Planar convex codes are decidable

Boris Bukh[*]   R. Amzi Jeffs[†]

December 2022

Abstract

We show that every convex code realizable by compact sets in the plane admits a realization consisting of polygons, and analogously every open convex code in the plane can be realized by interiors of polygons. We give factorial-type bounds on the number of vertices needed to form such realizations. Consequently we show that there is an algorithm to decide whether a convex code admits a closed or open realization in the plane.

Introduction

Prelude.   A basic way to represent relationships between several sets is to use *Venn diagrams*. In a Venn diagram, the sets are represented by planar regions, and elements are represented by points in such a way that the containment relation is preserved. For example, Figure 1(a) depicts the usual Venn diagram representing three sets A, B and C, with eight regions containing points of the eight sets $A \cap B \cap C, \overline{A} \cap B \cap C, \dots, \overline{A} \cap \overline{B} \cap \overline{C}$ respectively.

If certain intersections are empty, we may use Venn diagrams omitting them, as in Figure 1(b).
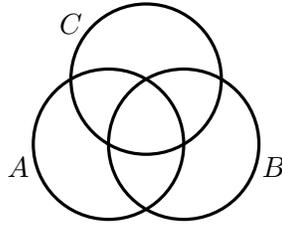
a) Eight regions are present   **b)** Region $A \cap B \cap C$ is omitted
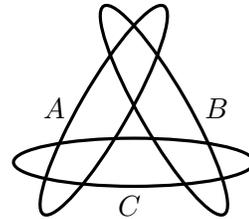
Figure 1: Three-set Venn diagrams.

This paper is about Venn diagrams realizable by convex regions. These are also known as *convex codes*, which is the terminology that we adopt. Convex codes have been studied extensively over the last decade, with most work following a 2013 paper of Curto, Itskov, Veliz-Cuba and Youngs [5]. The theory of convex codes can be viewed as a strengthening of the theory of d-representable complexes,

which are the simplicial complexes that record nonempty intersections of a collection of convex sets in $\mathbb{R}^d$. Convex codes record not only the nonempty intersections, but also how other sets cover these intersections.

It is well known that there is an algorithm to determine whether or not a given simplicial complex is d-representable [12, Section 4.1], but the analogous problem for convex codes has remained open. Chen, Frick and Shiu [2] showed that recognizing codes that can be realized by "good covers" is undecidable, and asked about recognizing convex codes realizable by bounded open sets. Our main result is an affirmative answer to this question in the case $d = 2$ (Theorem 1 below). Though our focus is on convex sets in the plane, we work in general $\mathbb{R}^d$ whenever possible.

Convex codes. Let $X_1, \ldots, X_n$ be n bounded convex sets in $\mathbb{R}^d$. The *intersection pattern* of the tuple $\mathcal{X} = (X_1, \ldots, X_n)$ at the point $p \in \mathbb{R}^d$ is $\mathrm{patt}_{\mathcal{X}}(p) \stackrel{\text{def}}{=} \{i \in [n] \mid p \in X_i\}$. The *convex code* of $\mathcal{X}$ is then defined as $\mathrm{code}(\mathcal{X}) \stackrel{\text{def}}{=} \{\mathrm{patt}_{\mathcal{X}}(p) \mid p \in \mathbb{R}^d\}$. If $\mathcal{C} = \mathrm{code}(\mathcal{X})$, we say that $\mathcal{C}$ is a convex code *realized* by $\mathcal{X}$, and that $\mathcal{X}$ is a d-realization of $\mathcal{C}$. Franke and Muthiah showed that every $\mathcal{C} \subseteq 2^{[n]}$ containing $\emptyset$ is d-realizable for some d [6, Theorem 1].

We shall say that a realization $\mathcal{X} = (X_1, \ldots, X_n)$ of a convex code $\mathcal{C}$ is *closed* if every set X_i is closed (which implies that it is compact). We similarly define an *open* realization. If we consider only the realizations in $\mathbb{R}^d$ for a fixed $d \geq 2$, these are different concepts: Some convex codes admit only open d-realizations, some admit only closed d-realizations, and some admit neither; the first example showing the difference between open and closed d-realizations is due to Lienkaemper, Shiu and Woodstock [10], and [3, 7, 8] contain more recent results.

Our results. The primary motivation for this work is to find a way to tell which convex codes are d-realizable. To that end, it is natural to wonder whether every convex code admits a not-too-complicated realization. We give a positive answer for $d = 2$, namely that every closed convex code in the plane can be realized by polygons (some of which might degenerate into line segments or points), and that the number of vertices among these polygons is bounded by a computable function. We obtain the analogous result for open realizations in the plane as well. As an immediate consequence, we conclude that there is an algorithm for determining if a convex code has a closed or open 2-realization.

Theorem 1. *Let $\mathcal{C} \subseteq 2^{[n]}$ be a convex code.*

    i) If $\mathcal{C}$ admits a closed convex realization in $\mathbb{R}^2$, then it admits a realization consisting of polygons with at most $6^n n!$ total vertices among them, and

    ii) If $\mathcal{C}$ admits an open convex realization in $\mathbb{R}^2$, then it admits a realization consisting of interiors of polygons with at most $4 \cdot 6^n (n+1)!$ total vertices among them.

As a consequence, there is an algorithm to decide whether or not $\mathcal{C}$ admits a closed (respectively, open) convex realization in the plane.

Our proof consists of three steps. First, starting with an arbitrary realization, we shrink its sets as much as possible. Second, we show that the resulting inclusion-minimal realizations are polygonal. Third, we bound the total number of vertices. The first step works for any d, but our methods

for the last two are specific to the plane. The situation for $d \geq 3$ remains unclear — we do not know if convex codes in higher dimensions admit polytopal realizations, and we do not know if there exists a computable bound on the number of vertices in inclusion-minimal polytopal realizations. In particular, it is possible that there is no algorithm for deciding whether an arbitrary convex code is d-realizable for fixed $d \geq 3$.

Sharpness. We do not know if the factorial-type bound in Theorem 1 is necessary. However, the number of vertices must grow with the size of the code.

Theorem 2. *Let $\mathcal{X}$ be a realization of $\mathcal{C} \subseteq 2^{[n]}$ consisting of closed polygons. Then the total number of vertices among the polygons in $\mathcal{X}$ is at least $|\mathcal{C}|/8n$.*

In particular, since the code $2^{[n]}$ admits a realization by closed polygons [4, Lemma 2.5], it follows that the $6^n n!$ in Theorem 1 cannot be improved to anything less than $\frac{1}{n} 2^{n-3}$.

