EMBEDDING DIMENSION PHENOMENA IN INTERSECTION COMPLETE CODES

R. AMZI JEFFS

ABSTRACT. Two tantalizing invariants of a combinatorial code $\mathcal{C} \subseteq 2^{[n]}$ are $\operatorname{cdim}(\mathcal{C})$ and $\operatorname{odim}(\mathcal{C})$, the smallest dimension in which $\mathcal{C}$ can be realized by convex closed or open sets, respectively. Cruz, Giusti, Itskov, and Kronholm showed that for intersection complete codes $\mathcal{C}$ with $m+1$ maximal codewords, $\operatorname{odim}(\mathcal{C})$ and $\operatorname{cdim}(\mathcal{C})$ are both bounded above by $\max\{2, m\}$. Results of Lienkaemper, Shiu, and Woodstock imply that odim and cdim may differ, even for intersection complete codes. We add to the literature on open and closed embedding dimensions of intersection complete codes with the following results:

- If $\mathcal{C}$ is a simplicial complex, then $\operatorname{cdim}(\mathcal{C}) = \operatorname{odim}(\mathcal{C})$,
- If $\mathcal{C}$ is intersection complete, then $\operatorname{cdim}(\mathcal{C}) \leq \operatorname{odim}(\mathcal{C})$,
- If $\mathcal{C} \subseteq 2^{[n]}$ is intersection complete with $n \geq 2$, then $\operatorname{cdim}(\mathcal{C}) \leq \min\{2d+1, n-1\}$, where d is the dimension of the simplicial complex of $\mathcal{C}$, and
- For each simplicial complex $\Delta \subseteq 2^{[n]}$ with $m \geq 2$ facets, the code $\mathcal{S}_\Delta := (\Delta * (n+1)) \cup \{[n]\} \subseteq 2^{[n+1]}$ is intersection complete, has $m+1$ maximal codewords, and satisfies $\operatorname{odim}(\mathcal{S}_\Delta) = m$. In particular, for each $n \geq 3$ there exists an intersection complete code $\mathcal{C} \subseteq 2^{[n]}$ with $\operatorname{odim}(\mathcal{C}) = \binom{n-1}{\lfloor (n-1)/2 \rfloor}$.

A key tool in our work is the study of sunflowers: arrangements of convex open sets in which the sets simultaneously meet in a central region, and nowhere else. We use Tverberg's theorem to study the structure of "k-flexible" sunflowers, and consequently obtain new lower bounds on $\operatorname{odim}(\mathcal{C})$ for intersection complete codes $\mathcal{C}$.

1. INTRODUCTION

In [6], Curto, Itskov, Veliz-Cuba, and Youngs introduced *convex codes* to mathematically model stimulus reconstruction from neural data, particularly in the context of hippocampal place cells. Classifying and understanding convex codes have been active areas of recent mathematical research, bringing together tools and perspectives from topology [4, 21], algebra [9, 11, 22], and discrete geometry [2, 12, 13, 16, 19]. A complete classification of convex codes is far out of reach for the moment, but progress can yield new techniques for analyzing neural data, as well as a deeper understanding of the mathematical theory of convex sets. In this paper, we give new bounds on the open and closed embedding dimensions of intersection complete codes, and families of examples where these bounds are tight. In particular, we provide infinite families of intersection complete codes for which open embedding dimension grows exponentially in the number of neurons, while closed embedding dimension grows only linearly.

Before stating our results we recall some definitions and frame our main questions of study. A convex code (see Definition 1.1 below) is a special case of a *combinatorial code*, which is a

Date: November 16, 2021.

2010 Mathematics Subject Classification. 32F27, 52A20, 52C99, 52A35.

Department of Mathematics. University of Washington, Seattle, Wa 98195.

Jeffs' research is supported by a graduate fellowship from NSF grant DGE-1761124.

collection of subsets of $[n] := \{1, \dots, n\}$. Due to the biological motivation behind our work, we think of the elements of $[n]$ as *neurons*, and each element of a code as recording a set of neurons which fired together in a small window of time.

The elements of a code are called *codewords*, and for concision we often omit braces and commas when writing codewords. For example, we may write 124 instead of $\{1, 2, 4\}$. The *weight* of a codeword is simply the number of neurons it contains. We will often think of a code as a partially ordered set under containment—for example, we may speak of maximal codewords, which are not properly contained in any other codeword. When writing down a specific code, we will bold the maximal codewords.

Codes can arise abstractly when one wishes to describe how a certain collection of sets covers a space, as follows. Let X be a set and $\mathcal{U} = \{U_1, \dots, U_n\}$ a collection of subsets of X. One may form the *code of* $\mathcal{U}$ *in* X, a combinatorial code whose codewords describe how the U_i intersect and cover one another:

$$\operatorname{code}(\mathcal{U}, X) := \left\{ \sigma \subseteq [n] \;\middle|\; \bigcap_{i \in \sigma} U_i \setminus \bigcup_{j \in [n] \setminus \sigma} U_j \neq \emptyset \right\}.$$

The region $\bigcap_{i\in\sigma} U_i \setminus \bigcup_{j\in[n]\setminus\sigma} U_j$ is called the *atom* of σ, and denoted by $\mathcal{A}^\sigma_\mathcal{U}$. The space X is called the *ambient space* or *stimulus space*, and the U_i are called *receptive fields* or *firing regions*. Note that the receptive fields are indexed by neurons. If $\mathcal{C} = \operatorname{code}(\mathcal{U}, X)$, then the collection $\mathcal{U}$ is called a *realization* of $\mathcal{C}$ in X. For concision, we will write U_σ for $\bigcap_{i\in\sigma} U_i$ (and similarly define V_σ when working with various V_i), and adopt the convention that $U_\emptyset = X$.

Unless otherwise specified, throughout this paper the ambient space will be $\mathbb{R}^d$, and the U_i will be (possibly empty) convex sets that are either all open, or all closed. We will write $\operatorname{code}(\mathcal{U})$ instead of $\operatorname{code}(\mathcal{U}, \mathbb{R}^d)$ when the ambient dimension is clear. We will also adopt the usual convention in the study of convex codes that $\emptyset$ is contained in all codes, i.e. that there is always a point in the ambient space not covered by any U_i. In our examples we illustrate open sets with a solid border so that our figures are clean and readable. The accompanying text will always specify whether we are regarding the sets in question as open or closed. Finally, we will always assume $n \geq 1$ (i.e. that we are not working with an empty set of neurons).

Definition 1.1. A code $\mathcal{C} \subseteq 2^{[n]}$ is called an *open convex* code if it has a realization consisting of convex open sets in $\mathbb{R}^d$. Similarly, $\mathcal{C}$ is called *closed convex* if it has a realization consisting of closed convex sets in $\mathbb{R}^d$.

Example 1.2. The figure below shows a realization in $\mathbb{R}^2$ of the (open/closed) convex code $\mathcal{C} = \{\mathbf{123}, 12, 23, 2, 3, \emptyset\}$. The atom $\mathcal{A}^{23}_\mathcal{U}$ is highlighted in grey.

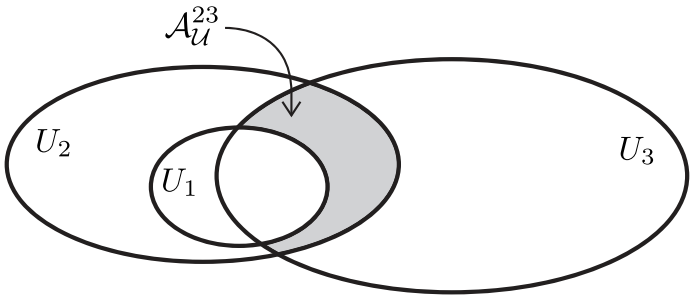

In the neuroscientific context mentioned above, open convex codes are of greater interest than closed convex codes, since receptive fields have been experimentally observed to be

full-dimensional (see [4, Figure 1] for example). However, we will also study closed convex codes in this paper to build on the work of [3], and to contrast their behavior to that of open convex codes. Moreover, it is of broad mathematical interest to develop our understanding of closed convex sets and their intersection patterns, since they are ubiquitous in fields such as optimization and discrete geometry (see for example [1, Section 3], [20, Chapters 6 and 8], and [23, Chapter 8]).

The study of convex codes asks two main questions. First, given a code $\mathcal{C} \subseteq 2^{[n]}$, **when can we find a (closed or open) convex realization of $\mathcal{C}$?** Second, if we can find a realization, **what is the smallest dimension in which we can do so?** Formally, we wish to investigate the open and closed embedding dimensions of combinatorial codes, described below.

Definition 1.3. Let $\mathcal{C} \subseteq 2^{[n]}$ be a code. The *open embedding dimension* of $\mathcal{C}$, denoted by $\operatorname{odim}(\mathcal{C})$, is the smallest d so that $\mathcal{C}$ has a realization in $\mathbb{R}^d$ consisting of convex open sets, or ∞ if no convex open realization exists. Similarly, the *closed embedding dimension*, denoted by $\operatorname{cdim}(\mathcal{C})$, is the smallest dimension in which $\mathcal{C}$ has a closed convex realization, or ∞ if none exist.

Remark 1.4. The study of open and closed embedding dimension generalizes the the study of d-representable simplicial complexes, a classical topic in discrete geometry (see [25] for a survey). A simplicial complex Δ is said to be *d-representable* if one can find a collection $\mathcal{U}$ of convex sets in $\mathbb{R}^d$ such that for each $\sigma \subseteq [n]$ we have $\sigma \in \Delta$ if and only if $U_\sigma \neq \emptyset$ (equivalently, Δ is the smallest simplicial complex containing $\operatorname{code}(\mathcal{U})$). In this context the requirement that all sets in $\mathcal{U}$ be open or closed is immaterial; one may work with open or closed convex sets interchangeably. Thus the differences that can arise between open and closed embedding dimensions of codes (which we will see starkly in Corollary 6.3) provide interesting evidence that the study of convex codes significantly generalizes the study of d-representable complexes.

Note that the convex realization in Example 1.2 is not minimal with respect to dimension, since we could flatten the U_i into (closed or open) intervals to obtain a realization in $\mathbb{R}^1$. Thus $\operatorname{odim}(\mathcal{C}) = \operatorname{cdim}(\mathcal{C}) = 1$ for $\mathcal{C} = \{\mathbf{123}, 12, 23, 2, 3, \emptyset\}$.

In this paper we will study codes that are *intersection complete*: the intersection of any two codewords is again a codeword. An important result of [3] is that intersection complete codes are always open and closed convex. More specifically, we have the following:

Theorem 1.5 (Special case of Theorem 1.2 of [3])**.** *Let $\mathcal{C} \subseteq 2^{[n]}$ be an intersection complete code with $m+1$ maximal codewords. Then*

$$\max\{\operatorname{odim}(\mathcal{C}), \operatorname{cdim}(\mathcal{C})\} \leq \max\{2, m\}.$$

We will prove a new bound on closed embedding dimension which improves the bound in Theorem 1.5 for many intersection complete codes (see Theorem 1.9). In contrast to the closed case, we will show that the bound $\operatorname{odim}(\mathcal{C}) \leq \max\{2, m\}$ can be tight for any choice of $m \geq 2$ (see Theorem 1.11). Except where stated otherwise, every code we work with in this paper is intersection complete.

A special case of intersection complete codes is that of a simplicial complex. For simplicial complexes, open and closed embedding dimensions are equal. Although this result is well known among the neural codes community, we are not aware of any written proofs. We provide one below.

Theorem 1.6. *Let $\mathcal{C} \subseteq 2^{[n]}$ be a simplicial complex. Then* $\operatorname{cdim}(\mathcal{C}) = \operatorname{odim}(\mathcal{C})$.

Proof. In Theorem 1.8, we will show that $\operatorname{cdim}(\mathcal{C}) \leq \operatorname{odim}(\mathcal{C})$. Thus we just need to prove that $\operatorname{odim}(\mathcal{C}) \leq \operatorname{cdim}(\mathcal{C})$. Let $\mathcal{V} = \{V_1, \ldots, V_n\}$ be a closed realization of $\mathcal{C}$ in $\mathbb{R}^{\operatorname{cdim}(\mathcal{C})}$. By intersecting all the V_i with a sufficiently large closed ball, we may assume that they are bounded, and hence compact. For each nonempty codeword $c \in \mathcal{C}$, choose a point $p_c \in \mathcal{A}^c_{\mathcal{V}}$. By compactness, each p_c has positive distance to any set V_i that does not contain it. Likewise, any choice of $\sigma, \tau \subseteq [n]$ for which V_σ and V_τ are disjoint, the sets V_σ and V_τ must have positive distance between one another. Thus we may choose ε such that replacing the V_i by their Minkowski sums with an open ε-ball neither causes any V_i to cover some p_c it did not before, nor causes disjoint V_σ and V_τ to intersect. This creates a collection of convex open sets whose code contains all the codewords of $\mathcal{C}$, and no new maximal codewords. Since $\mathcal{C}$ is a simplicial complex, this is exactly a convex open realization of $\mathcal{C}$. □

Example 1.7. Consider the code $\mathcal{C} = \{\mathbf{123}, \mathbf{34}, 12, 13, 23, 1, 2, 3, 4, \emptyset\}$, and note that $\mathcal{C}$ is a simplicial complex. The lefthand side of the figure below shows a realization of $\mathcal{C}$ in $\mathbb{R}^2$ with closed convex sets, as well as possible choices of points p_c for $c \in \mathcal{C}$ as used in the proof above. The righthand side shows the open realization given in the proof above, which results from adding a small ε-ball to each V_i.

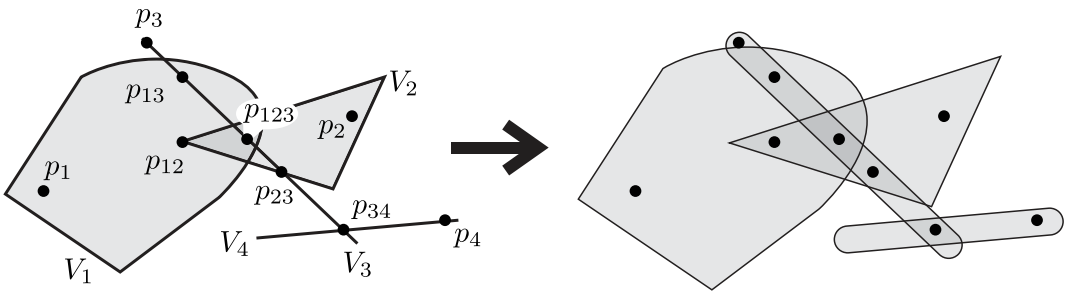

Can the above techniques be extended to realizations of codes that are not simplicial complexes? The answer in general is no, even for intersection complete codes, a fact which was first observed implicitly in the results of [13, 19]. Corollary 6.3 will yield a plethora of examples of intersection complete codes on n neurons that are closed convex in $\mathbb{R}^{n-1}$, but not open convex in $\mathbb{R}^{n-1}$. For such codes, adding an ε-ball to sets in a closed realization in $\mathbb{R}^{n-1}$ will always fail to produce an open realization.

The following theorems are the main contributions of this work, and give us a handle on how open and closed dimension behave for intersection complete codes.

Theorem 1.8. *Let $\mathcal{C} \subseteq 2^{[n]}$ be an intersection complete code. Then* $\operatorname{cdim}(\mathcal{C}) \leq \operatorname{odim}(\mathcal{C})$.

It is known that this inequality may be strict for intersection complete codes $\mathcal{C} \subseteq 2^{[n]}$, as mentioned above. In fact, the gap may be quite large: Theorem 1.9 implies that $\operatorname{cdim}(\mathcal{C}) \leq n - 1$, while Corollary 6.3 says that $\operatorname{odim}(\mathcal{C})$ **may be exponential in** n.

