{F}$ is there an $m_k = m(\mathcal{F}, k)$ such that any m_k-heavy hypergraph realizable with $\mathcal{F}$ has a proper/polychromatic k-coloring?*

Initially, this question has been mainly studied for polychromatic k-colorings (known in case of a dual range space as *cover-decomposition* problem), and it was shown that such an m_k exists if $\mathcal{F}$ is the family of translates of some convex polygon [22, 33, 28], or the family of all halfplanes [14, 32], or the homothetic[2] copies of a triangle [15] or of a square [2], while it was also shown that even m_2 does not exist if $\mathcal{F}$ is the family of translates of some appropriate concave polygon [26, 27] or any body[3] with a smooth boundary [23]. It was also shown that there is no m_k for proper k-colorings if $\mathcal{F}$ is the family of all lines [26] or all axis-parallel rectangles [10]; for these families, the same holds in case of dual realizations [26, 25]. For homothets of convex polygons other than triangles, it is known that there is no m_2 for dual realizations [19], unlike for primal realizations. Higher dimensional variants [15, 8] and improved bounds for m_k have been also studied [3, 13, 7, 16, 4, 9]. For other results, see also the decade old survey [24], or the up-to-date website `https://coge.elte.hu/cogezoo.html`.

If $\mathcal{F}$ is the translates or homothets of some planar convex body, it is an easy consequence of the properties of generalized Delaunay-triangulations and the Four Color Theorem that any hypergraph realizable with $\mathcal{F}$ is proper 4-colorable if every edge contains at least two vertices. We have recently shown that this cannot be improved for homothets.

▶ **Theorem 4** (Damásdi, Pálvölgyi [12]). *Let C be any convex body in the plane that has two parallel supporting lines such that C is strictly convex in some neighborhood of the two points of tangencies. For any positive integer m, there exists a 4-chromatic m-uniform hypergraph that is realizable with homothets of C.*

[2] A *homothetic copy*, or *homothet*, is a scaled and translated (but non-rotated) copy of a set. We always require the scaling factor to be positive. Note that this is sometimes called a positive homothet.

[3] By *body*, we always mean a compact subset of the plane with a non-empty interior, though our results (and most of the results mentioned) also hold for sets that are unbounded, or that contain an arbitrary part of their boundary, and are thus neither open, nor closed. This is because a realization of a hypergraph can be perturbed slightly to move the points off from the boundaries of the sets realizing the respective edges of the hypergraph.

For translates, we recall the following result.

▶ **Theorem 5** (Pach, Pálvölgyi [23]). *Let C be any convex body in the plane that has two parallel supporting lines such that C is strictly convex in some neighborhood of the two points of tangencies. For any positive integer m, there exists a 3-chromatic m-uniform hypergraph that is realizable with translates of C.*

This left only the following question open: Is it true for any planar convex body C that there is a positive integer m such that no 4-chromatic m-uniform hypergraph is realizable with translates of C? Our Theorem 1 answers this question affirmatively for all C by showing that all realizable m-heavy hypergraphs are three-colorable for some m. This has been hitherto known to hold only when C is a polygon (in which case 2 colors suffice [28], and 3 colors are known to be enough even for homothets [18]) and pseudodisk families that intersect in a common point [1] (which generalizes the case when C is unbounded, in which case 2 colors suffice [23]).

The proof of Theorem 1 relies on a surprising connection with two other famous results, the solution of the two dimensional case of the Illumination conjecture [20], and a recent solution of the Erdős-Sands-Sauer-Woodrow conjecture by Bousquet, Lochet and Thomassé [6]. In fact, we need a generalization of the latter result, which we prove with the addition of one more trick to their method; this can be of independent interest.

Note that the extended abstract of our first proof attempt appeared recently in the proceedings of EuroComb 2021 [11]. That proof did not use the above two results, however, it only worked when C was a disk, and while the generalization to other convex bodies with a smooth boundary seemed feasible, we saw no way to extend it to arbitrary convex bodies.

The rest of the paper is organized as follows.
In Section 2 we present the three main ingredients of our proof:

- the Union Lemma (Section 2.1),
- the Erdős-Sands-Sauer-Woodrow conjecture (Section 2.2) – the proof of our generalization of the Bousquet-Lochet-Thomassé theorem can be found in the full version of the paper,
- the Illumination conjecture (Section 2.3), which is a theorem of Levi in the plane.

In Section 3 we give the detailed proof of Theorem 1.
In Section 4 we give a general overview of the steps of the algorithm requiring computation to show that we can find a three-coloring in randomized polynomial time.
Finally, in Section 5, we pose some problems left open.

2 Tools

2.1 Union Lemma

Polychromatic colorability is a much stronger property than proper colorability. Any polychromatic k-colorable hypergraph is proper 2-colorable. We generalize this trivial observation to the following statement about unions of polychromatic k-colorable hypergraphs.

▶ **Lemma 6** (Union Lemma). *Let $\mathcal{H}_1 = (V, E_1), \dots, \mathcal{H}_{k-1} = (V, E_{k-1})$ be hypergraphs on a common vertex set V. If $\mathcal{H}_1, \dots, \mathcal{H}_{k-1}$ are polychromatic k-colorable, then the hypergraph $\bigcup_{i=1}^{k-1} \mathcal{H}_i = (V, \bigcup_{i=1}^{k-1} E_i)$ is proper k-colorable.*

Proof. Choose $c(v) \in \{1, \dots, k\}$ such that it differs from each $c_i(v)$. We claim that c is a proper k-coloring of $\bigcup_{i=1}^{k-1} \mathcal{H}_i$. To prove this, it is enough to show that for every edge $H \in \mathcal{H}_i$ and for every color $j \in \{1, \dots, k-1\}$, there is a $v \in H$ such that $c(v) \neq j$. We can pick $v \in H$ for which $c_i(v) = j$. This finishes the proof. ◀

Lemma 6 is sharp in the sense that for every k there are $k-1$ hypergraphs such that each is polychromatic k-colorable but their union is not properly $(k-1)$-colorable.