Step 1: Existence of inclusion-minimal realizations

Let $\mathcal{X} = (X_1, \ldots, X_n)$ be a tuple of compact convex sets. Recall that $\mathrm{code}(\mathcal{X}) \stackrel{\mathrm{def}}{=} \{\mathrm{patt}_{\mathcal{X}}(p) \mid p \in \mathbb{R}^d\}$. We say that $P \subseteq \mathbb{R}^d$ is a *set of representatives* for $\mathcal{X}$ if $\mathrm{code}(\mathcal{X}) = \{\mathrm{patt}_{\mathcal{X}}(p) \mid p \in P\}$. Obviously, every $\mathcal{X}$ admits a set of representatives of size $|\mathrm{code}(\mathcal{X})| \leq 2^n$ by picking a point in each region.

Lemma 3. *Let $\mathcal{X} = (X_1, \ldots, X_n)$ be a tuple of compact convex sets in $\mathbb{R}^d$, and P is a set of representatives for $\mathcal{X}$. Then there exists a tuple $\mathcal{Y} = (Y_1, \ldots, Y_n)$ of compact convex sets with the following properties:*

  i) $Y_i \subseteq X_i$ for all $i \in [n]$,

  ii) $Y_i \cap P = X_i \cap P$ for all $i \in [n]$,

  iii) $\mathrm{code}(\mathcal{Y}) = \mathrm{code}(\mathcal{X})$, and

  iv) The tuple $\mathcal{Y}$ is inclusion-minimal among all tuples satisfying (i), (ii) and (iii).

Proof. Let $\mathfrak{R}$ be the set consisting of tuples that satisfy conditions (i), (ii) and (iii). Introduce a partial order on $\mathfrak{R}$ by declaring $\mathcal{Y} \leq \mathcal{Y}'$ if and only if $Y_i' \subseteq Y_i$ for all $i \in [n]$. The poset $\mathfrak{R}$ is nonempty as it contains $\mathcal{X}$. The condition (iv) is equivalent to the tuple $\mathcal{Y}$ being a maximal element in this poset. We shall find it using Zorn's lemma.

Given a chain $\mathfrak{C}$ of tuples in $\mathfrak{R}$, define a tuple $\mathcal{Y} = (Y_1, \ldots, Y_n)$ by $Y_i = \bigcap_{\mathcal{Y}' \in \mathfrak{C}} Y_i'$. Each Y_i is an intersection of compact convex sets, and so is itself compact and convex. Moreover, each Y_i satisfies (i) and (ii) since these properties are preserved by intersections. It remains to argue that $\mathcal{Y}$ satisfies condition (iii).

Since $\mathrm{patt}_{\mathcal{Y}}(p) = \mathrm{patt}_{\mathcal{X}}(p)$ for every $p \in P$ and P is a set of representatives for $\mathcal{X}$, it follows that $\mathrm{code}(\mathcal{X}) \subseteq \mathrm{code}(\mathcal{Y})$. Conversely, let q be any point in $\mathbb{R}^d$. For every $i \notin \mathrm{patt}_{\mathcal{Y}}(q)$, there exists some $\mathcal{Y}^{(i)} \in \mathfrak{C}$ such that $q \notin Y_i^{(i)}$. For each $i \in [n]$, fix such a $\mathcal{Y}^{(i)}$. Let $\mathcal{Y}'$ be the smallest (with respect to inclusion) among the tuples $\mathcal{Y}^{(1)}, \ldots, \mathcal{Y}^{(n)}$. Then $i \notin \mathrm{patt}_{\mathcal{Y}'}(q)$ whenever $i \notin \mathrm{patt}_{\mathcal{Y}}(q)$.

Since $Y_i \subseteq Y_i'$, we also have $i \in \mathrm{patt}_{\mathcal{Y}'}(q)$ whenever $i \in \mathrm{patt}_{\mathcal{Y}}(q)$. Therefore, $\mathrm{patt}_{\mathcal{Y}}(q) = \mathrm{patt}_{\mathcal{Y}'}(q)$. Since $\mathcal{Y}' \in \mathfrak{C}$, this implies that $\mathrm{patt}_{\mathcal{Y}}(q) \in \mathrm{code}(\mathcal{X})$. Since $q \in \mathbb{R}^d$ is arbitrary, $\mathrm{code}(\mathcal{Y}) \subseteq \mathrm{code}(\mathcal{X})$ and hence $\mathrm{code}(\mathcal{Y}) = \mathrm{code}(\mathcal{X})$. We conclude that every chain $\mathfrak{C}$ in $\mathfrak{R}$ has an upper bound, so $\mathfrak{R}$ has a maximal element, and the result follows. $\qquad\square$

Step 2: Inclusion-minimal realizations in the plane are polygonal

Lemma 4. *Let $C \subseteq \mathbb{R}^2$ be a compact convex set with non-empty interior, and let p be a boundary point of C. Then there exists $\varepsilon > 0$ so that every circle S of positive radius $\delta < \varepsilon$ centered at p intersects C in a connected arc.*

Proof. Consider the set
$$C_p \overset{\mathrm{def}}{=} \{v \in \mathbb{S}^1 \mid p + tv \in C \text{ for some } t > 0\}.$$

The set C_p consists of the unit vectors which point "into" C from the boundary point p. Since C has non-empty interior, C_p is a connected arc (not necessarily closed or open) that contains more than one point. Up to affine transformation, we may assume that $p = (0,0)$, and that C_p is an arc in the upper half of the unit circle. Since C_p might not be closed, it might not contain its endpoints. So, instead we select points q_1 and q_2 in C_p that are within angle $\pi/4$ of the left and right endpoints of C_p, respectively. For convenience, we may also choose q_1 and q_2 to have positive y-coordinates. Finally, by scaling C up, we may assume that both q_1 and q_2 lie in C.

Let S be a circle of radius $\varepsilon > 0$ centered at p. The ε is chosen small enough to satisfy two conditions, which we call DISJOINTNESS and TANGENCY.

Disjointness: The ball $\mathrm{conv}(S)$ is disjoint from the line segment $q_1 q_2$.

Consider the leftmost tangent from q_1 to S, and let $t_1 = t_1(S)$ be the tangency point. Define the point $t_2 = t_2(S)$ symmetrically using the tangent from q_2.

Tangency: The angles $\angle q_1 p t_1$ and $\angle q_2 p t_2$ are both greater than $\pi/4$.

This situation is shown in Figure 2.

Note that any smaller circle centered at p also satisfies DISJOINTNESS and TANGENCY. We claim that any such circle intersects C in a connected arc. Without loss of generality, we prove this for S.