Theorem 1.9. *Let $\mathcal{C} \subseteq 2^{[n]}$ be an intersection complete code with $n \geq 2$, and d be one less than the weight of the largest codeword in $\mathcal{C}$ (i.e. $d = \dim(\Delta(\mathcal{C}))$). Then* $\operatorname{cdim}(\mathcal{C}) \leq \min\{2d + 1, n - 1\}$.

This bound is known to be tight. For every $d \geq 0$, [24] describes a d-dimensional simplicial complex on n vertices whose closed embedding dimension is exactly $2d + 1$ (which, in the

family given, is the same as $\min\{2d+1, n-1\}$). Our proof of Theorem 1.9 is inspired by a construction of Wegner and Perel'man which was originally used to show that every d-dimensional simplicial complex is $(2d+1)$-representable (see [25, Theorem 3.1]).

Interestingly, the bound in Theorem 1.9 does not hold for $\operatorname{odim}(\mathcal{C})$. In fact, Theorem 1.11 below gives us a way to construct numerous examples of intersection complete codes for which $\operatorname{odim}(\mathcal{C}) \gg \min\{2d+1, n-1\}$.

Definition 1.10. Let $\Delta \subseteq 2^{[n]}$ be a simplicial complex. Define $\mathcal{S}_\Delta \subseteq 2^{[n+1]}$ to be the code

$$\mathcal{S}_\Delta := (\Delta * (n+1)) \cup \{[n]\},$$

where $\Delta * (n+1)$ denotes the cone over Δ with apex $n+1$.

Theorem 1.11. *Let $\Delta \subseteq 2^{[n]}$ be a simplicial complex with $m \geq 2$ facets. Then $\mathcal{S}_\Delta$ is an intersection complete code with $m+1$ maximal codewords, and $\operatorname{odim}(\mathcal{S}_\Delta) = m$.*

A key tool in proving Theorem 1.11 is an application of a "sunflower theorem" that we proved in [13]. In this paper, we will generalize this theorem to "k-flexible" sunflowers of convex open sets, defined formally below. These are collections of convex open sets which have a common intersection, but no more than k of which overlap outside of this common intersection.

Definition 1.12. Let $\mathcal{U} = \{U_1, \ldots, U_n\}$ be a collection of convex sets in $\mathbb{R}^d$ and let $\mathcal{C} = \operatorname{code}(\mathcal{U})$. The collection $\mathcal{U}$ is called a *k-flexible sunflower* if $[n] \in \mathcal{C}$, and all other codewords have weight at most k. The U_i are called *petals* and $U_{[n]}$ is called the *center* of $\mathcal{U}$.

The following theorem tells us that if a k-flexible sunflower $\mathcal{U}$ in $\mathbb{R}^d$ has "enough" petals, then sampling a point from each petal and taking the convex hull always covers a point in the center of $\mathcal{U}$. Our proof of this theorem is given in Section 7 and relies on an application of Tverberg's theorem.

Theorem 1.13. *Let $\mathcal{U} = \{U_1, \ldots, U_n\}$ be an open k-flexible sunflower in $\mathbb{R}^d$. Suppose that $n \geq dk+1$, and for each $i \in [n]$ let $p_i \in U_i$. Then $\operatorname{conv}\{p_1, \ldots, p_n\}$ contains a point in the center of $\mathcal{U}$. Moreover, if $d \geq 2$ this result may fail when $n < dk+1$.*

By considering a set of line segments in $\mathbb{R}^2$ which meet at a point, one can see that this result does not hold for closed convex sets.

Example 1.14. Consider the open 2-flexible sunflower $\{U_1, U_2, U_3, U_4, U_5\}$ in $\mathbb{R}^2$ below. The center of this sunflower is the unit square highlighted in dark gray. Note that $d = 2$, $k = 2$, and $n = 5$. Thus $n \geq dk+1$, and so Theorem 1.13 applies. Indeed, any choice of $p_1 \in U_1, \ldots, p_5 \in U_5$ has the property that $\operatorname{conv}\{p_1, p_2, p_3, p_4, p_5\}$ intersects the center of the sunflower. One choice of such points is shown below.

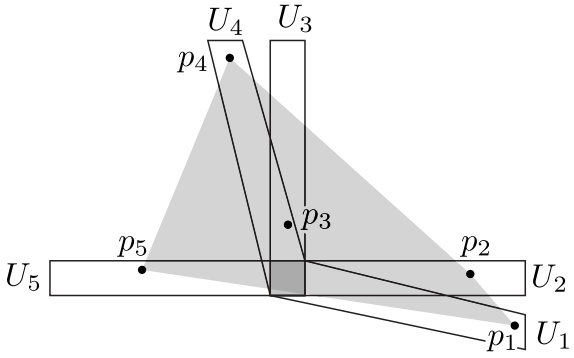

Observe that deleting U_5 yields a 2-flexible sunflower in $\mathbb{R}^2$ for which the conclusion of Theorem 1.13 does not hold: the set $\text{conv}\{p_1, p_2, p_3, p_4\}$ does not intersect the center of $\{U_1, U_2, U_3, U_4\}$.

In Section 2 we will recall some relevant background material. The subsequent sections are devoted to proving the theorems stated above, with one self-contained section per theorem. An exception to this is Section 5, which provides important supporting results and context for Section 6.

Section 8 describes a new family $\mathcal{T}_n$ of intersection complete codes, and initiates the study of their open embedding dimensions. The codes $\mathcal{T}_n$ are related to sunflowers, but the theorems that we prove regarding sunflowers are not sufficient to precisely determine $\text{odim}(\mathcal{T}_n)$.

Section 9 provides a unifying capstone to our results. We examine the families of codes from Sections 5, 6, and 8 in the context of a partially ordered set $\mathbf{P_{Code}}$ consisting of all neural codes, which was first introduced in [14]. We show that the bound on open embedding dimension from Theorem 1.11 can be proven combinatorially using this partial order. We also generalize Definition 1.10, and apply Theorem 1.13 to prove a generalization of Theorem 1.11, viewing these results through the lens of $\mathbf{P_{Code}}$.

2. Background and Preliminaries

Throughout this paper we will assume familiarity with standard concepts in topology and convex geometry; for example the interior, closure, and boundary of a set in $\mathbb{R}^d$, convex hulls, hyperplanes, and halfspaces (see [20, Chapter 1]). Recall that each hyperplane H in $\mathbb{R}^d$ may be given an orientation, so that we can speak of the (open) halfspaces $H^>$ and $H^<$ consisting of points lying on the positive and negative sides of H, respectively. We will also use $H^\geq$ and $H^\leq$ to denote the (closed) non-negative and non-positive respective halfspaces associated to H. For any convex set $U \subseteq \mathbb{R}^d$ and any boundary point p of U, one can find a *supporting hyperplane* through p: an oriented hyperplane H containing p with $U \subseteq H^\geq$.

Below, we provide additional background on convex codes, simplicial complexes, and polytopes.

2.1. Convex Codes.

In Section 1 we gave a brief overview of the theory of convex neural codes. We will need one additional concept related to neural codes, described below.

Definition 2.1. Let $\mathcal{C} \subseteq 2^{[n]}$ be a code, and let $\sigma \subseteq [n]$. The *trunk* of σ in $\mathcal{C}$ is

$$\text{Tk}_\mathcal{C}(\sigma) := \{c \in \mathcal{C} \mid \sigma \subseteq c\}.$$

A subset of $\mathcal{C}$ is called a trunk if it is empty, or equal to $\mathrm{Tk}_{\mathcal{C}}(\sigma)$ for some $\sigma \subseteq [n]$. When $\sigma = \{i\}$ we will call $\mathrm{Tk}_{\mathcal{C}}(\sigma)$ a *simple* trunk, and denote it $\mathrm{Tk}_{\mathcal{C}}(i)$.

We introduced trunks in [14] and used them (and a consequent notion of morphism) to define a convenient partial order on neural codes, in which convex codes form a down-set. We will make use of this partial order to contextualize our results in Section 9.

It is worth briefly justifying our requirement that realizations consist of all closed or all open sets. As mentioned in Section 1, openness is a natural requirement from the perspective of neuroscience, in which receptive fields are full-dimensional and do not terminate in sharp boundaries. From a mathematical perspective, requiring closed or open sets is also natural, so that we may think of the receptive fields as a collection of closed or open sets covering some topological subspace of $\mathbb{R}^d$. A further reason to place topological constraints on the sets in our realizations is the following: in [8], it was shown that every code has a realization consisting of convex sets (possibly neither open nor closed). Thus topological constraints are imperative to make the overall question of classifying convex codes meaningful.

2.2. **Simplicial Complexes.** For our purposes, an *(abstract) simplicial complex* is just a code that is closed under taking subsets (i.e. a subset of a codeword is again a codeword). If $\Delta \subseteq 2^{[n]}$ is a simplicial complex, the maximal codewords may be called *facets*, the codewords called *faces*, and elements of $[n]$ called *vertices*. Observe that every simplicial complex is uniquely specified by its facets together with the vertex set $[n]$. In contrast to the usual theory of simplicial complexes, we allow the case in which i is a vertex but $\{i\} \notin \mathcal{C}$.

The *dimension* of a simplicial complex Δ, denoted by $\dim(\Delta)$, is one less than the size of the largest face in Δ. If $\Delta \subseteq 2^{[n]}$ is a simplicial complex, and $m > n$, the *cone over* Δ *with apex* m is the simplicial complex

$$\Delta * m := \{\sigma \subseteq [m] \mid \sigma \setminus \{m\} \in \Delta\}.$$

That is, $\Delta * m$ is the simpicial complex whose facets are the facets of Δ with m added to them. Finally, for any code $\mathcal{C} \subseteq 2^{[n]}$, the *simplicial complex of* $\mathcal{C}$, denoted by $\Delta(\mathcal{C})$, is the smallest simplicial complex containing $\mathcal{C}$.

2.3. **Polytopes and Polytopal Complexes.** A *polytope* is the convex hull of a finite set of points in $\mathbb{R}^d$, or equivalently a bounded intersection of finitely many closed halfspaces. The *dimension* of a polytope is the dimension of its affine hull. The *(proper) faces* of a d-dimensional polytope in $\mathbb{R}^d$ are its intersections with supporting hyperplanes; faces consisting of a single point are called *vertices*, and maximal faces are called *facets*. We will also consider the empty set to be a proper face of any polytope P, and its associated supporting hyperplane to be any hyperplane that does not intersect P.

One can partially order the faces of a polytope by inclusion to form its *face poset*. Two polytopes are called *combinatorially equivalent* if their face posets are isomorphic. Every polytope $P \subseteq \mathbb{R}^d$ admits a *dual polytope* $P^* \subseteq \mathbb{R}^d$, which has the property that the face poset of P^* is isomorphic to the dual of the face poset of P (i.e. one obtains the face poset of P^* by turning the face poset of P upside down).

A polytope is called *d-neighborly* if the convex hull of any d of its vertices is a face. Conveniently, d-neighborly polytopes with an arbitrarily large number of vertices can always be found in $\mathbb{R}^{2d}$ (e.g. the cyclic polytope; see [26, Corollary 0.8]).

A *polytopal complex* in $\mathbb{R}^d$ is a finite set of polytopes $\mathcal{P}$ with the properties that (i) if $P \in \mathcal{P}$, then any face of P is also in $\mathcal{P}$, and (ii) the intersection of two polytopes $P_1, P_2 \in \mathcal{P}$ is a face of both P_1 and P_2. Polytopes in $\mathcal{P}$ are called *faces*.

Each polytopal complex $\mathcal{P}$ has a *face poset*, consisting of all faces in $\mathcal{P}$ partially ordered by containment. Two polytopal complexes are called *combinatorially equivalent* if their face posets are isomorphic. Maximal faces in $\mathcal{P}$ are called *facets*, and if all facets have the same dimension then $\mathcal{P}$ is called *pure*. Finally, we say that a polytopal complex $\mathcal{P}$ in $\mathbb{R}^d$ is *full-dimensional* if it has a facet of dimension d.

Given a d-dimensional polytope $P \subseteq \mathbb{R}^d$ and a facet F of P, one can form a pure, full-dimensional polytopal complex in $\mathbb{R}^{d-1}$ called the *Schlegel diagram of P based at F*. Roughly, one does this by "looking through" the facet F to project all other faces of P into $\mathbb{R}^{d-1}$. The key fact about Schlegel diagrams that we will need is the following: as a polytopal complex, the Schlegel diagram is combinatorially equivalent to the complex of all proper faces of P, but with F removed. For further background on polytopes and polyhedral complexes, we refer the reader to [26, Chapter 5].

3. Closed Embedding Dimension is Bounded by Open Embedding Dimension

To begin our investigation, we recall a useful characterization of intersection complete codes in terms of their realizations. This fact has been observed before in various forms, for example [5, Theorem 1.9].

Proposition 3.1. *A code $\mathcal{C} \subseteq 2^{[n]}$ is intersection complete if and only if the following holds: for all $\sigma \in \Delta(\mathcal{C}) \setminus \mathcal{C}$ and all (possibly non-convex) realizations $\mathcal{U} = \{U_1, \dots U_n\}$ of $\mathcal{C}$ there is some $i \in [n] \setminus \sigma$ with $U_\sigma \subseteq U_i$.*

Proof. First suppose that $\mathcal{C}$ is intersection complete, and has a realization $\mathcal{U} = \{U_1, \dots, U_n\}$. Let $\sigma \in \Delta(\mathcal{C}) \setminus \mathcal{C}$ and define $c_0 = \bigcap_{c \in \mathrm{Tk}_\mathcal{C}(\sigma)} c$. The trunk $\mathrm{Tk}_\mathcal{C}(\sigma)$ is nonempty since $\sigma \in \Delta(\mathcal{C})$, and $c_0 \in \mathcal{C}$ since $\mathcal{C}$ is intersection complete. By construction, σ is a subset of c_0. In fact, σ is a proper subset of c_0 since $\sigma \notin \mathcal{C}$. Thus we may choose $i \in c_0 \setminus \sigma$.

We claim that $U_\sigma \subseteq U_i$. Indeed, since c_0 is the unique minimal element of $\mathrm{Tk}_\mathcal{C}(\sigma)$, every codeword containing σ also contains i. This implies that $U_\sigma \subseteq U_i$.

For the converse, we prove the contrapositive. Suppose that $\mathcal{C}$ is not intersection complete, so there exist c_1 and c_2 in $\mathcal{C}$ such that $c_1 \cap c_2 \notin \mathcal{C}$. Define $\sigma = c_1 \cap c_2$ and note that $\sigma \in \Delta(\mathcal{C}) \setminus \mathcal{C}$. Then choose any (possibly non-convex) realization $\mathcal{U} = \{U_1, \dots, U_n\}$ of $\mathcal{C}$, and let $i \in [n] \setminus \sigma$. Observe that i is contained in at most one of c_1 and c_2. Since U_σ contains U_{c_1} and U_{c_2}, it follows that there is a point in U_σ that is not contained in U_i. This proves the result. □

In addition to Proposition 3.1, we will need the following "trimming" operation, which was also employed in [17].

Definition 3.2. Let $U \subseteq \mathbb{R}^d$ be any set and $\varepsilon > 0$. The *trim* of U by ε is the set

$$\mathrm{trim}(U, \varepsilon) := \{p \in U \mid B_\varepsilon(p) \subseteq U\},$$

where $B_\varepsilon(p)$ is the closed ball of radius ε centered at p.