We will apply the Union Lemma combined with the theorem below. A *pseudoline arrangement* is a collection of simple curves, each of which splits $\mathbb{R}^2$ into two unbounded parts, such that any two curves intersect at most once. A *pseudohalfplane* is the region on one side of a pseudoline in such an arrangement. For hypergraphs realizible by pseudohalfplanes the following was proved, generalizing a result of Smorodinsky and Yuditsky [32] about halfplanes.

▶ **Theorem 7** (Keszegh-Pálvölgyi [17]). *Any $(2k-1)$-heavy hypergraph realizable by pseudohalfplanes is polychromatic k-colorable, i.e., given a finite set of points and a pseudohalfplane arrangement in the plane, the points can be k-colored such that every pseudohalfplane that contains at least $2k-1$ points contains all k colors.*

Combining Theorem 7 with Lemma 6 for $k = 3$, we obtain the following.

▶ **Corollary 8.** *Any 5-heavy hypergraph realizable by two pseudohalfplane families is proper 3-colorable, i.e., given a finite set of points and two different pseudohalfplane arrangements in the plane, the points can be 3-colored such that every pseudohalfplane that contains at least 5 points contains two differently colored points.*

2.2 Erdős-Sands-Sauer-Woodrow conjecture

Given a quasi-order[4] $\prec$ on a set V, we interpret it as a digraph $D = (V, A)$, where the vertex set is V and a pair (x, y) defines an arc in A if $x \prec y$. The *closed in-neighborhood* of a vertex $x \in V$ is $N^-(x) = \{x\} \cup \{y | (y, x) \in A\}$. Similarly the *closed out-neighborhood* of a vertex x is $N^+(x) = \{x\} \cup \{y | (x, y) \in A\}$. We extend this to subsets $S \subset V$ as $N^-(S) = \bigcup_{x \in S} N^-(x)$ and $N^+(S) = \bigcup_{x \in S} N^+(x)$. A set of vertices S such that $N^+(S) = V$ is said to be *dominating.* For $A, B \subset V$ we will also say that A *dominates* B if $B \subset N^+(A)$.

A *complete multidigraph* is a digraph where parallel edges are allowed and in which there is at least one arc between each pair of distinct vertices. Let D be a complete multidigraph whose arcs are the disjoint union of k quasi-orders $\prec_1, \dots, \prec_k$ (parallel arcs are allowed). Define $N_i^-(x)$ (resp. $N_i^+(x)$) as the closed in-neighborhood (resp. out-neighborhood) of the digraph induced by $\prec_i$.

Proving the conjecture of Erdős, and of Sands, Sauer and Woodrow [31], Bousquet, Lochet and Thomassé recently showed the following.

▶ **Theorem 9** (Bousquet, Lochet, Thomassé [6]). *For every k, there exists an integer $f(k)$ such that if D is a complete multidigraph whose arcs are the union of k quasi-orders, then D has a dominating set of size at most $f(k)$.*

---

[4] A quasi-order $\prec$ is a reflexive and transitive relation, but it is not required to be antisymmetric, so $p \prec q \prec p$ is allowed, unlike for partial orders.

We show the following generalization of Theorem 9.

▶ **Theorem 10.** *For every pair of positive integers k and l, there exist an integer $f(k,l)$ such that if $D = (V, A)$ is a complete multidigraph whose arcs are the union of k quasi-orders $\prec_1, \dots, \prec_k$, then V contains a family of pairwise disjoint subsets S_i^j for $i \in [k]$, $j \in [l]$ with the following properties:*

- $|\bigcup_{i,j} S_i^j| \leq f(k,l)$
- *For each vertex $v \in V \setminus \bigcup_{i,j} S_i^j$ there is an $i \in [k]$ such that for each $j \in [l]$ there is an edge of $\prec_i$ from a vertex of S_i^j to v.*

Note that disjointness is the real difficulty here, without it the theorem would trivially hold from repeated applications of Theorem 9. We saw no way to derive Theorem 10 from Theorem 9, but with an extra modification the proof goes through. The proof of Theorem 10 can be found in the full version of the paper.

2.3 Hadwiger's Illumination conjecture and pseudolines

Hadwiger's Illumination conjecture has a number of equivalent formulations and names[5]. For a recent survey, see [5]. We will use the following version of the conjecture.

Let $\mathbb{S}^{d-1}$ denote the unit sphere in $\mathbb{R}^d$. For a convex body C, let ∂C denote the boundary of C and let $int(C)$ denote its interior. A direction (light) $u \in \mathbb{S}^{d-1}$ *illuminates* $b \in \partial C$ if $\{b + \lambda u : \lambda > 0\} \cap int(C) \neq \emptyset$.

▶ **Conjecture 11.** *The boundary of any convex body in $\mathbb{R}^d$ can be illuminated by 2^d or fewer directions. Furthermore, the 2^d lights are necessary if and only if the body is a parallelepiped.*

The conjecture is open in general. The $d = 2$ case was settled in affirmative by Levi [20] in 1955. For $d = 3$ the best result is due to Prymak [30], who showed that 16 lights are enough, improving the earlier method of Papadoperakis [29] with the help of a computer program.

In the following part we make an interesting connection between the Illumination conjecture for $d = 2$ and pseudolines. Roughly speaking, we show that the Illumination conjecture implies that for any convex body in the plane the boundary can be broken into three parts such that the translates of each part behave similarly to pseudolines, i.e., we get three pseudoline arrangements from the translates of the three parts.