Let q_1' (respectively, q_2') be the point where the line segment from q_1 (respectively, q_2) to p crosses S. There are two arcs of S connecting q_1' and q_2'. One of the arcs is contained in the quadrilateral $\mathrm{conv}(q_1, q_2, q_1', q_2')$. Since its vertices q_1, q_2, q_1', q_2' lie in C, it follows that the arc is contained in $S \cap C$. To see that $S \cap C$ is a connected arc, it will suffice to show that any $q' \in S \cap C$ is connected to q_1' or q_2' by an arc in $S \cap C$.

Let $q' \in S \cap C$, and suppose that q' does not lie between q_1' and q_2'. By the choice of q_1 and q_2, the point q' is at the angle at most $\pi/4$ away from either q_1' or q_2'. Consider the first case (the other case being symmetric), i.e., q' is to the left of q_1'. Since the angle $\angle q_1 p q'$ is at most $\pi/4$, and the angle $\angle q_1 p t_1$ exceeds $\pi/4$, the point q' lies in the cone bound by rays $\overrightarrow{pq_1}$ and $\overrightarrow{pt_1}$. Since the line $q_1 q'$ intersects S, it follows that q' in fact lies inside the triangle $\triangle p q_1 t_1$. From this it follows that the triangle $\triangle p q_1 q'$ is contained in $\triangle p q_1 t_1$, and so the angle $\angle p q' q_1$ is at least $\angle p t_1 q_1 = \pi/2$. From this

4
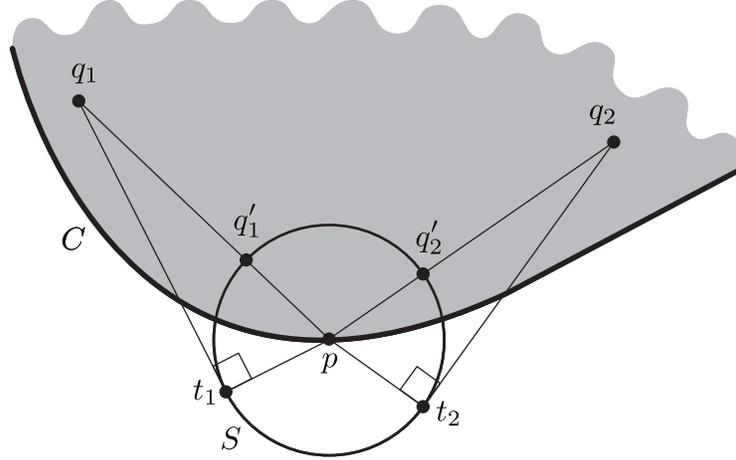
Figure 2: The small circle chosen in the proof of Lemma 4, and the relevant tangent segments.

it follows that the line segment $q_1 q'$ intersects S only in q', and so its radial projection from p onto S is contained in $S \cap C$. The projection is the desired arc from q' to q'_1 in $S \cap C$. □

a) Before   **b)** After
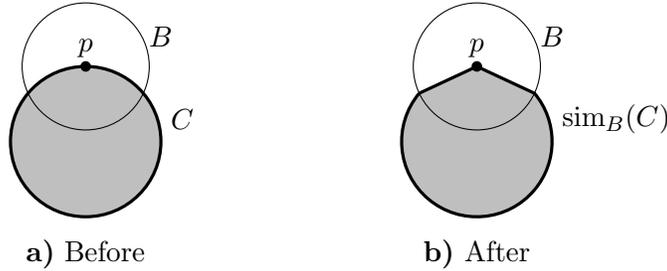
Figure 3: Simplification of a convex set C in the neighborhood of a point p.

For a set $A \subseteq \mathbb{R}^d$ and a point $p \in \mathbb{R}^d$, the *cone over A with apex p* is

$$\mathrm{cone}_p(A) \overset{\mathrm{def}}{=} \{\alpha p + (1 - \alpha)x \mid 0 \leq \alpha \leq 1 \text{ and } x \in A\} \cup \{p\}.$$

(The reason for including $\{p\}$ in this definition is to ensure that $\mathrm{cone}_p(\emptyset) = \{p\}$.)

Definition 5. *Let $C \subseteq \mathbb{R}^2$ be a compact convex set, and let p be a boundary point of C. Let B be a closed ball centered at p, and let ∂B be the boundary of B. The* simplification *of C relative to B is the set*

$$\mathrm{sim}_B(C) \overset{\mathrm{def}}{=} (C \setminus B) \cup \mathrm{cone}_p(\partial B \cap C).$$

Lemma 6. *Let B, C and p be as in Definition 5. If $\partial B \cap C$ is a connected arc, then $\mathrm{sim}_B(C)$ is the convex hull of $\{p\} \cup (C \setminus \mathrm{int}\, B)$.*

Proof. The definition of $\mathrm{sim}_B(C)$ implies that $\mathrm{sim}_B(C) \subseteq \mathrm{conv}\big(\{p\} \cup (C \setminus \mathrm{int}\, B)\big)$. To show the reverse inclusion it is enough to argue that $\mathrm{sim}_B(C)$ is convex. To that end, let $p_1, p_2 \in \mathrm{sim}_B(C)$ be arbitrary,

5

and consider two cases, according to how the line segment p_1p_2 is situated in relation to $\text{sim}_B(C)$. For brevity, denote the arc $\partial B \cap C$ by A.

Suppose that both p_1 and p_2 lie in $\text{cone}_p(A)$**:** If one of them is p, the line segment p_1p_2 is contained in $\text{sim}_B(C)$ by definition. Otherwise, neither is equal to p and we may consider the points p_1' and p_2' obtained by projecting p_1 and p_2 radially onto ∂B from p. Since p lies on the boundary of C, the angular measure of the arc A is at most π. As both p_1' and p_2' lie in A, this implies that the radial projection of the line segment p_1p_2 onto ∂B is an arc in A connecting p_1' to p_2'. So, the line segment p_1p_2 lies in $\text{sim}_B(C)$.

Suppose one of p_1 and p_2 does not lie in $\text{cone}_p(A)$**:** The line segment p_1p_2 lies in C, and so we need prove only that any portion of it that lies in B also lies in $\text{cone}_p(A)$. Note that if $p_1p_2 \cap B$ is nonempty, then it is a line segment whose endpoints lie in $\text{cone}_p(A)$. In our analysis of the first case, we showed that such a line segment is contained in $\text{cone}_p(A)$. Hence p_1p_2 is contained in $\text{sim}_B(C)$ and the result follows. $\qquad \square$

Proposition 7. *Let $\mathcal{X} = (X_1, \ldots, X_n)$ be a tuple of compact convex sets in $\mathbb{R}^2$, and let P be a finite set of representatives for $\mathcal{X}$. If $\mathcal{Y} = (Y_1, \ldots, Y_n)$ is an inclusion-minimal realization of $\text{code}(\mathcal{X})$ as in Lemma 3, then each Y_i is a polygon.*

Proof. Recall that x is an *extreme point* of a convex set C if x is not contained in $\text{conv}(C \setminus \{x\})$. Every compact convex set is the convex hull of its extreme points, and so it will suffice to show that every Y_i has finitely many extreme points. Suppose for contradiction that some Y_i has infinitely many extreme points. Each extreme point must lie on the boundary of Y_i, and since the boundary is compact there must exist a point p on the boundary of Y_i which is a limit point of the extreme points of Y_i. Fix this i and the point p for the remainder of the proof.