Proposition 3.3. *If $U \subseteq \mathbb{R}^d$ is convex and open, then $\mathrm{trim}(U, \varepsilon)$ is convex and open for any $\varepsilon > 0$. Moreover, $\mathrm{cl}(\mathrm{trim}(U, \varepsilon)) \subseteq U$.*

Proof. Let $p \in \text{trim}(U, \varepsilon)$. Since $B_\varepsilon(p)$ is a closed subset of U and U is open and convex, there exists $\delta > 0$ such that $B_{\varepsilon+\delta}(p) \subseteq U$. This implies that the open ball of radius δ centered at p is contained in $\text{trim}(U, \varepsilon)$. Thus p is an interior point of $\text{trim}(U, \varepsilon)$, so $\text{trim}(U, \varepsilon)$ is open.

Next let p and q be points in $\text{trim}(U, \varepsilon)$. By convexity of U, the Minkowski sum $C = \overline{pq} + B_\varepsilon(0)$ is contained in U. For any r on $\overline{pq}$, this implies that $B_\varepsilon(r) \subseteq C \subseteq U$. Thus r lies in $\text{trim}(U, \varepsilon)$, proving that $\text{trim}(U, \varepsilon)$ is convex.

For the final statement, observe that no boundary point of U is a boundary point of $\text{trim}(U, \varepsilon)$. Thus all boundary points of $\text{trim}(U, \varepsilon)$ lie in U, and so the closure $\text{cl}(\text{trim}(U, \varepsilon))$ must be a subset of U. □

Proposition 3.4. *Let U and V be sets in $\mathbb{R}^d$. Then $\text{trim}(U \cap V, \varepsilon) = \text{trim}(U, \varepsilon) \cap \text{trim}(V, \varepsilon)$. If $U \subseteq V$, then $\text{trim}(U, \varepsilon) \subseteq \text{trim}(V, \varepsilon)$.*

Proof. The first statement follows from the fact that $B_\varepsilon(p)$ is contained in both U and V if and only if it is contained in their intersection. The second statement is immediate from Definition 3.2. □

A notion of non-degeneracy for realizations was introduced in [3]. Intuitively, non-degeneracy requires that the different regions in the realization do not get too close to one another, unless they intersect. The formal definition is given below.

Definition 3.5 ([3])**.** A collection $\mathcal{U} = \{U_1, \ldots, U_n\}$ of convex sets in $\mathbb{R}^d$ is called *non-degenerate* if the following two conditions hold:

(i) For all $\sigma \in \text{code}(\mathcal{U})$, the atom $\mathcal{A}^\sigma_\mathcal{U}$ is top dimensional (i.e. its intersection with any open set is either empty, or has nonempty interior).
(ii) For all nonempty $\sigma \subseteq [n]$, we have $\bigcap_{i\in\sigma} \partial U_i \subseteq \partial U_\sigma$.

When $\mathcal{U}$ is a collection of convex open sets, [3] proved that (ii) implies (i). We will show that trimming a convex open realization of an intersection complete code $\mathcal{C}$ by a sufficiently small ε yields a non-degenerate realization of $\mathcal{C}$.

Lemma 3.6. *Let $\mathcal{C} \subseteq 2^{[n]}$ be an intersection complete code, and let $\mathcal{U} = \{U_1, \ldots, U_n\}$ be a convex open realization of $\mathcal{C}$. Then there exists $\varepsilon > 0$ such that the sets $V_i = \text{trim}(U_i, \varepsilon)$ form a non-degenerate convex open realization of $\mathcal{C}$.*

Proof. For each codeword $c \in \mathcal{C}$, choose a point $p_c \in \mathcal{A}^c_\mathcal{U}$. Observe that we may choose ε small enough that $B_\varepsilon(p_c) \subseteq U_c$ for all $c \in \mathcal{C}$. We claim that this suffices. Note that by choice of ε, $p_c \in V_c$ for all $c \in \mathcal{C}$. In particular, if U_σ is nonempty then so is V_σ.

To prove that $\mathcal{C} = \text{code}(\{V_1, \ldots, V_n\})$, we must show for all nonempty $\sigma \subseteq [n]$ that U_σ is covered by $\{U_i \mid i \in [n] \setminus \sigma\}$ if and only if V_σ is covered by $\{V_i \mid i \in [n] \setminus \sigma\}$. Suppose first that U_σ is covered by $\{U_i \mid i \in [n] \setminus \sigma\}$. Since $\mathcal{C}$ is intersection complete, Proposition 3.1 implies that there exists some $i \in [n] \setminus \sigma$ with $U_\sigma \subseteq U_i$. But by Proposition 3.4 trimming commutes with intersections and preserves containment, and so $V_\sigma \subseteq V_i$ as desired.

For the converse, we prove the contrapositive. Suppose that U_σ is not covered by $\{U_i \mid i \in [n] \setminus \sigma\}$. Then $\sigma \in \mathcal{C}$ and we may consider the point p_σ. By choice of ε, p_σ is in V_σ but not any U_i with $i \in [n] \setminus \sigma$. Since $V_i \subseteq U_i$, this implies that p_σ is not covered by $\{V_i \mid i \in [n] \setminus \sigma\}$. This proves that $\mathcal{C} = \text{code}(\{V_1, \ldots, V_n\})$.

To see that the V_i form a non-degenerate realization, we must check (ii) of Definition 3.5. For any nonempty $\sigma \subseteq [n]$, let p be a point in $\bigcap_{i\in\sigma} \partial V_i$. Observe that since the closure of any V_i is contained in U_i, the point p lies in U_σ. Hence U_σ is nonempty. We may choose a

point $q \in V_\sigma$, and consider the line segment $\overline{pq}$. Since p is a boundary point of all V_i with $i \in \sigma$, the line segment $\overline{pq}$ is contained in V_i except for the point p. But this implies that all points on the line segment except p lie in V_σ. Thus p is a boundary point of V_σ and the result follows. □

Remark 3.7. Note that we only used intersection completeness of $\mathcal{C}$ once in the proof of Lemma 3.6, namely in the second paragraph so that we could apply Proposition 3.1. The third paragraph of the proof shows that, independent of intersection completeness, the codewords of $\mathcal{C}$ all appear as codewords in the trimmed realization. Likewise, the fourth paragraph proves non-degeneracy of the trimmed realization without using intersection completeness. Thus if $\mathcal{C}$ is *any* code with a convex open realization $\mathcal{U}$, then we may trim $\mathcal{U}$ by some small $\varepsilon > 0$ to obtain a non-degenerate realization of a code $\mathcal{D}$ with $\mathcal{C} \subseteq \mathcal{D} \subseteq \Delta(\mathcal{C})$.

Example 3.8. Below we show the construction used in the proof of Lemma 3.6 (and in turn in the proof of Theorem 1.8) for two realizations of intersection complete codes.

The first is the code $\{\mathbf{123}, 1, 2, 3, \emptyset\}$. In the figure below the U_i already formed a non-degenerate open realization, but trimming them slightly does not change the realized code.

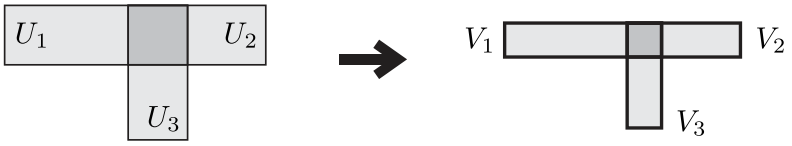

The figure below shows a degenerate open realization of the code $\{\mathbf{13}, \mathbf{14}, \mathbf{23}, \mathbf{24}, 1, 2, 3, 4, \emptyset\}$. This realization is degenerate since U_1 and U_2 are disjoint but share boundary points, and similarly for U_3 and U_4. On the lefthand side, we have labeled the regions corresponding to maximal codewords. On the right we have labeled each of the sets in the trimmed realization.

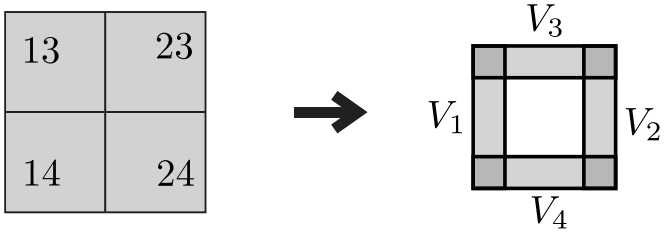

The importance of non-degeneracy is the following: when $\mathcal{U}$ is a non-degenerate collection of convex open sets, taking the closures of these sets does not change the code of the collection (see [3, Theorem 2.12]). With this, we are ready to prove Theorem 1.8.

Theorem 1.8. *Let $\mathcal{C} \subseteq 2^{[n]}$ be an intersection complete code. Then* $\operatorname{cdim}(\mathcal{C}) \leq \operatorname{odim}(\mathcal{C})$.

Proof. Let $\mathcal{U} = \{U_1, \ldots, U_n\}$ be an open realization of $\mathcal{C}$ in $\mathbb{R}^{\operatorname{odim}(\mathcal{C})}$. By Lemma 3.6, we may assume that $\mathcal{U}$ is non-degenerate by possibly trimming the sets in the realization. By [3, Theorem 2.12], the realization consisting of closures of the U_i is a closed convex realization of $\mathcal{C}$. Thus $\operatorname{cdim}(\mathcal{C}) \leq \operatorname{odim}(\mathcal{C})$. □

Example 3.9. Trimming an open realization may fail when a code is not intersection complete. The following shows an open convex realization of the code $\{\mathbf{123}, 12, 13, \emptyset\}$ with labeled atoms, and a trimming of that realization. One can observe that no matter how small we choose ε, trimming this realization always yields an arrangement in which part of V_1 is not

covered by V_2 and V_3.

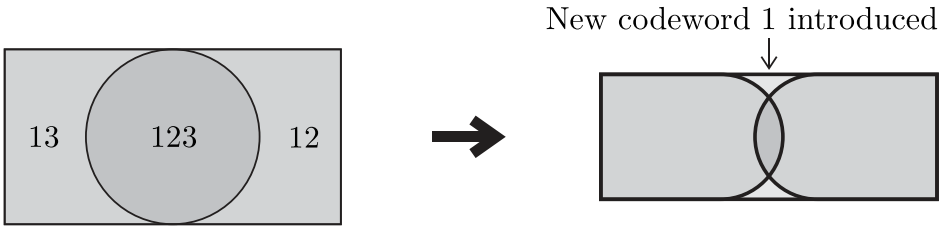

Of course, we could have drawn a better realization of this code. However, we are not always so lucky. There are examples of open convex codes where trimming will fail for *any* convex open realization—see for example [3, Section 2.3], which describes a code for which every convex open realization is degenerate.

4. Closed Embedding Dimension is Bounded by $\min\{2d+1, n-1\}$

Throughout this section, let us fix a (possibly not intersection complete) code $\mathcal{C} \subseteq 2^{[n]}$, and let $d = \dim(\Delta(\mathcal{C}))$. Let us also assume that $n \geq 2$. This avoids the possibility that $\mathcal{C} = \{\mathbf{1}, \emptyset\}$ (this code is intersection complete, but does not satisfy Theorem 1.9 because its closed embedding dimension is 1, while $\min\{2d+1, n-1\} = 0$).

We will attempt to build a realization of $\mathcal{C}$ using closed convex sets satisfying the bound of Theorem 1.9. As we will prove in Lemma 4.6, this construction will succeed if and only if $\mathcal{C}$ is intersection complete. This result echoes [3, Lemma 5.9], but our approach allows us stronger control over the dimension of the ambient space. Our approach is inspired by the construction described in [25, Theorem 3.1].

Throughout this section we will refer to the *intersection completion* of $\mathcal{C}$, which is the code containing all intersections of codewords in $\mathcal{C}$. Note that $\mathcal{C}$ is intersection complete if and only if it is equal to its intersection completion. To begin building our attempted realization, we need to introduce several combinatorial objects.

Lemma 4.1. *Let $m = \min\{2d+1, n-1\}$. There exists a pure, full-dimensional polytopal complex $\mathcal{P}$ in $\mathbb{R}^m$ with n facets $\{P_1, \ldots, P_n\}$ and the following property: each choice of $d+1$ distinct facets of $\mathcal{P}$ determines a unique nonempty face of $\mathcal{P}$ by taking the intersection of the chosen facets. In particular,* $\operatorname{code}(\{P_1, \ldots, P_n\})$ *contains all $\sigma \subseteq [n]$ with $|\sigma| \leq d+1$.*

Proof. First, recall that there exists a full-dimensional $(d+1)$-neighborly polytope in $\mathbb{R}^{m+1}$ with $n+1$ vertices. When $m = 2d+1$, one example is the cyclic polytope, and when $m = n-1$ the n-simplex suffices. Let $P \subseteq \mathbb{R}^{m+1}$ be a polytope dual to a $(d+1)$-neighborly polytope with $n+1$ vertices. Let $F_1, \ldots, F_n, F_{n+1}$ be the facets of P, and observe by neighborliness of the dual that any $d+1$ facets of P meet in a unique nonempty face of P (nonemptiness follows from the fact that in the neighborly polytope dual to P, any $d+1$ or fewer vertices are the vertices of a face which is not the whole polytope). Consider the Schlegel diagram of P in $\mathbb{R}^m$ based at the facet F_{n+1}. For $1 \leq i \leq n$, define P_i to be the image of F_i in the Schlegel diagram. We claim that the complex $\mathcal{P}$ with facets $\{P_1, \ldots, P_n\}$ is the desired polytopal complex.

Each P_i is full-dimensional since each F_i has dimension m. Hence $\mathcal{P}$ is pure and full-dimensional. Furthermore, for any nonempty $\sigma \subseteq [n]$ with $|\sigma| \leq d+1$, the facets P_i with $i \in \sigma$ meet at a unique nonempty face of $\mathcal{P}$ (since the Schlegel diagram preserves intersections and nonemptiness). A point in the relative interior of this face will not lie in any P_j with $j \notin \sigma$, and so $\sigma \in \operatorname{code}(\mathcal{P})$. This proves the result. □

For the remainder of this section, let us fix a polytopal complex $\mathcal{P}$ with facets $\{P_1, \dots, P_n\}$ as given by Lemma 4.1. So far we have a fixed code $\mathcal{C}$, and a fixed complex $\mathcal{P}$. We begin to relate these two objects to each other below. For $\sigma \subseteq [n]$, recall that P_σ denotes $\bigcap_{i\in\sigma} P_i$.

- For each nonempty $\sigma \subseteq [n]$ with $|\sigma| \leq d+1$, fix a point p_σ in the relative interior of P_σ, and
- For each $i \in [n]$, define $V_i := \operatorname{conv}\{p_c \mid c \in \operatorname{Tk}_{\mathcal{C}}(i)\}$.

These objects are illustrated in Example 4.7 below. For now, we observe two facts that we will make use of in the lemmas below:

(1) Since the various P_σ with $|\sigma| \leq d+1$ are distinct faces of $\mathcal{P}$, $p_\sigma \in P_i$ if and only if $i \in \sigma$, and
(2) For all $i \in [n]$, $V_i \subseteq P_i$, and as a consequence $V_\tau \subseteq P_\tau$ for all nonempty $\tau \subseteq [n]$.

The following lemmas build the connection between the sets V_i and the structure of our fixed code $\mathcal{C}$.