To put this into precise terms, we need some technical definitions and statements. Fix a body C and an injective parametrization of ∂C, $\gamma : [0,1] \to \partial C$, that follows ∂C counterclockwise. For each point p of ∂C there is a set of possible tangents touching at p. Let $g(p) \subset \mathbb{S}^1$ denote the Gauss image of p, i.e., $g(p)$ is the set of unit outernormals of the tangent lines touching at p. Note that $g(p)$ is an arc of $\mathbb{S}^1$ and $g(p)$ is a proper subset of $\mathbb{S}^1$.

Let $g_+ : \partial C \to \mathbb{S}^1$ be the function that assigns to p the counterclockwise last element of $g(p)$. (See Figure 1 left.) Similarly let g_- be the function that assigns to p the clockwise last element of $g(p)$. Thus, $g(p)$ is the arc of $\mathbb{S}^1$ from $g_-(p)$ to $g_+(p)$. Let $|g(p)|$ denote the length of $g(p)$.

▶ **Observation 12.** *$g_+ \circ \gamma$ is continuous from the right and $g_- \circ \gamma$ is continuous from the left.*

[5] These include names such as Levi–Hadwiger Conjecture, Gohberg–Markus Covering Conjecture, Hadwiger Covering Conjecture, Boltyanski–Hadwiger Illumination Conjecture.

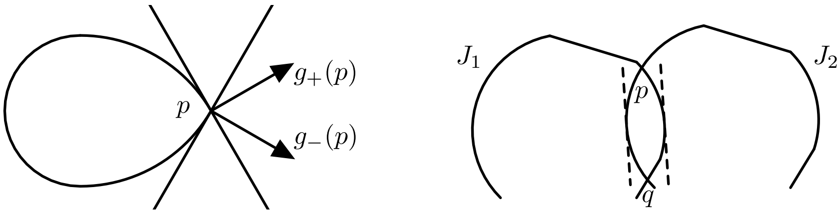

Figure 1 Extremal tangents at a boundary point (on the left) and parallel tangents on two intersecting translates (on the right).

For $t_1 < t_2$ let $\gamma_{[t_1,t_2]}$ denote the restriction of γ to the interval $[t_1, t_2]$. For $t_1 > t_2$ let $\gamma_{[t_1,t_2]}$ denote the concatenation of $\gamma_{[t_1,1]}$ and $\gamma_{[0,t_2]}$. When it leads to no confusion, we identify $\gamma_{[t_1,t_2]}$ with its image, which is a closed connected part of the boundary ∂C. For such a $J = \gamma_{[t_1,t_2]}$, let $g(J) = \bigcup_{p \in J} g(p)$. Clearly, $g(J)$ is an arc of $\mathbb{S}^1$ from $g_-(t_1)$ to $g_+(t_2)$; let $|g(J)|$ denote the length of this arc.

▶ **Lemma 13.** *Let C be a convex body and assume that J is a closed connected part of ∂C such that $|g(J)| < \pi$. Then there are no two translates of J that intersect in more than one point.*

Proof. Suppose J has two translates J_1 and J_2 such that they intersect in two points, p and q. Now both J_1 and J_2 have a tangent that is parallel to the segment pq, but since they lie on different sides of the pq line, they have opposite outer normal vectors. (See Figure 1 right.) This shows that J has two different tangents parallel to pq and therefore $|g(J)| \geq \pi$. ◀

▶ **Lemma 14.** *For a convex body C, which is not a parallelogram, and an injective parametrization γ of ∂C, we can pick $0 \leq t_1 < t_2 < t_3 \leq 1$ such that $|g(\gamma_{[t_1,t_2]})|, |g(\gamma_{[t_2,t_3]})|$ and $|g(\gamma_{[t_3,t_1]})|$ are each strictly smaller than π.*

Proof. We use the 2-dimensional case of the Illumination conjecture (proved by Levi [20]). If C is not a parallelogram, we can pick three directions, u_1, u_2 and u_3, that illuminate C. Pick t_1 such that $\gamma(t_1)$ is illuminated by both u_1 and u_2. To see why this is possible, suppose that the parts illuminated by u_1 and u_2 are disjoint. Each light illuminates a continuous open ended part of the boundary. So in this case there are two disjoint parts of the boundary that are not illuminated. If u_3 illuminates both, then it illuminates everything that is illuminated by u_1 or everything that is illuminated by u_2. This would mean that two lights illuminate the whole boundary but this is not possible for any convex body. Indeed, suppose that two lights u and v illuminate the whole body. Then there is a halfplane H through the origin that contains both vectors u and v. Take a translate of H that touches C. Clearly the touching point is not illuminated by either u or v, a contradiction.

Using the same argument, pick t_2 and t_3 such that $\gamma(t_2)$ is illuminated by both u_2 and u_3 and $\gamma(t_3)$ is illuminated by both u_3 and u_1.

Note that u_1 illuminates exactly those points for which $g_+(p) < u_1 + \pi/2$ and $g_-(p) > u_1 - \pi/2$. Therefore, $|g(\gamma_{[t_1,t_3]})| < u_1 + \pi/2 - (u_1 - \pi/2) = \pi$. Similarly $|g(\gamma_{[t_1,t_2]})| < \pi$ and $|g(\gamma_{[t_2,t_3]})| < \pi$. ◀

Observation 12 and Lemma 14 immediately imply the following statement.

▶ **Lemma 15.** *For a convex body C, which is not a parallelogram, and an injective parametrization γ of ∂C, we can pick $0 \leq t_1 < t_2 < t_3 \leq 1$ and $\varepsilon > 0$ such that $|g(\gamma_{[t_1-\varepsilon,t_2+\varepsilon]})|$, $|g(\gamma_{[t_2-\varepsilon,t_3+\varepsilon]})|$ and $|g(\gamma_{[t_3-\varepsilon,t_1+\varepsilon]})|$ are each strictly smaller than π.*

3 Proof of Theorem 1

3.1 Quasi-orders on planar point sets

Cones provide a natural way to define quasi-orders on point sets (see [33] for an example where this idea was used). A *cone* is a closed region in the plane that is bounded by two rays that emanate from the origin. For a cone K let $-K$ denote the cone that is the reflection of K across the origin and let $q+K$ denote the translate of K by the vector q.