Choose a closed ball B centered at p whose radius is small enough as to satisfy the following four conditions for each $j \in [n]$:

- B is disjoint from ∂Y_j unless $p \in \partial Y_j$,

- if Y_j has a non-empty interior and $p \in \partial Y_j$ then B satisfies the conclusion of Lemma 4, i.e., $\partial B \cap Y_j$ is a connected arc spanning an angle of no more than π,

- if Y_j is a line segment containing p, the set $\partial B \cap Y_j$ contains at least one point,

- B is disjoint from $P \setminus \{p\}$.

Now let $\mathcal{Y}' = (Y_1', \ldots, Y_n')$ where

$$Y_j' \overset{\text{def}}{=} \begin{cases} Y_j & \text{if } p \notin \partial Y_j, \\ \text{sim}_B(Y_j) & \text{if } p \in \partial Y_j. \end{cases}$$

Note that if Y_j is a line segment, then $\partial Y_j = Y_j$ and so $Y_j' = \text{sim}_B(Y_j)$. Clearly $Y_j' \subseteq Y_j$ for every $j \in [n]$. Moreover, the simplification $Y_i' = \text{sim}_B(Y_i)$ is a proper subset of Y_i as p is no longer a limit point of the extreme points.

We claim that $\mathcal{Y}'$ is a closed convex realization of $\mathrm{code}(\mathcal{Y})$. Observe that each Y'_j is compact and convex, for the sets $\mathrm{sim}_B(Y_j)$ are convex hulls of compact sets by Lemma 6. It remains to argue that $\mathrm{code}(\mathcal{Y}') = \mathrm{code}(\mathcal{Y})$. For this, consider any point $q \in \mathbb{R}^d$. If $q \notin B$ or $q = p$, then clearly $\mathrm{patt}_{\mathcal{Y}'}(q) = \mathrm{patt}_{\mathcal{Y}}(q)$. Otherwise, consider the radial projection centered at p onto ∂B. Let $q' \in \partial B$ be the projection of the point q. From the definition of simplification, $q \in Y'_j$ if and only if $q' \in Y_j$. Thus we have $\mathrm{patt}_{\mathcal{Y}'}(q) = \mathrm{patt}_{\mathcal{Y}}(q')$, and $\mathrm{code}(\mathcal{Y}') \subseteq \mathrm{code}(\mathcal{Y})$. For the reverse containment, recall that B avoids our set of representatives for $\mathcal{Y}$, except possibly p. The intersection pattern at p is the same in $\mathcal{Y}'$ as in $\mathcal{Y}$, so every member of $\mathrm{code}(\mathcal{Y})$ appears in $\mathrm{code}(\mathcal{Y}')$. Thus $\mathrm{code}(\mathcal{Y}') = \mathrm{code}(\mathcal{Y})$. However, $\mathcal{Y}'$ satisfies (i), (ii) and (iii) of Lemma 3, and so the fact that Y'_i is a proper subset of Y_i contradicts our choice of $\mathcal{Y}$ as an inclusion-minimal realization. Thus $\mathcal{Y}$ must be polygonal, and the result follows. $\qquad\square$

Step 3: Bounding the number of vertices

Our main task in this section is to prove the first part of Theorem 1, which says that every convex code admitting a closed 2-realization admits a polygonal realization in which the number of vertices is bounded by an explicit function of the code size. Indeed, one can express the question of whether a convex code $\mathcal{C}$ is 2-realizable by polygons with at most m vertices as a sentence in the language of fields. As the theory of real closed fields is decidable (originally shown by Tarski [13], see for example [1] for a modern exposition), this will imply Theorem 1.

Let $\mathcal{X} = (X_1, X_2, \ldots, X_n)$ be a closed 2-realization of a convex code $\mathcal{C}$. Let P be a set of representatives for $\mathcal{X}$ of size $|P| \leq 2^n$. We may assume that $\mathcal{X}$ is inclusion-minimal with respect to P, as in Lemma 3, which by Proposition 7 implies that each X_i is a closed polygon. We shall show that, if $\mathcal{X}$ has too many vertices, then we may shrink $\mathcal{X}$ by pulling a vertex inward. As $\mathcal{X}$ is inclusion-minimal, a bound on the number of vertices will follow.

Locating good vertices. The boundary of each X_i is made of *vertices* and *edges* that join them. We adopt the convention that the edges are relatively open, i.e., *edges do not contain their endpoints*. A point p is *good* with respect to $(X_1, \ldots, X_n)$ if p does not lie on an edge of any X_i. A point that is not good is called *bad*. A point p is *very good* with respect to $(X_1, \ldots, X_n)$ if it is good and $p \notin X_1 \setminus \mathrm{vertices}(X_1)$.

Lemma 8. *Suppose $X_1, \ldots, X_n$ are convex polygons in $\mathbb{R}^2$ and X_1 has K vertices. Then there is j such that X_j contains at least $\frac{1}{3^{n-1}(n-1)!}K$ very good vertices with respect to the tuple $(X_1, \ldots, X_n)$.*

Proof. We argue by induction on n. The base case $n = 1$ is trivial, so assume that $n \geq 2$. If at least $\frac{1}{3^{n-1}(n-1)!}K$ vertices of X_1 are good, then we are done with $j = 1$. Otherwise, let B be the set of all bad vertices in X_1. Note that $|B| \geq \left(1 - \frac{1}{3^{n-1}(n-1)!}\right)K \geq \frac{2}{3}K$. For each $i \neq 1$, let B_i be the set of those points in B that lie on some edge of X_i. Since $B = \bigcup_{i \geq 2} B_i$, it follows from the pigeonhole principle that one of $B_2, \ldots, B_n$ is large. Say, without loss of generality, that $|B_2| \geq \frac{2}{3}K/(n-1)$.