Lemma 4.2. *Let $\sigma \subseteq [n]$ with $|\sigma| \geq 2$, and let $i \in \sigma$. Let H be a supporting hyperplane for the proper face P_σ of P_i. Then $V_i \cap H = \operatorname{conv}\{p_c \mid c \in \operatorname{Tk}_{\mathcal{C}}(\sigma)\}$.*

Proof. Consider the points $\{p_c \mid c \in \operatorname{Tk}_{\mathcal{C}}(i)\}$, the convex hull of which is equal to V_i by definition. Since $V_i \subseteq P_i$, we see that $V_i \subseteq H^{\geq}$. Thus $V_i \cap H$ is the convex hull of all points in $\{p_c \mid c \in \operatorname{Tk}_{\mathcal{C}}(i)\}$ which lie in H. If $c \in \operatorname{Tk}_{\mathcal{C}}(i)$ but $\sigma \not\subseteq c$, then we may choose $j \in \sigma \setminus c$, noting that $p_c \notin P_j$. In particular, $p_c \in P_i$ but $p_c \notin P_\sigma$. Thus p_c lies in $H^{>}$ when $\sigma \not\subseteq c$. On the other hand, if $\sigma \subseteq c$ then $p_c \in P_\sigma \subseteq H$. Thus $V_i \cap H$ is the convex hull of $\{p_c \mid c \in \operatorname{Tk}_{\mathcal{C}}(\sigma)\}$ as desired. □

Lemma 4.3. *Let $\sigma \subseteq [n]$ be nonempty. Then $V_\sigma = \operatorname{conv}\{p_c \mid c \in \operatorname{Tk}_{\mathcal{C}}(\sigma)\}$.*

Proof. Let $C = \operatorname{conv}\{p_c \mid c \in \operatorname{Tk}_{\mathcal{C}}(\sigma)\}$. Then $C \subseteq V_\sigma$ since each p_c with $c \in \operatorname{Tk}_{\mathcal{C}}(\sigma)$ lies in V_j for all $j \in \sigma$. For the reverse inclusion, we consider two cases. If $\sigma = \{i\}$ then $C = V_i$ and the result is immediate. Otherwise, $|\sigma| \geq 2$ and we may choose $i \in \sigma$ and H a supporting hyperplane for the face P_σ of P_i. Observe that $V_\sigma \subseteq V_i \cap P_\sigma \subseteq V_i \cap H$, and by Lemma 4.2 $V_i \cap H = C$, proving the result. □

Lemma 4.4. *Let σ and τ be nonempty subsets of $[n]$. Then V_σ is a face of V_τ if and only if $\operatorname{Tk}_{\mathcal{C}}(\sigma) \subseteq \operatorname{Tk}_{\mathcal{C}}(\tau)$. Furthermore, V_σ is a proper face of V_τ if and only if $\operatorname{Tk}_{\mathcal{C}}(\sigma)$ is a proper subset of $\operatorname{Tk}_{\mathcal{C}}(\tau)$.*

Proof. First suppose that $\operatorname{Tk}_{\mathcal{C}}(\sigma) \subseteq \operatorname{Tk}_{\mathcal{C}}(\tau)$. This implies that every codeword that contains σ also contains τ, and so $\operatorname{Tk}_{\mathcal{C}}(\sigma) = \operatorname{Tk}_{\mathcal{C}}(\sigma \cup \tau)$. Lemma 4.3 then implies that $V_\sigma = V_{\sigma\cup\tau}$, and so it suffices to prove that $V_{\sigma\cup\tau}$ is a face of V_τ. We may reduce to the case in which $\tau \subseteq \sigma$, that is, it suffices to prove that V_σ is a face of all V_i with $i \in \tau$. If $\sigma = \{i\}$ then $\tau = \{i\}$ and the result is immediate. Otherwise, $|\sigma| \geq 2$, and for any $i \in \tau$ we may choose a hyperplane H supporting the face P_σ of P_i. Lemma 4.2 implies that $H \cap V_i = \operatorname{conv}\{p_c \mid c \in \operatorname{Tk}_{\mathcal{C}}(\sigma)\}$, and Lemma 4.3 implies that this set is V_σ. Thus $V_i \cap H = V_\sigma$ and V_σ is a face of V_i for all $i \in \tau$ as desired.

To prove that $\operatorname{Tk}_{\mathcal{C}}(\sigma) \subseteq \operatorname{Tk}_{\mathcal{C}}(\tau)$ whenever V_σ is a face of V_τ, we argue by contrapositive. If $\operatorname{Tk}_{\mathcal{C}}(\sigma) \not\subseteq \operatorname{Tk}_{\mathcal{C}}(\tau)$ then there exists $c \in \mathcal{C}$ with $\sigma \subseteq c$ but $\tau \not\subseteq c$. Consider the point p_c. Since $\tau \not\subseteq c$, there exists $i \in \tau \setminus c$, and we see that $p_c \notin P_i$. But $V_\tau \subseteq V_i \subseteq P_i$, so $p_c \notin V_\tau$. On the other hand, $p_c \in V_\sigma$, so V_σ is not contained in V_τ, and thus cannot be a face of V_τ.

We can now prove the final sentence in the lemma. If V_σ is a proper face of V_τ, then $\mathrm{Tk}_{\mathcal{C}}(\sigma)$ is a subset of $\mathrm{Tk}_{\mathcal{C}}(\tau)$. If this containment were not proper, then V_τ would be a face of V_σ, a contradiction. Conversely, if $\mathrm{Tk}_{\mathcal{C}}(\sigma)$ is a proper subset of $\mathrm{Tk}_{\mathcal{C}}(\tau)$ then V_σ is a face of V_τ. If V_σ were not a proper face, then we would have $V_\sigma = V_\tau$, which in turn implies $\mathrm{Tk}_{\mathcal{C}}(\tau) \subseteq \mathrm{Tk}_{\mathcal{C}}(\sigma)$, a contradiction. This proves the result. □

Lemma 4.5. *Let $\sigma \subseteq [n]$ be nonempty. Then σ lies in the intersection completion of $\mathcal{C}$ if and only if the following holds:* $\mathrm{Tk}_{\mathcal{C}}(\sigma)$ *is nonempty and properly contains* $\mathrm{Tk}_{\mathcal{C}}(\sigma \cup \{i\})$ *for all $i \in [n] \setminus \sigma$.*

Proof. If σ is an intersection of codewords in $\mathcal{C}$, then there must be a codeword containing σ, and thus $\mathrm{Tk}_{\mathcal{C}}(\sigma)$ is nonempty. If there exists $i \in [n] \setminus \sigma$ such that $\mathrm{Tk}_{\mathcal{C}}(\sigma) = \mathrm{Tk}_{\mathcal{C}}(\sigma \cup \{i\})$, then every codeword of $\mathcal{C}$ containing σ also contains i. This is a contradiction, since σ is the intersection of all codewords in $\mathcal{C}$ that contain it.

For the converse we consider two cases. If $\sigma = [n]$ and $\mathrm{Tk}_{\mathcal{C}}(\sigma)$ is nonempty then $[n] \in \mathcal{C}$ and the result follows. Otherwise σ is a proper subset of $[n]$. Since $\mathrm{Tk}_{\mathcal{C}}(\sigma)$ is nonempty and properly contains $\mathrm{Tk}_{\mathcal{C}}(\sigma \cup \{i\})$ for all $i \in [n] \setminus \sigma$, for every $i \in [n] \setminus \sigma$ we may choose a codeword c_i with $\sigma \subseteq c_i$ and $i \notin c_i$. The intersection of all such c_i is σ, proving the result. □

Lemma 4.6. *The set $\mathcal{V} = \{V_1, \ldots, V_n\}$ is a closed realization of the intersection completion of $\mathcal{C}$. In particular, $\mathcal{V}$ is a realization of $\mathcal{C}$ if and only if $\mathcal{C}$ is intersection complete.*

Proof. Let $\widehat{\mathcal{C}}$ denote the intersection completion of $\mathcal{C}$. We argue for each nonempty $\sigma \subseteq [n]$ that $\sigma \in \widehat{\mathcal{C}}$ if and only if $\sigma \in \mathrm{code}(\mathcal{V})$. By Lemma 4.5 it suffices to argue that $\sigma \in \mathrm{code}(\mathcal{V})$ if and only if $\mathrm{Tk}_{\mathcal{C}}(\sigma)$ is nonempty and $\mathrm{Tk}_{\mathcal{C}}(\sigma \cup \{i\})$ is a proper subset of $\mathrm{Tk}_{\mathcal{C}}(\sigma)$ for all $i \in [n] \setminus \sigma$. By Lemma 4.3 and Lemma 4.4, this condition is equivalent to the requirement that V_σ is nonempty, and $V_{\sigma \cup \{i\}}$ is a proper face of V_σ for all $i \in [n] \setminus \sigma$. This is in turn equivalent to the statement that V_σ is nonempty and not covered by $\{V_i \mid i \in [n] \setminus \sigma\}$, which happens if and only if $\sigma \in \mathrm{code}(\mathcal{V})$, proving the result. □

Example 4.7. To make the construction in Lemma 4.6 concrete, we give an example for the intersection complete code $\mathcal{C} = \{\mathbf{123}, 12, 1, 2, 3, \emptyset\}$. We choose $\mathcal{P}$ in $\mathbb{R}^2$ with facets P_1, P_2, P_3 which are triangles meeting at a common vertex. This is shown below on the left, and the various p_σ are shown with labeled dots. The righthand side of the figure illustrates the closed convex realization $\mathcal{V} = \{V_1, V_2, V_3\}$ constructed in the proof of Lemma 4.6. The sets V_1 and

V_2 are triangles, and V_3 is the line segment from p_3 to p_{123}.

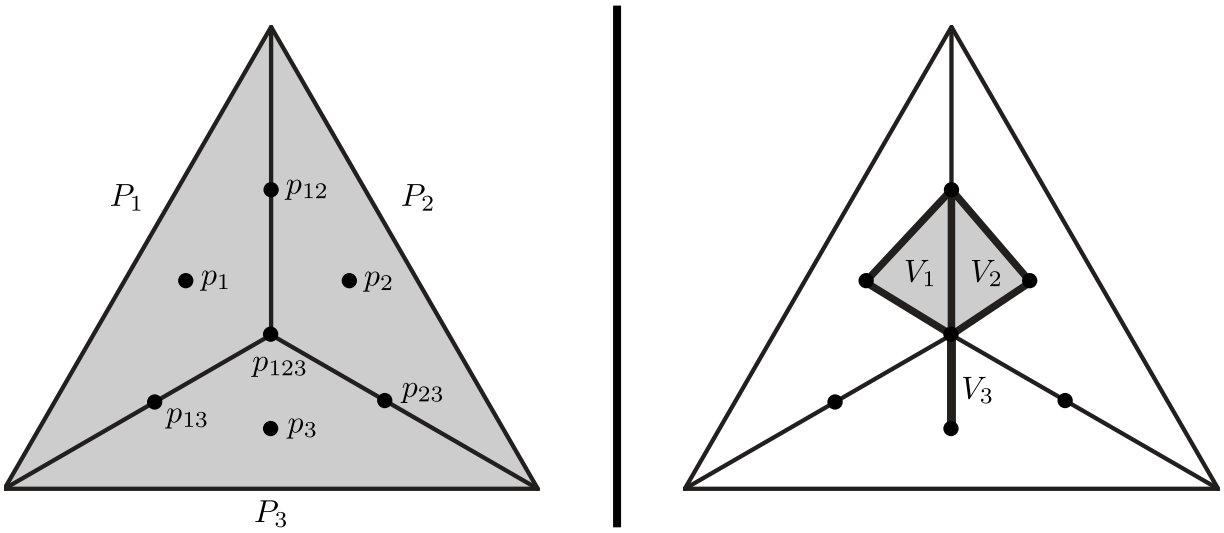

Theorem 1.9. *Let $\mathcal{C} \subseteq 2^{[n]}$ be an intersection complete code with $n \geq 2$, and let $d = \dim(\Delta(\mathcal{C}))$. Then $\operatorname{cdim}(\mathcal{C}) \leq \min\{2d+1, n-1\}$.*

Proof. In this section we have chosen a polytopal complex $\mathcal{P}$ in $\mathbb{R}^{\min\{2d+1,n-1\}}$, and used it to construct a collection $\mathcal{V} = \{V_1, \ldots, V_n\}$ of closed convex sets. Lemma 4.6 says that $\mathcal{V}$ realizes $\mathcal{C}$ if and only if $\mathcal{C}$ is intersection complete. This proves the result. □

In Section 6, we will see that this bound on closed embedding dimension may fail dramatically for open embedding dimension. Before proving this, we use Section 5 to recall a theorem from [13], and show that it is equivalent to a statement about the open embedding dimension of a family of intersection complete codes.

5. A Code Version of the Sunflower Theorem

In this section we recall a result regarding sunflowers of convex open sets. In Section 6, we will use this result to build a family of intersection complete codes with large open embedding dimension.

Definition 5.1. Let $\mathcal{U} = \{U_1, \ldots, U_n\}$ be a collection of convex sets in $\mathbb{R}^d$ and let $\mathcal{C} = \operatorname{code}(\mathcal{U})$. The collection $\mathcal{U}$ is called a *sunflower* if $[n] \in \mathcal{C}$, and $\mathcal{C} \setminus \{[n]\}$ contains only codewords of weight at most 1. That is, a sunflower is just a 1-flexible sunflower. As in Definition 1.12, we will call the U_i *petals* and $U_{[n]}$ will be called the *center* of $\mathcal{U}$.

Theorem 5.2 (Sunflower Theorem, [13])**.** *Let $d \geq 1$, let $\mathcal{U} = \{U_1, \ldots, U_{d+1}\}$ be a convex open sunflower in $\mathbb{R}^d$, and for each $i \in [d+1]$ choose a point $p_i \in U_i$. Then $\operatorname{conv}\{p_1, \ldots, p_{d+1}\}$ contains a point in the center of $\mathcal{U}$.*

Note that the result above fails when we consider a sunflower with d petals in $\mathbb{R}^d$. In particular, one may take an infinite rectangular prism about each coordinate axis to form a sunflower whose center is a hypercube at the origin. In this situation, choosing the p_i to be sufficiently large positive multiples of the coordinate basis vectors yields points in each petal whose convex hull does not touch the center of the sunflower.

The sunflower theorem may be restated purely in the language of convex codes. We do this below in order to simplify our discussion in the following section, and also to foreshadow our applications of Theorem 1.13 in Section 9.3.

Definition 5.3. For $n \geq 1$, define $\mathcal{S}_n \subseteq 2^{[n+1]}$ to be the code consisting of the following codewords: $[n]$, all singleton sets, all pairs $\{i, n+1\}$ for $1 \leq i \leq n$, and the empty set.

Example 5.4. Let us look at the first few $\mathcal{S}_n$:

$$\begin{aligned}\mathcal{S}_1 &= \{\mathbf{12}, 1, 2, \emptyset\},\\ \mathcal{S}_2 &= \{\mathbf{12}, \mathbf{13}, \mathbf{23}, 1, 2, 3, \emptyset\},\\ \mathcal{S}_3 &= \{\mathbf{123}, \mathbf{14}, \mathbf{24}, \mathbf{34}, 1, 2, 3, 4, \emptyset\}.\end{aligned}$$

These have convex open realizations in $\mathbb{R}^1, \mathbb{R}^2$, and $\mathbb{R}^3$ respectively, illustrated below. Theorem 5.5 below says that these realizations are minimal in dimension.

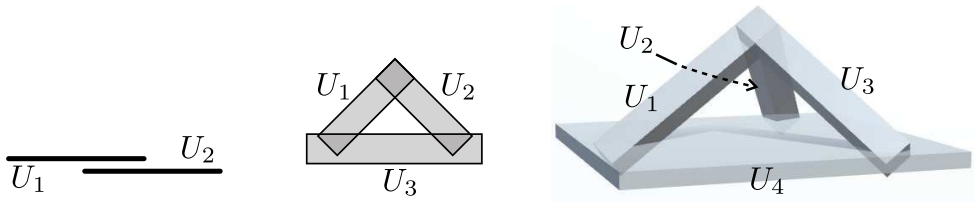

$\mathcal{S}_1$ in $\mathbb{R}^1$ $\mathcal{S}_2$ in $\mathbb{R}^2$ $\mathcal{S}_3$ in $\mathbb{R}^3$

Note that $\mathcal{S}_n$ is an intersection complete code. The sunflower theorem can be restated as follows:

Theorem 5.5 (Sunflower Theorem, Code Version)**.** *For all $n \geq 1$,* $\operatorname{odim}(\mathcal{S}_n) = n$.