▶ **Observation 16.** *For any $p, q \in \mathbb{R}^2$ and cone K, the following are equivalent (see Fig. 2):*
- $p \in q + K$
- $q \in p + (-K)$
- $p + K \subseteq q + K$

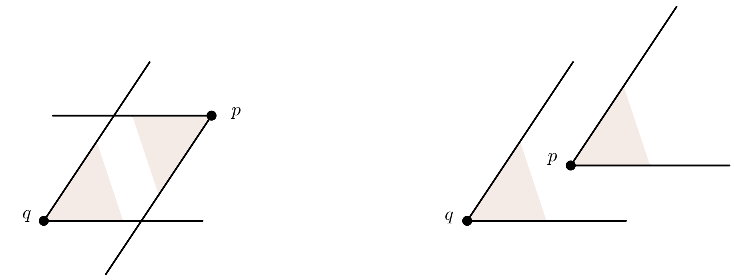

Figure 2 Basic properties of cones.

For a cone K let $\prec_K$ denote the relation on the points of the plane where a point p is bigger than a point q if and only if $p + K$ contains q. By Observation 16, this relation is transitive so it is a quasi-order. Recall that when $\prec_K$ is interpreted as a digraph, qp is an edge if and only if $q \prec_K p$.

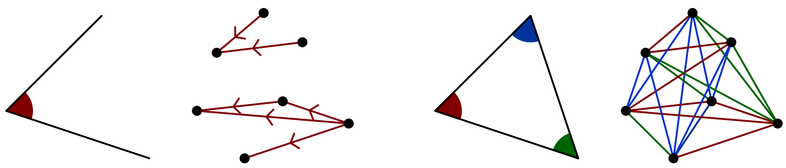

Figure 3 Quasi-order on a point set.

Suppose the cones K_1, K_2, K_3 are the translates of the three corners of a triangle so that all their apexes are in the origin, in other words the cones $K_1, -K_3, K_2, -K_1, K_3, -K_2$ partition the plane around the origin in this order. Then we will say that K_1, K_2, K_3 is a *set of tri-partition* cones. In this case the intersection of any translates of K_1, K_2, K_3 forms a (possibly degenerate) triangle.

▶ **Observation 17.** *Let K_1, K_2, K_3 be a set of tri-partition cones and let P be a planar point set. Then any two distinct points of P are comparable in either $\prec_{K_1}$, $\prec_{K_2}$ or $\prec_{K_3}$. (See Figure 3.)*

In other words, when interpreted as digraphs, the union of $\prec_{K_1}$, $\prec_{K_2}$ and $\prec_{K_3}$ forms a complete multidigraph on P. As a warm up for the proof of Theorem 1, we show the following theorem.

▶ **Theorem 18.** *There exists a positive integer m such that for any point set P, and any set of tri-partition cones K_1, K_2, K_3, we can three-color P such that no translate of K_1, K_2 or K_3 that contains at least m points of P is monochromatic.*

Proof. We set m to be $f(3,2)+13$ with the function of Theorem 10. Consider the three quasi-orders $\prec_{K_1}$, $\prec_{K_2}$ or $\prec_{K_3}$. Their union gives a complete multidigraph on P, hence we can apply Theorem 10 with $k=3$ and $l=2$, resulting in subsets S_i^j for $i \in [3], j \in [2]$. Let $S = \bigcup_{i\in[3],j\in[2]} S_i^j$. For each point $p \in P \setminus S$ there is an i such that $\prec_{K_i}$ has an edge from a vertex of $S_{i,1}$ and $S_{i,2}$ to p. Let P_1, P_2, P_3 be the partition of $P \setminus S$ according to this i value.

We start by coloring the points of S. Color the points of $S_{1,1} \cup S_{2,1} \cup S_{3,1}$ with the first color and color the points of $S_{1,2} \cup S_{2,2} \cup S_{3,2}$ with the second color.

Any translate of K_1, K_2 or K_3 that contains $f(3,2)+13$ points of P, must contain 5 points from either P_1, P_2 or P_3 by the pigeonhole principle. (Note that the cone might contain all points of S.) Therefore, it is enough to show that for each $i \in [3]$ the points of P_i can be three-colored such that no translate of K_1, K_2, or K_3 that contains at least 5 points of P_i is monochromatic.

Consider P_1; the proof is the same for P_2 and P_3. Take a translate of K_1 and suppose that it contains a point p of P_1. By Theorem 10, there is an edge of $\prec_{K_1}$ from a vertex of $S_{1,1}$ to p and another edge from a vertex of $S_{1,2}$ to p. Thus any such translate contains a point from $S_{1,1}$ and another point from $S_{1,2}$, and hence it cannot be monochromatic.

Therefore, we only have to consider the translates of K_2 and K_3. Two translates of a cone intersect at most once on their boundary. Hence, the translates of K_2 form a pseudohalfplane arrangement, and so do the translates of K_3. Therefore, by Corollary 8, there is a proper three-coloring for the translates of K_2 and K_3 together. ◀

▶ Remark 19. From Theorem 18, it follows using standard methods (see Section 3.2) that Theorem 1 holds for triangles. This was of course known before, even for two-colorings of homothetic copies of triangles. Our proof cannot be modified for homothets, but a two-coloring would follow if instead of Corollary 8 we applied a more careful analysis for the two cones.