For each $v \in B_2$ choose a direction for which v is an extremal point of X_1. Let $u(v)$ be the vertex of $X'_2 \stackrel{\mathrm{def}}{=} \mathrm{conv}(X_1 \cup X_2)$ that is extremal for this direction (if there are two such vertices, pick
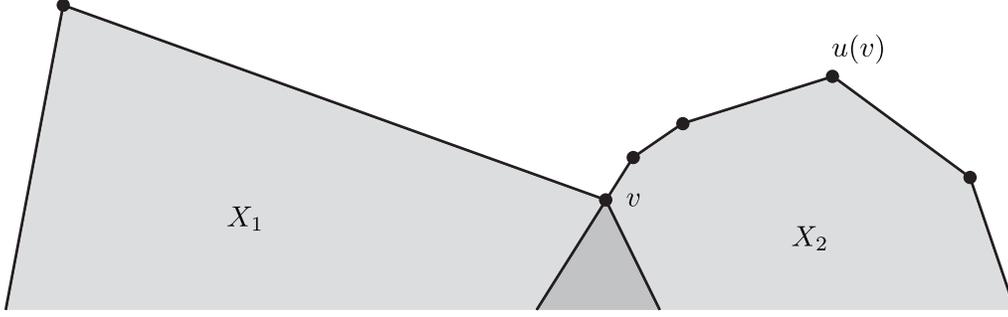
Figure 4: Associating bad vertices of X_1 to vertices of X_2'.

arbitrarily). Note that $u(v)$ is a vertex of X_2 that lies between two consecutive intersection points of ∂X_1 with ∂X_2. So, the map $v \mapsto u(v)$ from B_2 to vertices(X_2') is at most two-to-one. See Figure 4.

It follows that the polygon X_2' has at least $|B_2|/2 \geq \frac{K}{3(n-1)}$ vertices. By the induction hypothesis applied to the $(n-1)$-tuple $(X_2', X_3, X_4, \ldots, X_n)$ there is a $j \in \{2, 3, \ldots, n\}$ and a set of $\frac{1}{3^{n-2}(n-2)!} \cdot \frac{K}{3(n-1)}$ very good vertices in X_j (if $j > 2$) or X_2' (if $j = 2$) with respect to this $(n-1)$-tuple. Since $\frac{1}{3^{n-2}(n-2)!} \cdot \frac{K}{3(n-1)} = \frac{1}{3^{n-1}(n-1)!} K$, the proof is complete. $\qquad \square$

Corollary 9. *Suppose $\mathcal{X} = (X_1, \ldots, X_n)$ is an n-tuple of convex polygons in the plane, and P is a finite set of representatives for $\mathcal{X}$. If one of $X_1, \ldots, X_n$ has more than $3^n(n-1)!|P|$ vertices, then there exist two distinct $v, v' \in \mathbb{R}^2$ such that*

   (i) v is a vertex of some polygon in $\mathcal{X}$, and

   (ii) Both v and v' are good with respect to $\mathcal{X}$, and

   (iii) $v, v' \notin P$, and

   (iv) $\mathrm{patt}_{\mathcal{X}}(v) = \mathrm{patt}_{\mathcal{X}}(v')$.

Proof. By the previous lemma, some X_j contains more than $3|P|$ good vertices with respect to $\mathcal{X}$. More than $2|P|$ among them satisfy (iii). By the pigeonhole principle, a pair satisfies (iv) as well. $\qquad \square$

Pulling a vertex. Let $\mathcal{X}$, v and v' be as in Corollary 9. We will argue that $\mathcal{X}$ is not inclusion-minimal. By permuting the polygons in $\mathcal{X}$, we may assume that v is a vertex of $X_1, X_2, \ldots, X_m$ and is not a vertex of $X_{m+1}, \ldots, X_n$, for some $m \geq 1$. Let B be a ball around v small enough so that the only polygon vertex inside B is v. We shall shrink each of $X_1, X_2, \ldots, X_m$ by modifying inside B. Specifically, move the vertex v slightly toward v' along the line segment vv' to obtain a new vertex $\tilde{v}$. Then, for each $i \in [m]$, let a_i and b_i be the intersection points between the boundaries of X_i and B. If X_i is a line segment, then the points a_i and b_i are equal. Let $\widetilde{X}_i$ be obtained from X_i by replacing the line segments $a_i v$ and $b_i v$ with $a_i \tilde{v}$ and $b_i \tilde{v}$ respectively.

Since $\mathrm{patt}_{\mathcal{X}}(v) = \mathrm{patt}_{\mathcal{X}}(v')$ and $v \in X_1, \ldots, X_m$, it follows that $v' \in X_1, \ldots, X_m$ as well. Thus, the sets $\widetilde{X}_1, \ldots, \widetilde{X}_m$ are convex as long as $\tilde{v}$ is a sufficiently small perturbation of v. Now consider the tuple $\widetilde{X} = (\widetilde{X}_1, \ldots, \widetilde{X}_m, X_{m+1}, \ldots, X_n)$. Note that we may obtain $B \cap \widetilde{X}_1, \ldots, B \cap \widetilde{X}_m$ from

8

a) Before   **b)** After
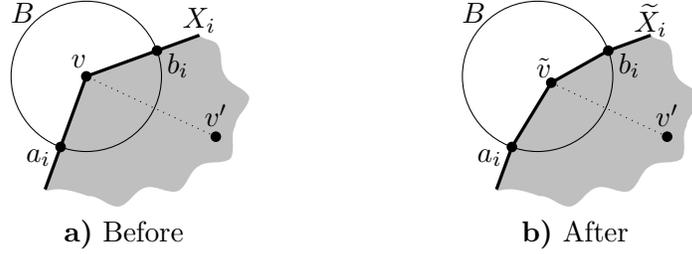
Figure 5: Pulling the vertex v toward v' in X_i.

$B \cap X_1, \ldots, B \cap X_m$ by a continuously deforming B while keeping its boundary fixed. Hence, the m-tuples $B \cap X_1, \ldots, B \cap X_m$ and $B \cap \widetilde{X}_1, \ldots, B \cap \widetilde{X}_m$ realize the same convex code. Furthermore, since the vertex v is good, the ball B can be chosen so small that that it does not intersect the boundary of $X_{m+1}, X_{m+2}, \ldots, X_n$, and so $\mathrm{code}(\mathcal{X}) = \mathrm{code}(\mathcal{X}')$. As $P \cap B = \emptyset$ for sufficiently small B, we conclude that $\mathcal{X}$ is not inclusion-minimal.

We immediately obtain the following corollary.

Corollary 10. *Let $\mathcal{Y}$ be an inclusion-minimal 2-realization of a convex code on $[n]$ with respect to a set of representatives P, as in Lemma 3. Then each set in $\mathcal{Y}$ is a polygon with no more than $3^n(n-1)!|P|$ vertices.*

As we may choose $|P| \le 2^n$, we conclude that every convex code that admits a closed convex realization in the plane admits a realization by polygons with at most $6^n(n-1)!$ vertices each.