Proof. When $n = 1$, we have $\mathcal{S}_n = \{12, 1, 2, \emptyset\}$, which can be realized by two overlapping intervals in $\mathbb{R}^1$. For $n \geq 2$, $\mathcal{S}_n$ has $n+1$ maximal codewords, and so by [3, Theorem 1.2] $\mathcal{S}_n$ has an open realization in $\mathbb{R}^n$. We will show that it does not have an open realization in $\mathbb{R}^{n-1}$. Suppose for contradiction that there exists an open realization $\mathcal{U} = \{U_1, \ldots, U_{n+1}\}$ of $\mathcal{S}_n$ in $\mathbb{R}^{n-1}$. Observe that $\{U_1, \ldots, U_n\}$ is a sunflower, and U_{n+1} intersects U_i for all $i \in [n]$. Thus for each $i \in [n]$ we may choose $p_i \in U_i \cap U_{n+1}$. The convex hull $\operatorname{conv}\{p_1, \ldots, p_n\}$ is contained in U_{n+1}, but by Theorem 5.2 this convex hull also meets $U_{[n]}$. Thus $U_{n+1} \cap U_{[n]}$ is nonempty. Since $[n+1]$ is not a codeword in $\mathcal{S}_n$ this is a contradiction. □

In the following section we build on the family $\mathcal{S}_n$ to construct a family of intersection complete codes on n neurons whose open embedding dimension is $\binom{n-1}{\lfloor (n-1)/2 \rfloor}$.

6. A Family of Codes with Large Open Embedding Dimension

In this section, we will associate to every simplicial complex $\Delta \subseteq 2^{[n]}$ an intersection complete code $\mathcal{S}_\Delta \subseteq 2^{[n+1]}$. As long as Δ has at least two facets, the open embedding dimension of $\mathcal{S}_\Delta$ is exactly the number of facets in Δ. To start, recall the following definition.

Definition 1.10. Let $\Delta \subseteq 2^{[n]}$ be a simplicial complex. Define $\mathcal{S}_\Delta \subseteq 2^{[n+1]}$ to be the code

$$\mathcal{S}_\Delta := (\Delta * (n+1)) \cup \{[n]\},$$

where $\Delta * (n+1)$ denotes the cone over Δ with apex $n+1$.

We start with some straightforward structural observations about the code $\mathcal{S}_\Delta$.

Proposition 6.1. *The code $\mathcal{S}_\Delta$ is intersection complete. If $\Delta \subsetneq 2^{[n]}$ and has m facets, then $\mathcal{S}_\Delta$ has $m+1$ maximal codewords. In particular,* $\operatorname{odim}(\mathcal{S}_\Delta) \leq \max\{2, m\}$.

Proof. First note that $\mathcal{S}_\Delta$ is a simplicial complex, plus the codeword $[n]$. Adding a single codeword to a simplicial complex always yields an intersection complete code, so $\mathcal{S}_\Delta$ is intersection complete.

Let $F_1, \ldots, F_m$ be the facets of Δ. Observe that each maximal codeword of $\mathcal{S}_\Delta$ is either a facet of $\Delta * (n+1)$, or equal to $[n]$. The facets of $\Delta * (n+1)$ are just $F_i \cup \{n+1\}$ for $i \in [m]$. Since $\Delta \subsetneq 2^{[n]}$, $[n]$ is an additional a maximal codeword of $\mathcal{S}_\Delta$, so $\mathcal{S}_\Delta$ has $m+1$ maximal codewords in total. The bound $\operatorname{odim}(\mathcal{S}_\Delta) \leq \max\{2, m\}$ then follows immediately from [3, Theorem 1.2]. □

Theorem 1.11. *Let $\Delta \subseteq 2^{[n]}$ be a simplicial complex with $m \geq 2$ facets. Then $\mathcal{S}_\Delta$ (as given by Definition 1.10) is an intersection complete code with $m+1$ maximal codewords, and* $\operatorname{odim}(\mathcal{S}_\Delta) = m$.

Proof. By Proposition 6.1 we know that $\mathcal{S}_\Delta$ is intersection complete, has $m+1$ maximal codewords, and satisfies $\operatorname{odim}(\mathcal{S}_\Delta) \leq m$. Thus it suffices to show that $\mathcal{S}_\Delta$ does not have an open realization in $\mathbb{R}^{m-1}$. Suppose for contradiction that we had such a realization $\mathcal{U} = \{U_1, \ldots, U_{n+1}\}$.

Label the facets of Δ as $F_1, \ldots, F_m$, and for each $i \in [m]$ define $V_i = U_{F_i}$. Lastly, define $V_{m+1} = U_{n+1}$. Now observe that the pairwise intersection of any two distinct V_i with $i \in [m]$ is $U_{[n]}$ since $[n]$ is the only codeword of $\mathcal{S}_\Delta$ that properly contains more than one F_i. Thus $\{V_1, \ldots, V_m\}$ is a sunflower. Note that V_{m+1} intersects each petal of this sunflower since $F_i \cup \{n+1\}$ is a codeword of $\mathcal{S}_\Delta$ for all $i \in [m]$. However, V_{m+1} does not intersect $U_{[n]}$ since $[n]$ is not a face of Δ (recall $m \geq 2$).

We have constructed an open sunflower $\{V_1, \ldots, V_m\}$ in $\mathbb{R}^{m-1}$ and a convex open set V_{m+1} which intersects each petal of the sunflower but not its center. Equivalently, $\{V_1, \ldots, V_{m+1}\}$ is a convex open realization of $\mathcal{S}_m$ in $\mathbb{R}^{m-1}$. This contradicts Theorem 5.5, and so $\operatorname{odim}(\mathcal{S}_\Delta)$ must be equal to m as desired. □

Corollary 6.2. *For any $n \geq 2$ and $1 \leq m \leq \binom{n-1}{\lfloor (n-1)/2 \rfloor}$, there exists an intersection complete code on n neurons with $m+1$ maximal codewords, and open embedding dimension equal to m.*

Proof. For $m = 1$, the code $\{\mathbf{1}, \emptyset\}$ suffices. For $m \geq 2$ we apply Theorem 1.11. Among all $\binom{n-1}{\lfloor (n-1)/2 \rfloor}$ subsets of $[n-1]$ with size $\lfloor (n-1)/2 \rfloor$, we may select m. Letting Δ be the simplicial complex with these subsets as its facets, we see that $\mathcal{S}_\Delta$ is the desired code. □

Corollary 6.3. *There is a family of codes $\mathcal{E}_n \subseteq 2^{[n]}$ such that* $\operatorname{odim}(\mathcal{E}_n)$ *grows at least as fast as $\frac{2^{n-1}}{n}$.*

Proof. By Corollary 6.2, we may choose $\mathcal{E}_n$ so that $\operatorname{odim}(\mathcal{E}_n) = \binom{n-1}{\lfloor (n-1)/2 \rfloor}$. Observe that $n\binom{n-1}{\lfloor (n-1)/2 \rfloor} \geq 2^{n-1}$ since there are no more than $\binom{n-1}{\lfloor (n-1)/2 \rfloor}$ subsets of $[n-1]$ of size i for $i = 0, 1, \ldots, n-1$. Thus $\binom{n-1}{\lfloor (n-1)/2 \rfloor} \geq \frac{2^{n-1}}{n}$, proving the result. □

Qualitatively, these results are very surprising. The codes $\mathcal{S}_\Delta$ are "almost" simplicial complexes (we have added the single codeword $[n]$ to a simplicial complex), but their open

embedding dimensions grow exponentially faster than that of any simplicial complex. Strikingly, these codes provide the first example of codes whose embedding dimension (open or closed) is larger than $n-1$.

Remark 6.4. From the perspective of the neuroscience which motivates the study of convex codes, Corollary 6.3 has the following interpretation: theoretically, n neurons may "recognize" dimensions that are exponentially large in n. Whether such a phenomenon ever occurs in experimental data could be an interesting avenue of investigation.

7. Flexible Sunflowers

In this section our goal is to investigate k-flexible sunflowers of convex open sets. These are a generalization of sunflowers in which we allow petals to overlap outside the center of the sunflower, but no more than k at a time. For sunflowers, we saw in Theorem 5.2 that sampling a point in each petal and taking the convex hull always yielded a point in the center of the sunflower if we had enough petals relative to our ambient dimension. We will see that the same holds for k-flexible sunflowers, and the minimum number of petals needed is proportional to k, as well as the ambient dimension (see Theorem 1.13 below). Qualitatively, the more flexibility we allow in a sunflower, the larger the number of petals we need to sample in order to guarantee that the convex hull of the sampled points intersects the center of the sunflower.

Our proofs rely heavily on the assumption that we are working with open sets, and the results of this section do not apply to sunflowers of closed convex sets. Indeed, taking n line segments in the positive quadrant of $\mathbb{R}^2$ which meet at the origin yields a closed sunflower for which the conclusion of Theorem 1.13 does not hold.

To begin, let us recall the definition of a k-flexible sunflower.

Definition 1.12. Let $\mathcal{U} = \{U_1, \ldots, U_n\}$ be a collection of convex sets in $\mathbb{R}^d$ and let $\mathcal{C} = \operatorname{code}(\mathcal{U})$. The collection $\mathcal{U}$ is called a *k-flexible sunflower* if $[n] \in \mathcal{C}$, and all other codewords have weight at most k. The U_i are called *petals* and $U_{[n]}$ is called the *center* of $\mathcal{U}$.

We start with a family of examples. For each $d \geq 2$ and $k \geq 1$, Proposition 7.1 describes a k-flexible sunflower $\mathcal{U}$ in $\mathbb{R}^d$ with dk petals in which we can sample points from each petal whose convex hull does not contain a point in the center of $\mathcal{U}$.

Proposition 7.1. *For all $d \geq 2$ and $k \geq 1$, there exists an open k-flexible sunflower $\mathcal{U} = \{U_1, \ldots, U_n\}$ in $\mathbb{R}^d$ with $n = dk$, and points $p_1, \ldots, p_n$ with $p_i \in U_i$, such that $\operatorname{conv}\{p_1, \ldots, p_n\}$ does not contain a point in the center of $\mathcal{U}$.*

Proof. For $k = 1$, we begin with an open unit hypercube in $\mathbb{R}^d$ centered at the origin, and let U_i be the Minkowski sum of this hypercube with a line segment from the origin to a large positive multiple of e_i. We can see that the U_i form a d-petal sunflower, and our desired p_i are just the large multiples of e_i.

For $k \geq 2$, we can take the sunflower described above and duplicate each of the d petals k times. This creates a k-flexible sunflower, and the same sampling of points (with each duplicated k times) satisfies the proposition. □

Remark 7.2. One might argue that the construction above is unsatisfying. Should we not stipulate that petals diverge in different directions, or at least are distinct? It turns out we can address these concerns. Start with the usual coordinate-direction sunflower whose center

is a unit hypercube, as described above. If $k = 1$ we are done. Otherwise, choose a cyclic permutation σ of $[d]$, for example $i \mapsto i + 1 \bmod d$. Then, we can duplicate each petal in our coordinate-direction sunflower k times, but when duplicating the i-th petal we "skew" it slightly in the direction of $-e_{\sigma(i)}$. If each duplicated petal is skewed a different amount, our petals will diverge from one another. As long as we skew a small enough amount, this yields a k-flexible sunflower from which we can sample the desired p_i.

This construction is illustrated below for $k = 3$ and $d = 2$:

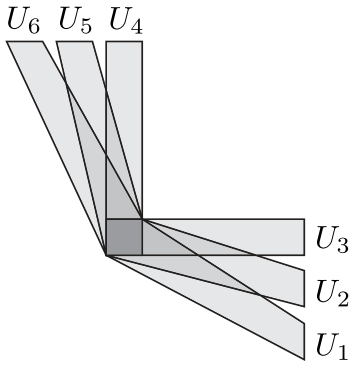

We now turn our attention to proving Theorem 1.13. We will see that some technical lemmas regarding k-flexible sunflowers together with Tverberg's theorem are enough to prove the theorem. We start by showing that the center of every open k-flexible sunflower admits a set of supporting halfspaces with two properties: first, the common intersection of the halfspaces closely approximates the center (in the sense that the common intersection is contained in the closure of the center), and second, any one of these halfspaces contains all but at most k of the petals.

Definition 7.3. Let $\mathcal{U} = \{U_1, \dots, U_n\}$ be a k-flexible sunflower in $\mathbb{R}^d$ with center U. A point $b \in \partial U$ is called *well-supported* if it is not in the boundary of $U_i \cap \partial U$ (considered as a subset of the topological space ∂U) for any $i \in [n]$.

Lemma 7.4. *Let $\mathcal{U} = \{U_1, \dots, U_n\}$ be a k-flexible sunflower in $\mathbb{R}^d$ with center U. The set of well-supported points is dense in ∂U.*

Proof. Consider the sets $U_i \cap \partial U$ in ∂U. For each of these sets, the set of non-boundary points in ∂U is dense and open when considered as a subset of ∂U. The set of well-supported points is just the intersection of non-boundary points of $U_i \cap \partial U$ in $\partial U$for all i, and a finite intersection of dense open sets is again open and dense. Thus the well-supported points are dense in ∂U. □

Lemma 7.5. *Let $U \subseteq \mathbb{R}^d$ be a convex open set. Let B be a dense subset of the boundary of U, and for each $b \in B$ let H_b be a supporting hyperplane to U at b. Then $\bigcap_{b \in B} H_b^>$ is contained in $\overline{U}$.*

Proof. Consider any point $p \notin \overline{U}$. Since U is open and p lies a positive distance away from U, the intersection of $\operatorname{int}(\operatorname{conv}(\{p\} \cup U))$ with ∂U is a relatively open nonempty subset of ∂U, and thus contains some $b \in B$ because B is dense. Since $\operatorname{int}(\operatorname{conv}(\{p\} \cup U))$ is open, the line segment $\overline{pb}$ can be extended so that it ends at a point $q \in U$, as shown in the following

figure.

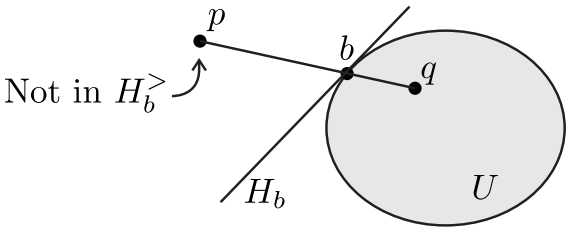

Now, consider the supporting hyperplane H_b. We have $U \subseteq H_b^>$. In particular, $H_b^>$ contains q but not b. Since b lies between q and p, we see that $H_b^>$ does not contain p. Thus $p \notin \bigcap_{b \in B} H_b^>$ and the lemma follows. □

Lemma 7.6. *Let $\mathcal{U} = \{U_1, \ldots, U_n\}$ be an open k-flexible sunflower in $\mathbb{R}^d$ with center U, and let $b \in \partial U$ be well-supported. Let H_b be a supporting hyperplane for U at b. Then $U_i \subseteq H_b^>$ for all $i \in [n]$ such that $b \notin U_i$. In particular, $H_b^>$ contains all but at most k petals of $\mathcal{U}$.*

Proof. Suppose not, so that there exists $i \in [n]$ such that $b \notin U_i$ and U_i is not contained in $H_b^>$. Since U_i is open, we may assume that there exists a point $p \in U_i$ strictly on the negative side of H_b. Then choose any point $q \in U$, and consider the line segment $\overline{qb}$. All points on this line segment other than b lie in U. For each $r \in \overline{qb}$ with $r \neq b$, note that the line segment $\overline{pr}$ is contained in U_i. Moreover, since U is convex and q lies in U while p lies outside of U, this line segment intersects the boundary of U in a unique point. The set of these intersection points forms a subset of $U_i \cap \partial U$ whose closure contains b. This is illustrated in the figure below, with the points in $U_i \cap \partial U$ converging to b shown in the bold curved line segment.