3.2 Proof of Theorem 1

If C is a parallelogram, then our proof method fails. Luckily, translates of parallelograms (and other symmetric polygons) were the first for which it was shown that even two colors are enough [22]; in fact, by now we know that two colors are enough even for homothets of parallelograms [2]. So from now on we assume that C is not a parallelogram.

The proof of Theorem 1 relies on the same ideas as we used for Theorem 18. We partition P into several parts, and for each part P_i, we divide the translates of C into three families such that two of the families each form a pseudohalfplane arrangement over P_i, while the third family will only contain translates that are automatically non-monochromatic. Then Corollary 8 gives us a proper three-coloring. As in the proof of Theorem 18, this is not done directly. First, we divide the plane using a grid, and then in each small square we will use Theorem 10 to discard some of the translates of C at the cost of a bounded number of points.

Now we start the proof of Theorem 1. The first step is a classic divide and conquer idea [22]. We chose a constant $r = r(C)$ depending only on C and divide the plane into a grid of squares of side length r. Since each translate of C intersects some bounded number of squares, by the pigeonhole principle we can find for any positive integer m another integer m' such that the following holds: each translate $\hat{C}$ of C that contains at least m' points intersects a square Q such that $\hat{C} \cap Q$ contains at least m points. For example, we can choose $m' = m(diam(C)/r + 2)^2$, where $diam(C)$ denotes the diameter of C. Therefore, it is enough to show the following localized version of Theorem 1, since applying it separately for the points in each square of the grid provides a proper three-coloring of the whole point set.

▶ **Theorem 20.** *There is a positive integer m such that for any convex body C there is a positive real r such that any finite point set P in the plane that lies in a square of side length r can be three-colored in a way that there is no translate of C containing at least m points of P, all of the same color.*

We will show that m can be chosen to be $f(3, 2) + 13$ with the function of Theorem 10, independently of C.

Proof. We pick r the following way. First we fix an injective parametrization γ of ∂C and then fix t_1, t_2, t_3 and ε according to Lemma 15. Let ℓ_1, ℓ_2, ℓ_3 be the tangents of C touching at $\gamma(t_1), \gamma(t_2)$ and $\gamma(t_3)$. Let $K_{1,2}$, $K_{2,3}$, $K_{3,1}$ be the set of tri-partition cones bordered by ℓ_1, ℓ_2, ℓ_3, such that $K_{i,i+1}$ is bordered by ℓ_i on its counterclockwise side, and by ℓ_{i+1} on its clockwise side (see Figure 4 left, and note that we always treat $3 + 1$ as 1 in the subscript).

For a translate $\hat{C}$ of C we will denote by $\hat{\gamma}$ the translated parametrization of $\partial\hat{C}$, i.e., $\hat{\gamma}(t) = \gamma(t) + v$ if $\hat{C}$ was translated by vector v. Our aim is to choose r small enough to satisfy the following two properties for each $i \in [3]$.

- **(A)** Let $\hat{C}$ be a translate of C, and Q be a square of side length r such that $\partial\hat{C} \cap Q \subset \hat{\gamma}_{[t_i+\varepsilon/2, t_{i+1}-\varepsilon/2]}$ (see Figure 4 right). Then for any translate K of $K_{i,i+1}$ whose apex is in $Q \cap \hat{C}$, we have $K \cap Q \subset \hat{C}$. (I.e., r is small with respect to C.)
- **(B)** Let $\hat{C}$ be a translate of C, and Q be a square of side length r such that $\hat{\gamma}_{[t_i-\varepsilon/2, t_{i+1}+\varepsilon/2]}$ intersects Q. Then $\partial\hat{C} \cap Q \subset \hat{\gamma}_{[t_i-\varepsilon, t_{i+1}+\varepsilon]}$. (I.e., r is small compared to ε.)

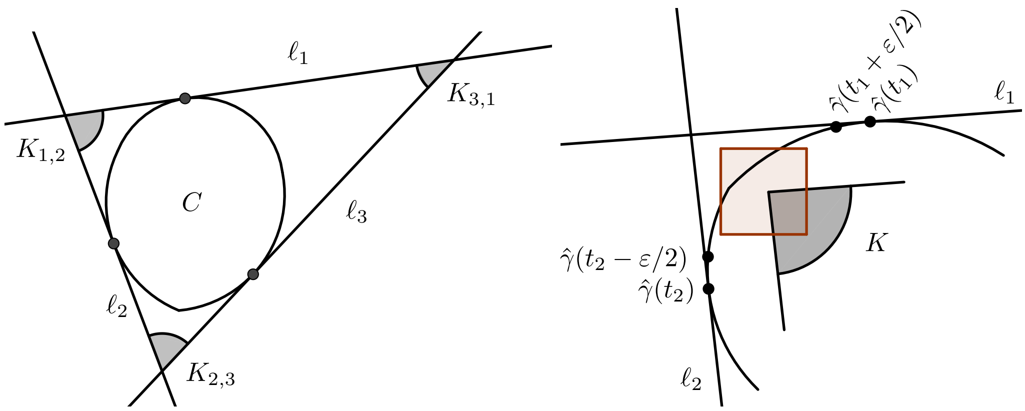

Figure 4 Selecting the cones (on the left) and Property (A) (on the right).

We show that an r satisfying properties (A) and (B) can be found for $i = 1$. The argument is the same for $i = 2$ and $i = 3$, and we can take the smallest among the three resulting values of r.

First, consider property (A). Since the sides of K are parallel to ℓ_1 and ℓ_2, the portion of K that lies "above" the segment $\overline{\hat{\gamma}(t_1)\hat{\gamma}(t_2)}$ is in $\hat{C}$. Hence, if we choose r small enough so that Q cannot intersect $\overline{\hat{\gamma}(t_1)\hat{\gamma}(t_2)}$, then property (A) is satisfied. We can choose r to be smaller than $\frac{1}{\sqrt{2}}$ times the distance of the segments $\overline{\hat{\gamma}(t_1)\hat{\gamma}(t_2)}$ and $\overline{\hat{\gamma}(t_1 + \varepsilon/2)\hat{\gamma}(t_2 - \varepsilon/2)}$.