Closed 2-realizability is decidable. Let $N = 6^n(n-1)!$. Every polygon with at most N vertices is an intersection of N closed halfspaces. Writing a polygon X_i as $\bigcap_{j=1}^{N}\{a_{i,j}x + b_{i,j}y \le c_{i,j}\}$ with $a_{i,j}, b_{i,j}, c_{i,j} \in \mathbb{R}$, we can write a formula asserting that a point (x, y) is in X_i,

$$\Psi_i(\vec{a}, \vec{b}, \vec{c}, x, y) \stackrel{\text{def}}{=} \forall j \in [N] \, (a_{i,j}x + b_{i,j}y \le c_{i,j}),$$

where we abbreviated $\vec{a} \stackrel{\text{def}}{=} (a_{i,j})_{i \in [n], j \in [N]}$, and similarly for $\vec{b}$ and $\vec{c}$. We can then write a formula asserting that $S \in \mathrm{code}(X_1, \ldots, X_n)$,

$$\Phi(S; \vec{a}, \vec{b}, \vec{c}) \stackrel{\text{def}}{=} \exists x, y \, \forall i \in [n] \Big(i \in S \iff \Psi_i(\vec{a}, \vec{b}, \vec{c}, x, y)\Big).$$

We can also write a formula asserting that X_i is bounded,

$$\Gamma_i(\vec{a}, \vec{b}, \vec{c}) \stackrel{\text{def}}{=} \exists r \left(x^2 + y^2 > r \implies \neg\Psi_i(\vec{a}, \vec{b}, \vec{c}, x, y)\right).$$

Then the assertion that a convex code $\mathcal{C} \subseteq 2^{[n]}$ admits a closed 2-realization can be written as

$$\exists \vec{a}, \vec{b}, \vec{c} \bigwedge_{i \in [n]} \Gamma_i(\vec{a}, \vec{b}, \vec{c}) \wedge \bigwedge_{S \in \mathcal{C}} \Phi(S; \vec{a}, \vec{b}, \vec{c}) \wedge \bigwedge_{S \notin \mathcal{C}} \neg\Phi(S; \vec{a}, \vec{b}, \vec{c}).$$

As mentioned above, the theory of real closed fields is decidable, and hence so is the question of whether a convex code admits a closed 2-realization.

Remark 11. Kunin, Lienkaemper and Rosen [9] were the first to explain that, for any fixed computable function of n and d, there is an algorithm to decide whether a convex code has a polytopal d-realization with the total number of vertices bounded by this function (in fact they considered a bound on the number of facets, but this is the same up to a slight modification of the computable function in question). Seeking a combinatorial argument that d-realizability is decidable, they reduced the problem of finding a d-realization with a bounded number of vertices to the problem of determining whether certain oriented matroids are representable. We have opted for the direct statements above, since deciding representability of oriented matroids is complete with respect to the theory of real closed fields, and hence the reduction to oriented matroids does not improve the computational complexity of the decision problem.

Complexity: Realizations requiring many vertices

In this section we prove Theorem 2, which gives a lower bound on the number of vertices in a realization $\mathcal{X}$ of any code $\mathcal{C}$ by closed polygons.

Let ℓ be some vertical line. The intersection pattern of $\mathcal{X}$ at $p \in \ell$ is determined by the position of p relative to the intersection points of $\partial X_1, \ldots, \partial X_n$ with ℓ. Let $P(\ell)$ be the set of intersection patterns that occur on ℓ.

Imagine starting with ℓ to the left of all $X_1, \ldots, X_n$, and sweeping ℓ to the right across the plane. The relative positions of the intersection points between ℓ and the various ∂X_i changes only as ℓ passes through some certain *special* points: leftmost and rightmost points of the polygons, and the points where two polygon edges intersect transversely. By rotating $\mathbb{R}^2$ slightly, we may ensure that the x-coordinates of special points are distinct. However, it might still happen that a point p is special for several different reasons: for example, if several different pairs of polygons intersect at p. In this case, we speak of *multiplicity* of p, and denote it by $m(p)$.

As the line ℓ passes through a point p of multiplicity $m(p)$, the set $P(\ell)$ changes, and we may obtain intersection patterns that did not appear to the left of p. We obtain a new pattern when ℓ reaches p, and then at most one new pattern per each polygonal segment through p once ℓ passes over p. Hence, the total number of new patterns is at most $1 + 2m(p) \leq 4m(p)$. Letting M be the number of special points (with multiplicity), it follows that $|\mathcal{C}| \leq 4M$.

Let N be the total number of vertices among the polygons $X_1, \ldots, X_n$. Since a line segment belonging to one of $\partial X_1, \ldots, \partial X_n$ intersects $\partial X_1, \ldots, \partial X_n$ transversely in a total of at most $2n - 2$ points, it follows that $M \leq 2n + (2n - 2)N$. Therefore, $|\mathcal{C}| \leq 8nN$, as desired.

Bonus: From closed to open realizations

In this section we will prove the second part of Theorem 1. We first show that when we replace sets by their closures in an open realization, new intersection patterns can only arise inside convex regions with empty interior.

Lemma 12. *Let $\mathcal{U} = (U_1, \ldots, U_n)$ be a tuple of convex sets in $\mathbb{R}^d$, and define $\mathcal{X} = (X_1, \ldots, X_n)$ where $X_i \stackrel{\text{def}}{=} \text{cl}\, U_i$. If $c \in \text{code}(\mathcal{X}) \setminus \text{code}(\mathcal{U})$, then $\bigcap_{i \in c} X_i$ has empty interior.*

Proof. Suppose for contradiction that there exists $c \in \mathrm{code}(\mathcal{X}) \setminus \mathrm{code}(\mathcal{U})$ so that $U \stackrel{\mathrm{def}}{=} \mathrm{int}\left(\bigcap_{i \in c} X_i\right)$ is nonempty. Let $p \in U$, and let q be a point with $\mathrm{patt}_{\mathcal{X}}(q) = c$. Since $q \in \bigcap_{i \in c} X_i$, the line segment pq is contained in U except possibly for the point q. Since all X_i are closed, q lies a positive distance from X_i for all $i \notin c$. Hence we may choose $q' \in U$ on the segment pq so that $\mathrm{patt}_{\mathcal{X}}(q') = c$. Since U_i is convex, $\mathrm{int}\, X_i \subseteq U_i$. Hence, $q' \in \mathrm{int}\, X_i \subseteq U_i$ for all $i \in c$, and so $\mathrm{patt}_{\mathcal{U}}(q') = c$. This contradicts the assumption that $c \notin \mathrm{code}(\mathcal{U})$. $\qquad\square$

Reduction from open to closed realizations. Let $\mathcal{U} = (U_1, \ldots, U_n)$ be a tuple of bounded convex open sets in $\mathbb{R}^2$. Let $\mathcal{X} = (X_1, \ldots, X_n)$ be the tuple of compact convex sets in $\mathbb{R}^2$ with $X_i \stackrel{\mathrm{def}}{=} \mathrm{cl}\, U_i$ for all $i \in [n]$. For every nonempty $\sigma \subseteq [n]$ for which $\bigcap_{i \in \sigma} X_i$ is nonempty but has empty interior, let L_σ be a line containing the set $\bigcap_{i \in \sigma} X_i$. Fix a set of representatives P for $\mathcal{X}$ with the following properties:

i) for every $c \in \mathrm{code}(\mathcal{U})$, the set P contains the vertices of a triangle contained in $\bigcap_{i \in c} U_i$ whose interior contains a point p with $\mathrm{patt}_{\mathcal{U}}(p) = c$,

ii) for every pair (L_σ, X_i) with $L_\sigma \cap U_i \neq \emptyset$, the set P contains the vertices of a quadrilateral contained in X_i which has one of its diagonals equal to $L_\sigma \cap X_i$.