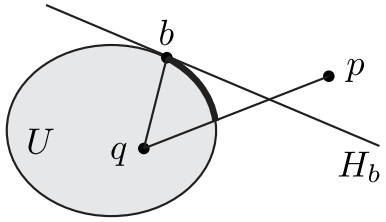

But since $b \notin U_i$, this implies that b is a boundary point of $U_i \cap \partial U$ in ∂U. This contradicts the fact that b is well-supported, and so $U_i \subseteq H_b^>$ whenever $b \notin U_i$. Since $\mathcal{U}$ is k-flexible, b belongs to no more than k petals of $\mathcal{U}$, and thus $H_b^>$ contains all but at most k petals. □

Finally, we recall Tverberg's theorem. After stating this theorem, we are ready to prove Theorem 1.13.

Theorem 7.7 (Tverberg's theorem)**.** *Let $d \geq 1$, $r \geq 2$, and $n = (d+1)(r-1)+1$. For any set of n points $P = \{p_1, \ldots, p_n\}$ in $\mathbb{R}^d$, there is a partition of P into r parts $P_1, \ldots, P_r$ such that $\bigcap_{i=1}^r \operatorname{conv}(P_i) \neq \emptyset$.*

Theorem 1.13. *Let $\mathcal{U} = \{U_1, \ldots, U_n\}$ be an open k-flexible sunflower in $\mathbb{R}^d$. Suppose that $n \geq dk+1$, and for each $i \in [n]$ let $p_i \in U_i$. Then $\operatorname{conv}\{p_1, \ldots, p_n\}$ contains a point in the center of $\mathcal{U}$. Moreover, if $d \geq 2$ this result may fail when $n < dk+1$.*

Proof. It suffices to prove the first statement for $n = dk+1$. Let U denote the center of $\mathcal{U}$. Suppose for contradiction that the theorem does not hold, so that $\operatorname{conv}\{p_1, \ldots, p_n\}$ does not contain a point in U. Since the U_i are open, we may uniformly translate the p_i away from U by a small positive distance, and choose a separating hyperplane H between $\operatorname{conv}\{p_1, \ldots, p_n\}$

and U such that H does not contain any boundary point of U. Moreover, we can replace each p_i by the intersection of the line segment $\overline{p_i p}$ with H, so that all p_i lie inside H.

Now, H has dimension $d-1$, so we may apply Tverberg's theorem to our points p_i with $r = k+1$. We obtain a partition $P_1, \dots, P_{k+1}$ such that $\bigcap_{i=1}^{k+1} \operatorname{conv}(P_i) \neq \emptyset$. Choose any point p lying in this intersection, and observe that $p \in H$.

Let B be the set of well-supported points in ∂U, and choose supporting halfspaces $\{H_b^> \mid b \in B\}$ as per Lemma 7.6. Recall by Lemma 7.6, that for any $b \in B$, the halfspace $H_b^>$ contains all U_j (and thus p_j) except for at most k. In particular, there must be some P_i such that $H_b^>$ contains all points in P_i, and hence also their convex hull. Thus $p \in H_b^>$ for all $b \in B$. But by Lemma 7.4 and Lemma 7.5, this implies that $p \in \overline{U}$. Since $p \in H$ and H was constructed not to contain U or any of its boundary points, this is a contradiction.

To prove the second part of the theorem, recall that Proposition 7.1 shows that when $d \geq 2$ and $n = dk$, we can choose an open k-flexible sunflower $\mathcal{U}$ in $\mathbb{R}^d$ and points in each petal whose convex hull does not intersect the center of $\mathcal{U}$. This proves the result. □

Remark 7.8. Note that when $k = 1$, Theorem 1.13 is the same as Theorem 5.2 (the usual Sunflower Theorem), and the application of Tverberg's theorem in the proof above reduces to an application of Radon's Theorem. Thus the fact that Theorem 1.13 generalizes Theorem 5.2 is directly analogous to the fact that Tverberg's theorem generalizes Radon's.

Remark 7.9. In terms of neuroscientific motivation, flexible sunflowers are natural to investigate. Allowing some codewords beyond singletons, but of a fixed weight, accounts for some tolerance to error in data gathering and also captures a wider range of possibilities. We hope that flexible sunflowers may yield meaningful bounds on dimensions in experimental data.

Theorem 1.13 has implications regarding the open embedding dimensions of intersection complete codes, which we will illustrate in Section 9.3, in particular by generalizing the families $\mathcal{S}_n$ and $\mathcal{S}_\Delta$ that were defined in Sections 5 and 6 respectively.

We conclude with a corollary which examines the extremal case in which we have a k-flexible sunflower $\mathcal{U}$ with $n = dk$ petals for which Theorem 1.13 fails. In this case Theorem 1.13 implies $\operatorname{code}(\mathcal{U})$ must contain at least one codeword of weight k, but we can actually say something slightly stronger:

Corollary 7.10. *Let $\mathcal{U} = \{U_1, \dots, U_n\}$ be an open k-flexible sunflower in $\mathbb{R}^d$. Suppose that $n = dk$, and there exist points $p_1, \dots, p_n$ such that $p_i \in U_i$ and $\operatorname{conv}\{p_1, \dots, p_n\}$ does not contain a point in the center of $\mathcal{U}$. Then $\operatorname{code}(\mathcal{U})$ contains at least d distinct codewords of weight k.*

Proof. We work by induction on k. When $k = 1$ the result is clear since if there are fewer than d codewords of weight k in $\operatorname{code}(\mathcal{U})$ then some U_i is equal to the center of $\mathcal{U}$, and so some p_i lies in the center of $\mathcal{U}$, a contradiction. For $k \geq 2$, suppose for contradiction that $\operatorname{code}(\mathcal{U})$ contains fewer than d codewords of weight k. For each of these codewords c, select some petal U_i with $i \in c$. Deleting these U_i yields a $(k-1)$-flexible sunflower, and since we have deleted fewer than d petals our new $(k-1)$-flexible sunflower has more than $d(k-1)$ petals. But the same choice of p_i yields a collection of points whose convex hull does not contain a point in the center of this $(k-1)$-flexible sunflower, contradicting Theorem 1.13. □

8. Tangled Sunflowers

For $n \geq 1$ we construct a family of intersection complete codes $\mathcal{T}_n \subseteq 2^{[2n]}$, and investigate the open embedding dimensions of these codes. The code $\mathcal{T}_n$ is constructed so that its realizations consist of two sunflowers with n petals that are "tangled" in the sense that the i-th petal of the first sunflower meets the i-th petal of the second, and no other incidences occur. We use the sunflower theorem (Theorem 5.2) to prove the following: for any $d \geq 1$ there exists n such that $\operatorname{odim}(\mathcal{T}_n) = d$. Thus for every $d \geq 1$, one of the $\mathcal{T}_n$ codes describes an arrangement of convex open sets which can be achieved in $\mathbb{R}^d$ but not a smaller dimension. Beyond this statement and some basic bounds, however, determining the exact open embedding dimension of $\mathcal{T}_n$ remains an open problem, ripe for future investigation.

Definition 8.1. Let $n \geq 1$. Define $\mathcal{T}_n \subseteq 2^{[2n]}$ to be the code consisting of the following codewords:

(i) $\{2k-1, 2k\}$ for $k = 1, 2, \ldots, n$,
(ii) $\{1, 3, 5, \ldots, 2n-1\}$ and $\{2, 4, 6, \ldots, 2n\}$,
(iii) all singletons, and
(iv) the empty set.

For each n define $t_n := \operatorname{odim}(\mathcal{T}_n)$.

Observe that codewords of type (i) and (ii) are the maximal codewords in $\mathcal{T}_n$ for $n \geq 2$; in particular $\mathcal{T}_n$ has $n+2$ maximal codewords. Furthermore observe that $\mathcal{T}_n$ is intersection complete, and hence open convex. Thus t_n is finite for all n.

Moreover, note that the odd-numbered sets in any realization of $\mathcal{T}_n$ form an n-petal sunflower, as do the even-numbered sets. These two sunflowers are "tangled," in that their petals are matched and the matched petals overlap each other.

Example 8.2. The first four $\mathcal{T}_n$ are given below:

$$\mathcal{T}_1 = \{\mathbf{12}, 1, 2, \emptyset\},$$
$$\mathcal{T}_2 = \{\mathbf{13}, \mathbf{24}, \mathbf{12}, \mathbf{34}, 1, 2, 3, 4, \emptyset\},$$
$$\mathcal{T}_3 = \{\mathbf{135}, \mathbf{246}, \mathbf{12}, \mathbf{34}, \mathbf{56}, 1, 2, 3, 4, 5, 6, \emptyset\}.$$
$$\mathcal{T}_4 = \{\mathbf{1357}, \mathbf{2468}, \mathbf{12}, \mathbf{34}, \mathbf{56}, \mathbf{78}, 1, 2, 3, 4, 5, 6, 7, 8, \emptyset\}.$$

These have open convex realizations in $\mathbb{R}^1, \mathbb{R}^2$, $\mathbb{R}^3$, and $\mathbb{R}^3$ respectively.

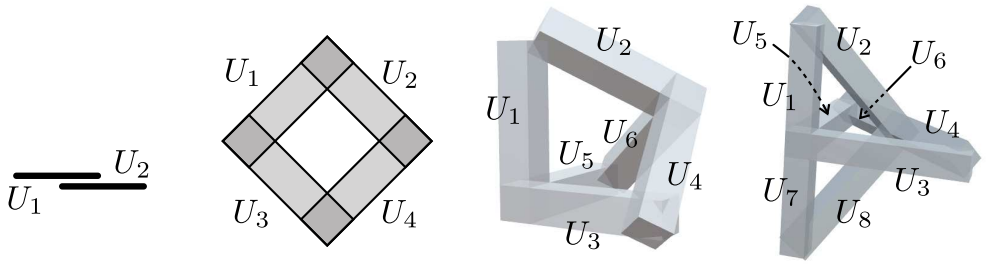

$\mathcal{T}_1$ in $\mathbb{R}^1$ $\mathcal{T}_2$ in $\mathbb{R}^2$ $\mathcal{T}_3$ in $\mathbb{R}^3$ $\mathcal{T}_4$ in $\mathbb{R}^3$

We will see that in fact each of the realizations in Example 8.2 is minimal with respect to dimension. That is, $t_1 = 1, t_2 = 2, t_3 = t_4 = 3$. To build towards this result, we first prove some general results about the minimal embedding dimensions $\{t_n \mid n \geq 1\}$.

Proposition 8.3. *For all $n \geq 1$, $t_n \leq t_{n+1} \leq t_n + 1$. That is, the sequence $\{t_n \mid n \geq 1\}$ is weakly increasing and changes by at most 1 at each step.*

Proof. The inequality $t_n \leq t_{n+1}$ follows from the fact that a realization of $\mathcal{T}_n$ can be obtained from a realization of $\mathcal{T}_{n+1}$ by simply deleting U_{2n+1} and U_{2n+2}. To prove the inequality $t_{n+1} \leq t_n + 1$ we argue that if $\mathcal{T}_n$ is open convex in $\mathbb{R}^d$, then $\mathcal{T}_{n+1}$ is open convex in $\mathbb{R}^{d+1}$.

Since $\mathcal{T}_n$ is intersection complete, we may apply Lemma 3.6 to obtain an open convex realization $\mathcal{U} = \{U_1, \ldots, U_{2n}\}$ of $\mathcal{T}_n$ in $\mathbb{R}^d$ in which disjoint U_σ have positive distance between them. We will use this to create an open convex realization of $\mathcal{T}_{n+1}$ in $\mathbb{R}^{d+1}$. To start, identify $\mathbb{R}^d$ with the subspace of $\mathbb{R}^{d+1}$ in which $x_{d+1} = 0$, and define $W_1 = U_1 \cap U_3 \cap \cdots \cap U_{2n-1}$ and $W_2 = U_2 \cap U_4 \cap \cdots \cap U_{2n}$. We may assume that the origin lies in W_1. Now choose a vector $w \in W_2$ and a small positive ε, and define a collection $\mathcal{V} = \{V_1, V_2, \ldots, V_{2n+2}\}$ as follows:

$$V_i = \begin{cases} \{v + \gamma e_{d+1} \mid v \in U_i \text{ and } 0 < \gamma < \varepsilon\} & \text{for } i = 1, 3, \ldots, 2n-1, \\ \{v + \gamma(e_{d+1} - w) \mid v \in U_i \text{ and } 0 < \gamma < \varepsilon\} & \text{for } i = 2, 4, \ldots, 2n, \\ \{v + \gamma e_{d+1} \mid v \in W_1 \text{ and } \gamma > 0\} & \text{for } i = 2n+1, \\ \{v + \gamma(e_{d+1} - w) \mid v \in W_2 \text{ and } \gamma > 0\} & \text{for } i = 2n+2. \end{cases}$$

This construction is shown below when $d = 2$. The set V_{2n+1} is a vertical prism over W_1, and the set V_{2n+2} is the skewed prism over W_2. The remaining V_i are ε-thick prisms over the corresponding U_i, with even V_i skewed at the same angle as V_{2n+2}. The origin is represented by the black dot.

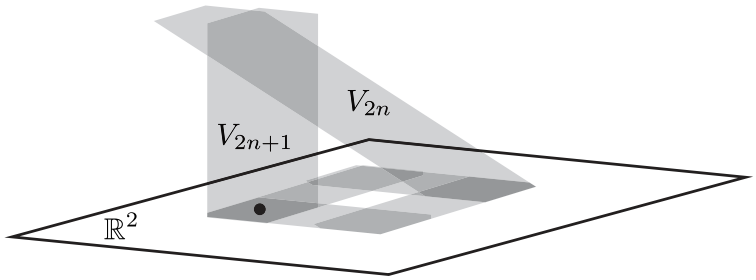

We claim that the collection $\mathcal{V}$ is an open convex realization of $\mathcal{T}_{n+1}$. First, observe that all V_i are open and convex in $\mathbb{R}^{d+1}$. To see that they form a realization of $\mathcal{T}_{n+1}$, we must check that the odd V_i and even V_i both form sunflowers, and that only the appropriate petals intersect one another.

For the odd V_i, note that $\{V_1, V_3, \ldots, V_{2n-1}\}$ is a sunflower since the odd U_i form a sunflower. Adding V_{2n+1} to this collection preserves the sunflower property since V_{2n+1} is simply the product of W_1 with an open ray. Similar logic holds for the even V_i: we see that $\{V_2, V_4, \ldots, V_{2n}\}$ forms a sunflower, and the additional petal V_{2n+2} only overlaps any other petal in the region $\{v + \gamma(e_{d+1} - w) \mid v \in W_2 \text{ and } 0 < \gamma < \varepsilon\}$, which is the intersection of all the petals.

To see that the petals overlap in the correct manner, first note that $V_{2i-1} \cap V_{2i}$ is nonempty for $i = 1, \ldots, n$ since the same holds for $U_{2i-1} \cap U_{2i}$. For $V_{2n+1} \cap V_{2n+2}$, simply note that $e_{d+1} \in V_{2n+1} \cap V_{2n+2}$ so the intersection is nonempty. Thus we have at least the appropriate overlapping between the petals of our two sunflowers, and it remains to show that no additional overlap has been introduced.