Using that γ is a continuous function on a compact set, we can pick r such that property (B) is satisfied. Therefore, there is an r satisfying properties (A) and (B).

The next step is a subdivision of the point set P using Theorem 10, like we did in the proof of Theorem 18. The beginning of our argument is exactly the same.

Apply Theorem 10 for the graph given by the union of $\prec_{K_{1,2}}$, $\prec_{K_{2,3}}$ and $\prec_{K_{3,1}}$. By Observation 16, this is indeed a complete multidigraph on P.

We apply Theorem 10 with $k = 3$ and $l = 2$, resulting in subsets S_i^j for $i \in [3], j \in [2]$. Let $S = \bigcup_{i\in[3],j\in[2]} S_i^j$. For each point $p \in P \setminus S$ there is an i such that $\prec_{K_{i,i+1}}$ has an edge from a vertex of $S_{i,1}$ and $S_{i,2}$ to p. Let P_1, P_2, P_3 be the partition of $P \setminus S$ according to this i value.

We start by coloring the points of S. Color the points of $S_{1,1} \cup S_{2,1} \cup S_{3,1}$ with the first color and color the points of $S_{1,2} \cup S_{2,2} \cup S_{3,2}$ with the second color.

Note that m is at least $f(3, 2) + 13$. Any translate of C that contains $f(3, 2) + 13$ points of P must contain 5 points from either P_1, P_2 or P_3. (Note that the cone might contain all points of S). Thus, it is enough to show that for each $i \in [3]$ the points of P_i can be 3-colored so that no translate of C that contains at least 5 points of P_i is monochromatic.

Consider P_1, the proof is the same for P_2 and P_3. We divide the translates of C that intersect Q into four (not necessarily disjoint) groups. Let $\mathcal{C}_0$ denote the translates where $\hat{C} \cap Q = Q$. Let $\mathcal{C}_1$ denote the translates for which $\partial\hat{C} \cap Q \subset \hat{\gamma}_{[t_1+\varepsilon/2,t_2-\varepsilon/2]}$. Let $\mathcal{C}_2$ denote the translates for which $\partial\hat{C} \cap Q \cap \hat{\gamma}_{[t_2-\varepsilon/2,t_3]} \neq \emptyset$. Let $\mathcal{C}_3$ denote the remaining translates for which $\partial\hat{C} \cap Q \cap \hat{\gamma}_{[t_3,t_1+\varepsilon/2]} \neq \emptyset$.

We do not need to worry about the translates in $\mathcal{C}_0$, as Q itself will not be monochromatic.

Take a translate $\hat{C}$ from $\mathcal{C}_1$ and suppose that it contains a point $p \in P_1$. By Theorem 10, there is an edge of $\prec_{K_{1,2}}$ from a vertex of $S_{1,1}$ to p and another edge from a vertex of $S_{1,2}$ to p. I.e., the cone $p + K_{1,2}$ contains a point from $S_{1,1}$ and another point from $S_{1,2}$, and hence it is not monochromatic. From property (A) we know that every point in $(p + K_{1,2}) \cap P$ is also in $\hat{C}$. Therefore, $\hat{C}$ is not monochromatic.

Now consider the translates in $\mathcal{C}_2$. From property (B) we know that for these translates we have $\partial\hat{C} \cap Q \subset \hat{\gamma}_{[t_2-\varepsilon,t_3+\varepsilon]}$. By the definition of t_1, t_2 and t_3, we know that this implies that any two translates from $\mathcal{C}_2$ intersect at most once on their boundary within Q, i.e., they behave as pseudohalfplanes. To turn the translates in $\mathcal{C}_2$ into a pseudohalfplane arrangement as defined earlier, we can do as follows. For a translate $\hat{C}$, replace it with the convex set whose boundary is $\hat{\gamma}_{[t_2-\varepsilon,t_3+\varepsilon]}$ extended from its endpoints with two rays orthogonal to the segment $\overline{\hat{\gamma}(t_2 - \varepsilon)\hat{\gamma}(t_3 + \varepsilon)}$. This new family provides the same intersection pattern in Q and forms a pseudohalfplane arrangement. We can do the same with the translates in $\mathcal{C}_3$. Therefore, by Corollary 8 there is a proper three-coloring for the translates in $\mathcal{C}_2 \cup \mathcal{C}_3$. ◀

4 Overview of the computational complexity of the algorithm

In this section we show that given a point set P and a convex set C, we can determine some $m = m(C)$ and calculate a three-coloring of P efficiently if C is given in a natural way, for example, if C is a disk. Our algorithm is randomized and its expected running time is a polynomial of the number of points, $n = |P|$.

- First, we need to fix three points on the boundary, $\tau_1, \tau_2, \tau_3 \subset \partial C$ such that Lemma 15 is satisfied with $\tau_i = \gamma(t_i)$ for some t_i and $\varepsilon > 0$ for each i. Note that we do not need to fix a complete parametrization γ of ∂C or $\varepsilon > 0$; instead, it is enough to choose some points τ_i^{--} and τ_i^{++} that satisfy the conclusion of Lemma 15 if we assume $\tau_i^{--} = \gamma(t_i - \varepsilon)$ and $\tau_i^{++} = \gamma(t_i + \varepsilon)$ for each i. If C has a smooth boundary, like a disk, we can pick τ_1, τ_2, τ_3 to be the touching points of an equilateral triangle with C inscribed in it. If the boundary of C contains vertex-type sharp turns, the complexity of finding these turns depends on how C is given, but for any reasonable input method, this should be straight-forward. After that, one can follow closely the steps of the proof of the Illumination conjecture in the plane to get an algorithm, but apparently, this has not yet been studied in detail.