To satisfy the conditions (i) and (ii), first select, for each $c \in \mathrm{code}(\mathcal{U})$, a point p such that $\mathrm{patt}_{\mathcal{U}}(p) = c$, and then choose an appropriate small triangle or quadrilateral containing p. This way we may guarantee

$$|P| \leq |\mathrm{code}(\mathcal{X})| + 3|\mathrm{code}(\mathcal{U})| + 4n2^n$$
$$\leq 2^n + 3 \cdot 2^n + 4n2^n = 4(n+1)2^n.$$

Now let $\mathcal{X}' = (X_1', \ldots, X_n')$ be an inclusion-minimal realization of $\mathrm{code}(\mathcal{X})$ with respect to the set of representatives P, as in Lemma 3. By Corollary 10 the number of vertices among all X_i' is at most $3^n (n-1)!|P| \leq 4(n+1) \cdot 6^n (n-1)!$. Define $\mathcal{U}' = (U_1', \ldots, U_n')$ to be the tuple with $U_i' \stackrel{\mathrm{def}}{=} \mathrm{int}\, X_i'$ for all $i \in [n]$. We claim that $\mathrm{code}(\mathcal{U}') = \mathrm{code}(\mathcal{U})$, which will establish the second part of Theorem 1. The relationships between the tuples $\mathcal{U}, \mathcal{X}, \mathcal{X}'$ and $\mathcal{U}'$ is depicted on the right.

$$\mathcal{X} \xrightarrow{\mathrm{incl.min.}} \mathcal{X}'$$
$$\mathrm{cl}\,\Big\uparrow \qquad\qquad \Big\downarrow\, \mathrm{int}$$
$$\mathcal{U} \qquad\qquad \mathcal{U}'$$
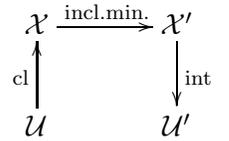
The inclusion $\mathrm{code}(\mathcal{U}') \subseteq \mathrm{code}(\mathcal{U})$**:** Note that the condition (ii) implies that

$$L_\sigma \cap U_i' = L_\sigma \cap U_i \quad \text{for every } i \in [n]. \tag{1}$$

Indeed, let a, b be the end points of the line segment $L_\sigma \cap X_i$. Let c, d be the other two vertices of the quadrilateral from the condition (ii). Suppose first that p is an interior point of the line segment $L_\sigma \cap X_i$. Then $p \in \mathrm{int}\,\mathrm{conv}\{a, b, c, d\}$. Since $a, b, c, d \in X_i \cap P$, it follows that $a, b, c, d \in X_i' \cap P$ and hence $p \in \mathrm{int}\,\mathrm{conv}\{a, b, c, d\} \subseteq U_i'$. Suppose next that p is either a or b. In this case, since U_i is open it follows that $p \notin U_i$. Also, $p \in \partial X_i'$ and so $p \notin U_i'$. So, neither a nor b is in U_i' implying that the intersection of the line L_σ with U_i' is the open line segment (a, b). Because $L_\sigma \cap U_i = (a, b)$, this proves (1).

From (1) it follows that $\mathrm{patt}_{\mathcal{U}'}(p) = \mathrm{patt}_{\mathcal{U}}(p)$ if $p \in L_\sigma$. Let $L = \bigcup_\sigma L_\sigma$ be the union of all L_σ's. It remains to show that $\mathrm{patt}_{\mathcal{U}'}(p) \in \mathrm{code}(\mathcal{U})$ for $p \notin L$. So, fix a point $p \notin L$, and put $c \overset{\mathrm{def}}{=} \mathrm{patt}_{\mathcal{X}'}(p)$. Since any intersection pattern in $\mathrm{code}(\mathcal{X}) \setminus \mathrm{code}(\mathcal{U})$ arise only inside L by Lemma 12, it follows any intersection patterns in $\mathrm{code}(\mathcal{X}') \setminus \mathrm{code}(\mathcal{U})$ also arise only inside L. As $p \notin L$, this implies that $c \in \mathrm{code}(\mathcal{U})$. So by the choice of P in the condition (i) the set $\bigcap_{i \in c} X_i'$ has nonempty interior. Let q be a point in the interior of $\bigcap_{i \in c} X_i'$. For sufficiently small $\varepsilon > 0$, consider the point $p_\varepsilon \overset{\mathrm{def}}{=} p + \varepsilon(p - q)$. Since $q \in \mathrm{int}\, X_i'$ for every $i \in c$, it follows that $p_\varepsilon \in X_i'$ if and only if $p_\varepsilon \in \mathrm{int}\, X_i'$, for small enough ε. As $\mathrm{int}\, X_i' = U_i'$, we conclude that $\mathrm{patt}_{\mathcal{U}'}(p) = \mathrm{patt}_{\mathcal{X}'}(p_\varepsilon)$. Furthermore, since $p \notin L$ and $\varepsilon > 0$ is small, we also have $p_\varepsilon \notin L$. Again using the fact that the intersection patterns in $\mathrm{code}(\mathcal{X}') \setminus \mathrm{code}(\mathcal{U})$ arise only inside L, we have $\mathrm{patt}_{\mathcal{U}'}(p) = \mathrm{patt}_{\mathcal{X}'}(p_\varepsilon) \in \mathrm{code}(\mathcal{U})$, as promised.

The inclusion $\mathrm{code}(\mathcal{U}) \subseteq \mathrm{code}(\mathcal{U}')$: Let $c \in \mathrm{code}(\mathcal{U})$. Consider the point p and the triangle $\triangle$ in $\bigcap_{i \in c} U_i$ from the condition (i). Since P contains the vertices of $\triangle$, the same is true of the X_i' for $i \in c$. Because $p \in \mathrm{int}\, \triangle$ it follows that $p \in U_i'$ for all $i \in c$. As $U_i' \subseteq U_i$ for all $i \in [n]$, we also have $p \notin U_i'$ for all $i \notin c$. Therefore, $\mathrm{patt}_{\mathcal{U}'}(p) = \mathrm{patt}_{\mathcal{U}}(p)$, and so $c \in \mathrm{code}(\mathcal{U}')$.