For this it suffices to argue that for all $j < k$ with different parity, the sets V_j and V_k are disjoint unless $j = 2i - 1$ and $k = 2i$. We know that this property holds for the U_i, and

since we chose a nondegenerate realization we know that disjoint U_i have positive distance between them. Except for V_{2n+1} and V_{2n+2}, all the V_i are simply a slightly thickened U_i, possibly with a small skew by the vector w. By choosing ε small enough, we can assume that the skew does not overcome the distance between disjoint U_i, so the V_i satisfy the same disjointness for $i = 1, 2, \ldots, 2n$. This leaves the case of V_{2n+1} and V_{2n+2}. For these, observe that all V_i with $i \leq 2n$ contain only points whose $(d+1)$-st coordinate is between 0 and ε. As discussed previously, the only points in V_{2n+1} and V_{2n+2} whose $(d+1)$-st coordinate satisfies these constraints are those in the center of the respective sunflowers. Thus neither of these sets overlap any petals they should not, and we have indeed formed an open convex realization of $\mathcal{T}_{n+1}$ in $\mathbb{R}^{d+1}$. This proves the result. □

Theorem 8.4. *For all n, $t_n \geq \lceil n/2 \rceil$. In particular, the sequence $\{t_n \mid n \geq 1\}$ is unbounded.*

Proof. Let $d = \lceil n/2 \rceil - 1$. We must show that $\mathcal{T}_n$ does not have an open convex realization in $\mathbb{R}^d$. Suppose for contradiction that such a realization existed, consisting of sets $\{U_1, \ldots, U_{2n}\}$. Define $V_1 = U_1 \cap U_3 \cap \cdots \cap U_{2n-1}$ and $V_2 = U_2 \cap U_4 \cap \cdots \cap U_{2n}$. Observe that V_1 and V_2 are disjoint, nonempty, convex, and open. Thus we can choose a hyperplane H separating V_1 and V_2.

Choose $p_1 \in V_1$ and $p_2 \in V_2$, and for $k \in [n]$ choose a point $q_k \in U_{2k-1} \cap U_{2k}$ (this intersection is nonempty since $\{2k-1, 2k\}$ is a codeword in $\mathcal{T}_n$). Now, for $k \in [n]$ consider the line segments $L_k = \overline{p_1 q_k}$ and $M_k = \overline{q_k p_2}$. The union $L_k \cup M_k$ forms a path that begins on one side of H and ends on the other, so for all k either L_k or M_k contains a point in H (and possibly both do). By choice of d and pigeonhole principle, either at least $d+1$ of the line segments $\{L_k\}$ contain a point in H, or at least $d+1$ of the line segments $\{M_k\}$ contain a point in H.

Without loss of generality, we may assume that at least $d+1$ of the $\{L_k\}$ contain a point p_k in H. The convex hull of these p_k lies in H, and therefore does not intersect the center V_1 of the sunflower $\{U_1, U_3, \ldots, U_{2n-1}\}$. But $L_k \subseteq U_{2k-1}$, so each p_k lies in the petal U_{2k-1}. Since there are at least $d+1$ points p_k, Theorem 5.2 implies that their convex hull (and thus H) must intersect V_1, a contradiction. □

Corollary 8.5. *The sequence $\{t_n \mid n \geq 1\}$ takes on all positive integer values.*

Proof. We know that $t_1 = 1$. Theorem 8.4 implies that the sequence is unbounded, and Proposition 8.3 tells us that it increases by at most 1 at each step. Thus it must achieve every positive integer value. □

In the remainder of this section, we determine t_n for all $n \leq 5$. The arguments used below are concrete, but seem difficult to generalize.

Proposition 8.6. *The code $\mathcal{T}_3$ does not have an open convex realization in $\mathbb{R}^2$, but does have an open convex realization in $\mathbb{R}^3$.*

Proof. An open convex realization of $\mathcal{T}_3$ in $\mathbb{R}^3$ is given in Example 8.2. Thus we just have to argue that $\mathcal{T}_3$ does not have an open convex realization in $\mathbb{R}^2$. Suppose for contradiction that $\{U_1, U_2, U_3, U_4, U_5, U_6\}$ is such a realization of $\mathcal{T}_3$ in $\mathbb{R}^2$. Choose points $q_1 \in U_1 \cap U_2$, $q_2 \in U_3 \cap U_4$, and $q_3 \in U_5 \cap U_6$. Note that $\{U_1, U_3, U_5\}$ and $\{U_2, U_4, U_6\}$ are both sunflowers and that $\{q_1, q_2, q_3\}$ is a set containing one point from each petal for both of these sunflowers. By Theorem 5.2 the triangle $\operatorname{conv}\{q_1, q_2, q_3\}$ contains a point $p_1 \in U_1 \cap U_3 \cap U_5$ and $p_2 \in U_2 \cap U_4 \cap U_6$. Since all the U_i are open sets, we may assume that $\{p_1, p_2, q_1, q_2, q_3\}$ is in

general position. The set of points $\{p_1, q_1, q_2, q_3\}$ can be visualized as follows:

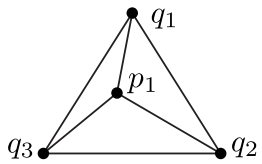

Now, p_2 falls in the interior of one of the three triangular regions surrounding p_1. Suppose that p_2 lies in the interior of $\text{conv}\{p_1, q_1, q_2\}$ (i.e. the top right triangle above). Then consider the line segment $L = \overline{p_2 q_3}$, observing that L is contained in U_6. The line segment L must cross either the line segment $\overline{p_1 q_1} \subseteq U_1$ or $\overline{p_1 q_2} \subseteq U_3$. In the former case we see that $U_6 \cap U_1 \neq \emptyset$, and in the latter $U_6 \cap U_3 \neq \emptyset$. But there is no codeword in $\mathcal{T}_3$ containing $\{1, 6\}$ or $\{3, 6\}$, so both of these situations lead to a contradiction. Thus $\mathcal{T}_3$ is not convex in $\mathbb{R}^2$. □

The lemma below will allow us to prove that $t_5 \geq 4$ by showing that if $\mathcal{T}_5$ has an open convex realization in $\mathbb{R}^3$, then $\mathcal{T}_3$ has an open convex realization in $\mathbb{R}^2$, contradicting Proposition 8.6.

Lemma 8.7. *Given five points in $\mathbb{R}^3$ in general position, there exists a plane H containing three of the points and with the remaining two points on opposite sides of H.*

Proof. Up to affine transformation we may assume that our set of points is $\{0, e_1, e_2, e_3, p\}$ where p is a point none of whose coordinates are zero. We consider two cases. First suppose that one of the coordinates of p is negative. By permuting our coordinates we can assume this is the last coordinate. Then choose $H = \text{span}\{e_1, e_2\}$. This contains the three points $0, e_1$, and e_2. Moreover since e_3 has positive last coordinate and p has negative last coordinate, they lie on opposite sides of H and the lemma follows.

Otherwise every coordinate of p is positive. In this case, write $p = (x, y, z)$ and choose $H = \text{span}\{e_3, p\}$. Observe that H contains the three points $0, e_3$, and p, and that $v = (y, -x, 0)$ is a normal vector to H. We see that $v \cdot e_1 > 0$ and $v \cdot e_2 < 0$, so the remaining two points e_1 and e_2 lie on opposite sides of H. This proves the result. □

Proposition 8.8. *The code $\mathcal{T}_5$ does not have an open convex realization in $\mathbb{R}^3$.*

Proof. Suppose for contradiction that we have an open convex realization $\mathcal{U} = \{U_1, U_2, \ldots, U_{10}\}$ of $\mathcal{T}_5$ in $\mathbb{R}^3$. For $i = 1, \ldots, 5$, choose a point p_i in the open set $U_{2i-1} \cap U_{2i}$, such that all p_i are in general position. Applying Lemma 8.7 to these five points, we obtain a hyperplane H which contains three of them, and with the remaining two on opposite sides. By permuting the labels on our realization of $\mathcal{T}_5$, we may assume that p_1, p_2, and p_3 all lie in H.

Now, consider the two tetrahedra $\Delta_1 = \text{conv}\{p_1, p_2, p_3, p_4\}$ and $\Delta_2 = \text{conv}\{p_1, p_2, p_3, p_5\}$. The vertices of these tetrahedra belong to distinct petals of the sunflowers $\{U_1, U_3, \ldots, U_9\}$ and $\{U_2, U_4, \ldots, U_{10}\}$, so by Theorem 5.2 each of these tetrahedra contain a point in the center of both of these sunflowers. Since the tetrahedra lie on opposite sides of H, each of the centers of these two sunflowers contains a point on each side of H. But the center of a sunflower is convex, and so H itself must contain a point in the center of each of the two sunflowers.

With this observation, consider the set $\mathcal{V} = \{V_1, \ldots, V_6\}$ where $V_i = U_i \cap H$. Since $H \cong \mathbb{R}^2$, we can regard this set as an open convex realization of a code in $\mathbb{R}^2$. We claim that in fact this code is $\mathcal{T}_3$. To verify this, it suffices to show that (i) $\{V_1, V_3, V_5\}$ and $\{V_2, V_4, V_6\}$ are

both sunflowers and (ii) that $V_1 \cap V_2$, $V_3 \cap V_4$, and $V_5 \cap V_6$ are nonempty, and that (iii) no other petals overlap.

Condition (i) follows from the fact that the V_i are subsets of the U_i and that the sunflowers making up the realization of $\mathcal{T}_5$ both have centers that intersect H. Condition (ii) follows by considering the points p_1, p_2, and p_3, which all lie in the desired respective intersections. Condition (iii) is a consequence of the fact that the the petals of the U_i sunflowers overlap appropriately.

However, this is a contradiction: $\mathcal{T}_3$ is not convex in $\mathbb{R}^2$ by Proposition 8.6. Thus $\mathcal{T}_5$ cannot be convex in $\mathbb{R}^3$. □

Corollary 8.9. *The sequence t_n begins as follows:*

n	1	2	3	4	5
t_n	1	2	3	3	4

Proof. Clearly $t_1 = 1$ since $\mathcal{T}_1$ is convex in $\mathbb{R}^1$ but has more than one codeword, so is not convex in $\mathbb{R}^0$. The code $\mathcal{T}_2$ has a realization in $\mathbb{R}^2$ as given in Example 8.2, but has no realization in $\mathbb{R}^1$ since any realization contains a non-crossing loop. Thus $t_2 = 2$.

Note that $t_3 \leq 3$ and $t_4 \leq 3$ by Example 8.2, and both bounds are tight by Proposition 8.6 and monotonicity of the t_n. By Proposition 8.8 we know that $t_5 \geq 4$, and simultaneously Proposition 8.3 implies that $t_5 \leq t_4 + 1 = 4$. This proves the result. □

The proofs presented in Propositions 8.6 and 8.8 are both somewhat ad hoc and do not seem ripe for generalization. Determining t_n for $n \geq 6$ remains an open problem, perhaps of significant difficulty.

9. Contextualizing Our Results via Code Minors

In this section we will situate our results in the framework of code morphisms and minors. We begin in Section 9.1 by recalling several basic definitions and results regarding morphisms and minors, following the most recent treatment of this material which appears in [15, Chapters 3 and 4]. In Section 9.2 we explain how the families of codes $\mathcal{S}_n$, $\mathcal{S}_\Delta$, and $\mathcal{T}_n$ relate to one another in the framework of minors. We then generalize the codes $\mathcal{S}_n$ and $\mathcal{S}_\Delta$ as well as their accompanying results in Section 9.3.

9.1. Minors of Codes.

In [14], we introduced a framework of morphisms for neural codes. Morphisms allow us to define a notion of "minors" for codes, analogous to (but not a generalization of) minors of graphs or matroids. Morphisms and minors have a strong relationship to convexity (see Theorem 9.3 below), and provide a useful context in which to state and compare results about convex neural codes.

For the definition below, recall from Definition 2.1 that a trunk in a code $\mathcal{C}$ is a set of the form $\mathrm{Tk}_{\mathcal{C}}(\sigma) := \{c \in \mathcal{C} \mid \sigma \subseteq c\}$ for some $\sigma \subseteq [n]$.

Definition 9.1. Let $\mathcal{C}$ and $\mathcal{D}$ be codes. A function $f : \mathcal{C} \to \mathcal{D}$ is called a *morphism* if the preimage of any proper trunk in $\mathcal{D}$ under f is a proper trunk in $\mathcal{C}$. An *isomorphism* is a morphism with an inverse function that is also a morphism.

Definition 9.2. We say that a code $\mathcal{D}$ is a *minor* of a code $\mathcal{C}$ if there exists a surjective morphism $f : \mathcal{C} \to \mathcal{D}$.

The relation "$\mathcal{D}$ is a minor of $\mathcal{C}$" forms a partial order on isomorphism classes of codes. We denote the resulting partially ordered set by $\mathbf{P}_{\mathbf{Code}}$, and write $\mathcal{D} \leq \mathcal{C}$ when $\mathcal{D}$ is a minor of $\mathcal{C}$. Note that morphisms and minors are defined in a purely combinatorial manner. However, minors can be used to understand geometric information about realizations of codes, as the following theorem indicates.

Theorem 9.3. *The following properties are* minor-closed *(that is, if $\mathcal{C}$ has one of the properties below, then so does every $\mathcal{D} \leq \mathcal{C}$):*

- *Open convexity in $\mathbb{R}^d$,*
- *Closed convexity in $\mathbb{R}^d$,*
- *Non-degenerate open/closed convexity in $\mathbb{R}^d$ (see Definition 3.5),*
- *Intersection completeness.*

Theorem 9.3 follows from a more general observation, first noted by Caitlin Lienkaemper: if $\mathcal{C} \subseteq 2^{[n]}$ is a code with a (possibly not convex or open) realization $\mathcal{U} = \{U_1, \ldots, U_n\}$ in a space X, then there is a bijection

$$\{ \text{ Minors of } \mathcal{C} \ \} \quad \longleftrightarrow \quad \left\{ \begin{array}{c} \text{Codes that can be} \\ \text{realized in } X \text{ using} \\ \text{sets of the form } U_\sigma \end{array} \right\}.$$

For details on this result, see [15, Section 4.2] and [18, Section 4].

Theorem 9.3 implies that intersection completeness is an isomorphism invariant, and that restricting our attention to intersection complete codes amounts to restricting to a minor-closed family in $\mathbf{P}_{\mathbf{Code}}$. Throughout the rest of this section, we will examine exclusively intersection complete codes, with the partial order inherited from $\mathbf{P}_{\mathbf{Code}}$. We will focus on the open embedding dimensions of these codes.

One can visualize minors as stratifying intersection complete codes into different "layers" according to their open embedding dimensions, as sketched in the figure below. In the figure, "Open convex in $\mathbb{R}^d$" simply refers to codes whose open embedding dimension is equal to d.

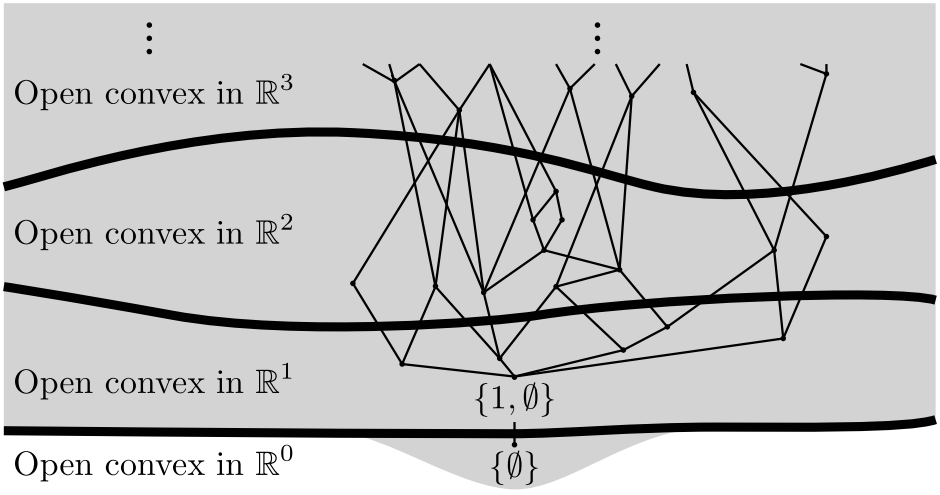

Note that $\{\emptyset\}$ is the only intersection complete code whose open embedding dimension is zero. For any $d \geq 1$, however, there are infinitely many intersection complete codes with open embedding dimension d.