- To pick r, the side length of the squares of the grid, we can fix some arbitrary points τ_i^- between τ_i^{--} and τ_i, and points τ_i^+ between τ_i and τ_i^{++}, to play the roles of $\gamma(t_i - \varepsilon/2)$ and $\gamma(t_i + \varepsilon/2)$, respectively, for each i. It is sufficient to pick r so that $r\sqrt{2}$, the diameter of the square of side length r, is less than
  - the distance of τ_i^+ and τ_{i+1}^- from the segment $\overline{\tau_i \tau_{i+1}}$,
  - the distance of τ_i^- from τ_i^{--}, and
  - the distance of τ_i^+ from τ_i^{++},

  for each i, to guarantee that properties (A) and (B) are satisfied.
- Set $m = f(3,2) + 13$, which is an absolute constant given by Theorem 10. We need to construct the complete multidigraph given by the tri-partition cones determined by τ_1, τ_2, τ_3, which needs a comparison for each pair of points. To obtain the subsets $S_i^j \subset P$ for $i \in [3], j \in [2]$, where P is the set of points that are contained in a square of side length r, we randomly sample the required number of points from each of the constantly many $T_{j_1,\dots,j_i}$ according to the probability distributions $w_{j_1,\dots,j_i}$ given in the proof. These probability distributions can be computed by LP. With high probability, all the S_i^j-s will be disjoint – otherwise, we can resample until we obtain disjoint sets.
- To find the three-coloring for the two pseudohalfplane arrangements, given by Corollary 8, it is enough to determine the two-coloring given by Theorem 7 for one pseudohalfplane arrangement. While not mentioned explicitly in [17], the polychromatic k-coloring can be found in polynomial time if we know the hypergraph determined by the range space, as this hypergraph can only have a polynomial number of edges, and the coloring algorithm only needs to check some simple relations among a constant number of vertices and edges.
- Finally, to compute a suitable m' for Theorem 1 from the m of Theorem 20, it is enough to know any upper bound B for the diameter of C, and let $m' = m(B/r + 2)^2$.

5 Open questions

It is a natural question whether there is a universal m that works for all convex bodies in Theorem 1, like in Theorem 20. This would follow if we could choose r to be a universal constant. While the r given by our algorithm can depend on C, we can apply an appropriate affine transformation to C before choosing r; this does not change the hypergraphs that can be realized with the range space determined by the translates of C. To ensure that properties (A) and (B) are satisfied would require further study of the Illumination conjecture.

Our bound for m is quite large, even for the unit disk, both in Theorems 1 and 20, which is mainly due to the fact that $f(3,2)$ given by Theorem 10 is huge. It has been conjectured that in Theorem 9 the optimal value is $f(3) = 3$, and a similarly small number seems realistic for $f(3,2)$ as well.

While Theorem 1 closed the last question left open for primal hypergraphs realizable by translates of planar bodies, the respective problem is still open in higher dimensions. While it is not hard to show that some hypergraphs with high chromatic number often used in constructions can be easily realized by unit balls in $\mathbb{R}^5$, we do not know whether the chromatic number is bounded or not in $\mathbb{R}^3$. From our Union Lemma (Lemma 6) it follows that to establish boundedness, it would be enough to find a polychromatic k-coloring for pseudohalfspaces, whatever this word means.

References

1 Eyal Ackerman, Balázs Keszegh, and Dömötör Pálvölgyi. Coloring hypergraphs defined by stabbed pseudo-disks and ABAB-free hypergraphs. *SIAM J. Discrete Math.*, 34(4):2250–2269, 2020. doi:10.1137/19M1290231.

2 Eyal Ackerman, Balázs Keszegh, and Máté Vizer. Coloring points with respect to squares. *Discrete Comput. Geom.*, 58(4):757–784, 2017. doi:10.1007/s00454-017-9902-y.

3 Greg Aloupis, Jean Cardinal, Sébastien Collette, Stefan Langerman, David Orden, and Pedro Ramos. Decomposition of multiple coverings into more parts. *Discrete Comput. Geom.*, 44(3):706–723, 2010. doi:10.1007/s00454-009-9238-3.

4 Andrei Asinowski, Jean Cardinal, Nathann Cohen, Sébastien Collette, Thomas Hackl, Michael Hoffmann, Kolja Knauer, Stefan Langerman, MichałLasoń, Piotr Micek, Günter Rote, and Torsten Ueckerdt. Coloring hypergraphs induced by dynamic point sets and bottomless rectangles. In *Algorithms and data structures*, volume 8037 of *Lecture Notes in Comput. Sci.*, pages 73–84. Springer, Heidelberg, 2013. doi:10.1007/978-3-642-40104-6_7.

5 Károly Bezdek and Muhammad A. Khan. The geometry of homothetic covering and illumination. In *Discrete geometry and symmetry*, volume 234 of *Springer Proc. Math. Stat.*, pages 1–30. Springer, Cham, 2018. doi:10.1007/978-3-319-78434-2_1.

6 Nicolas Bousquet, William Lochet, and Stéphan Thomassé. A proof of the Erdős-Sands-Sauer-Woodrow conjecture. *J. Combin. Theory Ser. B*, 137:316–319, 2019. doi:10.1016/j.jctb.2018.11.005.

7 Jean Cardinal, Kolja Knauer, Piotr Micek, and Torsten Ueckerdt. Making triangles colorful. *J. Comput. Geom.*, 4(1):240–246, 2013. doi:10.20382/jocg.v4i1a10.