Open 2-realizability is decidable. This is similar to the decidability of closed 2-realizability. There are two minor differences. First, the number of polygon vertices is bounded by $4(n+1) \cdot 6^n (n-1)!$ instead of $6^n (n-1)!$. Second, the formula Ψ_i should use strict inequalities:

$$\Psi_i(\vec{a}, \vec{b}, \vec{c}, x, y) \overset{\mathrm{def}}{=} \forall j \in [N]\, (a_{i,j} x + b_{i,j} y < c_{i,j}).$$

The rest of the argument is the same.

Epilogue: comments about the general case

In this paper we restricted our attention to realizations by open or closed bounded sets in the plane. Extending our techniques to arbitrary convex sets in the plane does not seem to pose serious challenges, but would make the argument more technical. However, the situation in higher dimensions appears more difficult, perhaps fundamentally so.

For example, there exist convex sets in $\mathbb{R}^3$ that are not polyhedral at a neighborhood of a point, but which cannot be made locally polyhedral using the simpification process in Definition 5 centered at the point. Namely, consider the set C, which is the convex hull of

$$\{(x, y, z) \mid z = 0, x^2 + y^2 = 1\} \cup \{(x, y, z) \mid y = 0, x^2 + z^2 = 1\}.$$

The set C is the convex hull of two circles which pass through the point $p = (1, 0, 0)$ transversely (see Figure 6). If B is a small ball centered at p, then $\mathrm{cone}_p(\partial B \cap C)$ is not even a convex set, no matter how small B is. One could attempt to remedy this by replacing B with a neighborhood of p that is shaped differently than a ball, but the fact that we need a simplification process that works simultaneously for many sets introduces further difficulties. Even if one does obtain polytopal realizations in higher dimensions, a computable bound on the number of vertices is not immediate, as our Lemma 8 does not have an obvious generalization beyond the plane.
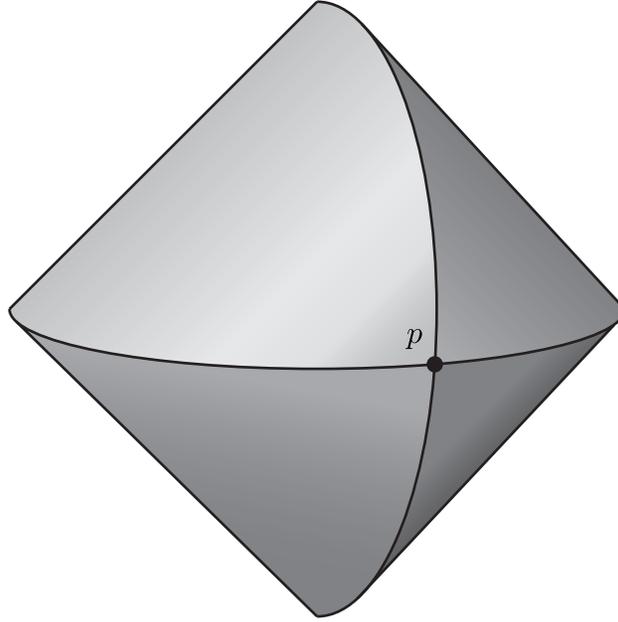
Figure 6: A convex set in $\mathbb{R}^3$ that is difficult to locally simplify.

These difficulties suggest that recognizing 3-realizable convex codes is harder than recognizing 2-realizable convex codes. The results of Kunin, Lienkaemper and Rosen [9] suggest that recognizing 2-realizable convex codes is already at least NP-hard. In addition, Tancer [11, Section 6] showed that recognizing d-representable complexes (for any fixed $d \geq 2$) is NP-hard. It is entirely possible that recognizing 3-realizable convex codes is in fact an undecidable problem.

References

[1] Saugata Basu, Richard Pollack, and Marie-Françoise Roy. *Algorithms in real algebraic geometry*, volume 10 of *Algorithms and Computation in Mathematics*. Springer-Verlag, Berlin, 2003.

[2] Aaron Chen, Florian Frick, and Anne Shiu. Neural codes, decidability, and a new local obstruction to convexity. *SIAM Journal on Applied Algebra and Geometry*, 3(1):44–66, 2019.

[3] Joshua Cruz, Chad Giusti, Vladimir Itskov, and Bill Kronholm. On open and closed convex codes. *Discrete & Computational Geometry*, 61:247–270, 2016.

[4] Carina Curto, Elizabeth Gross, Jack Jeffries, Katherine Morrison, Mohamed Omar, Zvi Rosen, Anne Shiu, and Nora Youngs. What makes a neural code convex? *SIAM J. Appl. Algebra Geom.*, 1(1):222–238, 2017.

[5] Carina Curto, Vladimir Itskov, Alan Veliz-Cuba, and Nora Youngs. The neural ring: an algebraic tool for analyzing the intrinsic structure of neural codes. *Bulletin of Mathematical Biology*, 75(9):1571–1611, 2013.

[6] Megan K. Franke and Samuel Muthiah. Every binary code can be realized by convex sets. *Advances in Applied Mathematics*, 99:83–93, 2017.

[7] R. Amzi Jeffs. Embedding dimension phenomena in intersection complete codes. *Selecta Mathematica*, 28, 2022. arXiv:1909.13406.

[8] R. Amzi Jeffs. Open, closed, and non-degenerate embedding dimensions of neural codes. *Discrete & Computational Geometry*, 2022. (To appear) arXiv:2111.01010.

[9] Alexander Kunin, Caitlin Lienkaemper, and Zvi Rosen. Oriented matroids and combinatorial neural codes. 2020. arXiv:2002.03542.

[10] Caitlin Lienkaemper, Anne Shiu, and Zev Woodstock. Obstructions to convexity in neural codes. *Advances in Applied Mathematics*, 85:31–59, 2017.

[11] Martin Tancer. d-collapsibility is NP-complete for $d \geq 4$. *Chic. J. Theoret. Comput. Sci.*, pages Article 3, 32, 2010.

[12] Martin Tancer. Intersection patterns of convex sets via simplicial complexes: a survey. In *Thirty essays on geometric graph theory*, pages 521–540. Springer, New York, 2013. arXiv:1102.0417.

[13] Alfred Tarski. *A decision method for elementary algebra and geometry*. University of California Press, Berkeley-Los Angeles, Calif., 1951. 2nd ed.