8 Jean Cardinal, Kolja Knauer, Piotr Micek, and Torsten Ueckerdt. Making octants colorful and related covering decomposition problems. *SIAM J. Discrete Math.*, 28(4):1948–1959, 2014. doi:10.1137/140955975.

9 Jean Cardinal, Piotr Micek Kolja Knauer, Dömötör Pálvölgyi, Torsten Ueckerdt, and Narmada Varadarajan. Colouring bottomless rectangles and arborescences. *To appear*, 2020.

10 Xiaomin Chen, János Pach, Mario Szegedy, and Gábor Tardos. Delaunay graphs of point sets in the plane with respect to axis-parallel rectangles. *Random Structures Algorithms*, 34(1):11–23, 2009. doi:10.1002/rsa.20246.

11 Gábor Damásdi and Dömötör Pálvölgyi. Unit disks hypergraphs are three-colorable. *Extended Abstracts EuroComb 2021, Trends in Mathematics*, 14:483–489, 2021.

12 Gábor Damásdi and Dömötör Pálvölgyi. Realizing an m-uniform four-chromatic hypergraph with disks. *Combinatorica, to appear*, 2022.

13 Matt Gibson and Kasturi Varadarajan. Optimally decomposing coverings with translates of a convex polygon. *Discrete Comput. Geom.*, 46(2):313–333, 2011. doi:10.1007/s00454-011-9353-9.

14 Balázs Keszegh. Coloring half-planes and bottomless rectangles. *Comput. Geom.*, 45(9):495–507, 2012. doi:10.1016/j.comgeo.2011.09.004.

15 Balázs Keszegh and Dömötör Pálvölgyi. Octants are cover-decomposable. *Discrete Comput. Geom.*, 47(3):598–609, 2012. doi:10.1007/s00454-011-9377-1.

16 Balázs Keszegh and Dömötör Pálvölgyi. Convex polygons are self-coverable. *Discrete Comput. Geom.*, 51(4):885–895, 2014. doi:10.1007/s00454-014-9582-9.

17 Balázs Keszegh and Dömötör Pálvölgyi. An abstract approach to polychromatic coloring: shallow hitting sets in ABA-free hypergraphs and pseudohalfplanes. *J. Comput. Geom.*, 10(1):1–26, 2019. doi:10.20382/jocg.v10i1a1.

18 Balázs Keszegh and Dömötör Pálvölgyi. Proper coloring of geometric hypergraphs. *Discrete Comput. Geom.*, 62(3):674–689, 2019. doi:10.1007/s00454-019-00096-9.

19 István Kovács. Indecomposable coverings with homothetic polygons. *Discrete Comput. Geom.*, 53(4):817–824, 2015. doi:10.1007/s00454-015-9687-9.

20 F. W. Levi. Überdeckung eines Eibereiches durch Parallelverschiebung seines offenen Kerns. *Arch. Math. (Basel)*, 6:369–370, 1955. doi:10.1007/BF01900507.

21 János Pach. Decomposition of multiple packing and covering. *Diskrete Geometrie, 2. Kolloq. Math. Inst. Univ. Salzburg*, pages 169–178, 1980.

22 János Pach. Covering the plane with convex polygons. *Discrete Comput. Geom.*, 1(1):73–81, 1986. doi:10.1007/BF02187684.

23 János Pach and Dömötör Pálvölgyi. Unsplittable coverings in the plane. *Adv. Math.*, 302:433–457, 2016. doi:10.1016/j.aim.2016.07.011.

24 János Pach, Dömötör Pálvölgyi, and Géza Tóth. Survey on decomposition of multiple coverings. In *Geometry—intuitive, discrete, and convex*, volume 24 of *Bolyai Soc. Math. Stud.*, pages 219–257. János Bolyai Math. Soc., Budapest, 2013. doi:10.1007/978-3-642-41498-5_9.

25 János Pach and Gábor Tardos. Coloring axis-parallel rectangles. *J. Combin. Theory Ser. A*, 117(6):776–782, 2010. doi:10.1016/j.jcta.2009.04.007.

26 János Pach, Gábor Tardos, and Géza Tóth. Indecomposable coverings. In *Discrete geometry, combinatorics and graph theory*, volume 4381 of *Lecture Notes in Comput. Sci.*, pages 135–148. Springer, Berlin, 2007. doi:10.1007/978-3-540-70666-3_15.

27 Dömötör Pálvölgyi. Indecomposable coverings with concave polygons. *Discrete Comput. Geom.*, 44(3):577–588, 2010. doi:10.1007/s00454-009-9194-y.

28 Dömötör Pálvölgyi and Géza Tóth. Convex polygons are cover-decomposable. *Discrete Comput. Geom.*, 43(3):483–496, 2010. doi:10.1007/s00454-009-9133-y.

29 Ioannis Papadoperakis. An estimate for the problem of illumination of the boundary of a convex body in E^3. *Geom. Dedicata*, 75(3):275–285, 1999. doi:10.1023/A:1005056207406.

30 A Prymak. Every 3-dimensional convex body can be covered by 14 smaller homothetic copies. *arXiv preprint*, 2021. arXiv:2112.10698.

31 Bill Sands, Norbert W. Sauer, and Robert E. Woodrow. On monochromatic paths in edge-coloured digraphs. *J. Combin. Theory Ser. B*, 33(3):271–275, 1982. doi:10.1016/0095-8956(82)90047-8.

32 Shakhar Smorodinsky and Yelena Yuditsky. Polychromatic coloring for half-planes. *J. Combin. Theory Ser. A*, 119(1):146–154, 2012. doi:10.1016/j.jcta.2011.07.001.

33 Gábor Tardos and Géza Tóth. Multiple coverings of the plane with triangles. *Discrete Comput. Geom.*, 38(2):443–450, 2007. doi:10.1007/s00454-007-1345-4.