The figure above is slightly misleading: each "layer" of codes with open embedding dimension d is not finitely thick. Indeed, each layer may contain chains that are infinitely long ("tall"), and antichains that are infinitely large ("wide").

Our aim in Section 9.2 will be to understand where the codes we have constructed in this paper sit inside this partial order. In Section 9.3, we will provide some more general examples and results using Theorem 1.13.

We will make heavy use of the following definition and proposition, which give a combinatorial description of all morphisms. For details, see [15, Section 3.2].

Definition 9.4. Let $\mathcal{C} \subseteq 2^{[n]}$ be a code, and for $i \in [m]$ let $T_i \subseteq \mathcal{C}$ be a proper trunk. The *morphism determined by the trunks* $\{T_1, \dots, T_m\}$ is the map $f : \mathcal{C} \to 2^{[m]}$ given by $f(c) = \{i \in [m] \mid c \in T_i\}$.

Proposition 9.5. *The map described in Definition 9.4 is a morphism from $\mathcal{C}$ to $2^{[m]}$. Moreover, every morphism arises in this way. Formally, for codes $\mathcal{C} \subseteq 2^{[n]}$ and $\mathcal{D} \subseteq 2^{[m]}$, and any morphism $f : \mathcal{C} \to \mathcal{D}$, f is the morphism determined by the trunks $\{T_i := f^{-1}(\mathrm{Tk}_{\mathcal{D}}(i)) \mid i \in [m]\}$ (restricted to the range $\mathcal{D}$). Equivalently, for all $c \in \mathcal{C}$,*

$$f(c) = \{i \in [m] \mid c \in f^{-1}(\mathrm{Tk}_{\mathcal{D}}(i))\}.$$

9.2. **The codes $\mathcal{S}_n, \mathcal{S}_\Delta$, and $\mathcal{T}_n$ in $\mathbf{P_{Code}}$.** Let us begin by establishing a relationship between codes of the type $\mathcal{S}_n$ and the type $\mathcal{S}_\Delta$. Recall that $\mathcal{S}_n$ is a special case of $\mathcal{S}_\Delta$—in particular, $\mathcal{S}_n = \mathcal{S}_\Delta$ where Δ is n points. More generally, we have the following:

Proposition 9.6. *Let $\Delta \subseteq 2^{[n]}$ be a simplicial complex with m facets. Then there exists a surjective morphism $\mathcal{S}_\Delta \to \mathcal{S}_m$. In particular, $\mathcal{S}_m \leq \mathcal{S}_\Delta$.*

Proof. Let $F_1, \dots, F_m$ be the facets of Δ. For $i \in [m]$ define $T_i = \mathrm{Tk}_{\mathcal{S}_\Delta}(F_i)$, and define $T_{m+1} = \mathrm{Tk}_{\mathcal{S}_\Delta}(n+1)$. We claim that $\mathcal{S}_m$ is the image of $\mathcal{S}_\Delta$ under the morphism f determined by the trunks $\{T_1, T_2, \dots, T_{m+1}\}$. Recall from Definition 1.10 that the codewords of $\mathcal{S}_\Delta$ are:

- σ for $\sigma \in \Delta$,
- $\sigma \cup \{n+1\}$ for $\sigma \in \Delta$, and
- $[n]$.

The images of these codewords under f are as follows:

- $f(\sigma)$ is equal to $\emptyset$ if σ is not a facet of Δ, and equal to $\{i\}$ if $\sigma = F_i$,
- $f(\sigma \cup \{n+1\})$ is equal to $\{m+1\}$ if σ is not a facet of Δ, and equal to $\{i, m+1\}$ if $\sigma = F_i$, and
- $f([n]) = [m]$ since $[n]$ contains all facets of Δ, but does not contain $m+1$.

But comparing these images to Definition 5.3, we see that these are exactly the codewords of $\mathcal{S}_m$, proving the result. □

Remark 9.7. One way to think of Proposition 9.6 is as follows. The set

$$\{\mathcal{S}_\Delta \mid \Delta \text{ is a simplicial complex with } m \text{ facets}\}$$

inherits a partial order from $\mathbf{P_{Code}}$, and with this inherited order $\mathcal{S}_m$ is the unique minimal element of the set. Theorem 1.11 says that for $m \geq 2$ all of these live in the "layer" of codes

with open embedding dimension m. We can visualize this situation as follows.

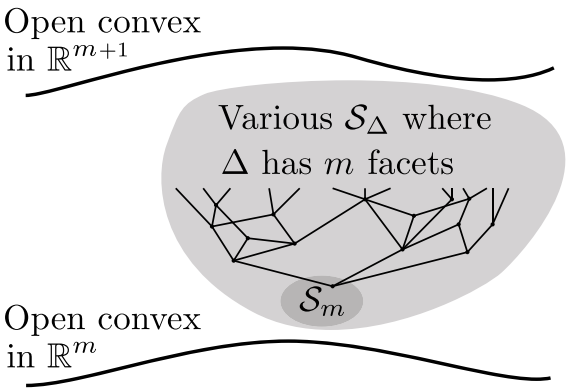

It is also worth noting the following, regarding the codes $\mathcal{T}_n$ described in Section 8.

Proposition 9.8. *For any $n \geq 1$, $\mathcal{T}_n \leq \mathcal{T}_{n+1}$. In particular, the codes $\{\mathcal{T}_n \mid n \geq 1\}$ form a chain in $\mathbf{P}_{\mathbf{Code}}$.*

Proof. Given a set of neurons $\sigma \subseteq [n]$, one can form a natural "restriction" of a code $\mathcal{C} \subseteq 2^{[n]}$ by mapping $c \mapsto c \cap \sigma$. This restriction is a morphism (see [14, Section 2]). In the case of the codes $\mathcal{T}_n \subseteq 2^{[2n]}$, one can note that $\mathcal{T}_n$ is the image of $\mathcal{T}_{n+1}$ under the restriction map with $\sigma = [2n] \subseteq [2n+2]$. This surjective morphism from $\mathcal{T}_{n+1}$ to $\mathcal{T}_n$ implies that $\mathcal{T}_n \leq \mathcal{T}_{n+1}$ as desired. □

9.3. **Generalizing $\mathcal{S}_n$ and $\mathcal{S}_\Delta$ using Theorem 1.13.** We begin with a definition generalizing that of $\mathcal{S}_\Delta$.

Definition 9.9. Let $\mathcal{D} \subseteq \mathcal{C} \subseteq 2^{[n]}$ be intersection complete codes. We define

$$\mathcal{S}_{\mathcal{C}/\mathcal{D}} := \mathcal{C} \cup \{[n]\} \cup \{d \cup \{n+1\} \mid d \in \mathcal{D}\} \subseteq 2^{[n+1]}.$$

Note that choosing $\mathcal{D} = \{\text{minimal nonempty codewords in } \mathcal{C}\} \cup \{\emptyset\}$ always satisfies the above conditions. In this case, we will let $\mathcal{S}_{\mathcal{C}/\min}$ denote $\mathcal{S}_{\mathcal{C}/\mathcal{D}}$.

Qualitatively, $\mathcal{S}_{\mathcal{C}/\mathcal{D}}$ is the result of forming a flexible sunflower using the codewords in $\mathcal{C}$, and then "gluing" the petals of that sunflower to a new set U_{n+1} along codewords in $\mathcal{D}$. Observe that $\mathcal{S}_\Delta$ of Definition 1.10 is equal to $\mathcal{S}_{\Delta/\Delta}$ in this notation. Also, if $\mathcal{C} = \{\{1\}, \{2\}, \ldots, \{n\}, \emptyset\}$, then we see $\mathcal{S}_n$ of Definition 5.3 is equal to $\mathcal{S}_{\mathcal{C}/\min}$.

Proposition 9.10. *Let $\mathcal{D} \subseteq \mathcal{C} \subseteq 2^{[n]}$ be intersection complete codes. The code $\mathcal{S}_{\mathcal{C}/\mathcal{D}}$ is intersection complete. If $\mathcal{D}$ has m maximal codewords and does not contain $[n]$, then $\mathcal{S}_{\mathcal{C}/\mathcal{D}}$ has $m+1$ maximal codewords. In particular,* $\operatorname{odim}(\mathcal{S}_{\mathcal{C}/\mathcal{D}}) \leq \max\{2, m\}$.

Proof. Codewords in $\mathcal{S}_{\mathcal{C}/\mathcal{D}}$ come in three types: codewords from $\mathcal{C}$, the codeword $[n]$, and those of the form $d \cup \{n+1\}$ where $d \in \mathcal{D}$. Since $\mathcal{C}$ and $\mathcal{D}$ are intersection complete, the intersection of two codewords of the same type always yields another codeword of that type (and hence lying in $\mathcal{S}_{\mathcal{C}/\mathcal{D}}$). This leaves the intersections of codewords of different types. The intersection of a codeword in $\mathcal{C}$ with $[n]$ is simply the same codeword in $\mathcal{C}$. The intersection of $d \cup \{n+1\}$ with $[n]$ is just d, which lies in $\mathcal{S}_{\mathcal{C}/\mathcal{D}}$ since $\mathcal{D} \subseteq \mathcal{C}$. Finally, the intersection of $c \in \mathcal{C}$ with $d \cup \{n+1\}$ is $c \cap d$, which lies in $\mathcal{C}$ since since $\mathcal{C}$ is intersection complete.

For the second part of the statement, note that if d is a maximal codeword of $\mathcal{D}$, then $d \cup \{n+1\}$ is a maximal codeword of $\mathcal{S}_{\mathcal{C}/\mathcal{D}}$. Since $[n] \notin \mathcal{D}$, the codeword $[n]$ is also a maximal

codeword of $\mathcal{S}_{\mathcal{C}/\mathcal{D}}$, yielding $m+1$ total maximal codewords. The bound on $\operatorname{odim}(\mathcal{S}_{\mathcal{C}/\mathcal{D}})$ follows immediately from [3, Theorem 1.2]. □

The following proposition provides a generalization of Theorem 5.5 to the codes $\mathcal{S}_{\mathcal{C}/\min}$.

Proposition 9.11. *Let $\mathcal{C} \subseteq 2^{[n]}$ be an intersection complete code that contains every singleton set. Then*

$$\operatorname{odim}(\mathcal{S}_{\mathcal{C}/min}) \geq \left\lceil \frac{n}{\dim(\Delta(\mathcal{C})) + 1} \right\rceil.$$

Proof. We start with a degenerate case: if $n = 1$, then $\mathcal{C} = \{\emptyset, 1\}$ and $\mathcal{S}_{\mathcal{C}/\min} = \{12, 1, 2, \emptyset\}$. In this case $\operatorname{odim}(\mathcal{S}_{\mathcal{C}/\min}) = 1$, while $n = 1$ and $\dim(\Delta(\mathcal{C})) + 1 = 1$. We see that the bound given above is satisfied as desired.

Otherwise, $n \geq 2$. In this case, let $\{U_1, \ldots, U_{m+1}\}$ be an open convex realization of $\mathcal{S}_{\mathcal{C}/\min}$ in $\mathbb{R}^d$. Since the minimal nonempty codewords of $\mathcal{C}$ are all singletons, the code $\mathcal{S}_{\mathcal{C}/\min}$ consists of codewords from $\mathcal{C}$, the codeword $[n]$, codewords of the form $\{i, n+1\}$ where $i \in [n]$, and lastly the codeword $\{n+1\}$. Since $[n]$ is a codeword, the sets $\{U_1, \ldots, U_n\}$ all meet in a central point. In particular, $\{U_1, \ldots, U_n\}$ is a k-flexible sunflower, where k is the largest weight of a codeword in $\mathcal{C}$ other than possibly $[n]$. In particular $k \leq \dim(\Delta(\mathcal{C})) + 1$, with equality if $[n] \notin \mathcal{C}$.

But consider the set U_{n+1}. This set does not meet $U_{[n]}$ since $[n+1]$ is not a codeword of $\mathcal{S}_{\mathcal{C}/\min}$. However, it does touch each U_i since $\{i, n+1\}$ is a codeword. If we choose $p_i \in U_i \cap U_{n+1}$, then the convex hull of $\{p_1, \ldots, p_n\}$ is contained in U_{n+1} and therefore does not contain a point in the center of $\{U_1, \ldots, U_n\}$. By Theorem 1.13, such a sampling of p_i cannot be chosen if $n \geq dk + 1$. Therefore we must have $n \leq dk$. Rearranging, this implies $d \geq \lceil n/k \rceil$. Using the inequality $k \leq \dim(\Delta(\mathcal{C})) + 1$ yields the result. □

The added assumption in Proposition 9.11 that $\mathcal{C}$ contains all singletons is not too restrictive, since adding singletons to an intersection complete code always maintains intersection completeness.

Continuing our pattern of generalizations, the proposition below is analogous to Theorem 1.11 and its second part generalizes Proposition 9.6.

Proposition 9.12. *Let $\mathcal{D} \subseteq \mathcal{C} \subseteq 2^{[n]}$ be intersection complete codes. Let $m \geq 2$ be the number of maximal codewords in $\mathcal{D}$, and let k be the largest number of maximal codewords in $\mathcal{D}$ whose union lies in $\Delta(\mathcal{C})$. Then there exists an intersection complete code $\mathcal{E} \subseteq 2^{[m]}$ containing all singleton sets such that (i) $k = \dim(\Delta(\mathcal{E}))+1$, and (ii) there exists a surjective morphism $\mathcal{S}_{\mathcal{C}/\mathcal{D}} \to \mathcal{S}_{\mathcal{E}/min}$. In particular, $\mathcal{S}_{\mathcal{E}/min} \leq \mathcal{S}_{\mathcal{C}/\mathcal{D}}$ and $m \geq \operatorname{odim}(\mathcal{S}_{\mathcal{C}/\mathcal{D}}) \geq \left\lceil \frac{m}{k} \right\rceil$.*

Proof. We will mirror the proof of Proposition 9.6. Let $F_1, \ldots, F_m$ be the maximal codewords of $\mathcal{D}$. For $i \in [m]$ define $T_i = \operatorname{Tk}_{\mathcal{S}_{\mathcal{C}/\mathcal{D}}}(F_i)$, and define $T_{m+1} = \operatorname{Tk}_{\mathcal{S}_{\mathcal{C}/\mathcal{D}}}(n+1)$. Let us consider the image of $\mathcal{S}_{\mathcal{C}/\mathcal{D}}$ under the morphism f determined by $\{T_1, \ldots, T_{m+1}\}$. Recall from Definition 9.9 that the codewords of $\mathcal{S}_{\mathcal{C}/\mathcal{D}}$ come in the following types:

- c for $c \in \mathcal{C}$,
- $d \cup \{n+1\}$ for $d \in \mathcal{D}$, and
- $[n]$.

The images of these codewords under f are as follows:

- $f(c)$ is equal to $\{i \in [m] \mid c \text{ contains } F_i\}$,

- $f(d \cup \{n+1\})$ is equal to $\{m+1\}$ if d is not equal to some F_i, and is equal to $\{i, m+1\}$ if $d = F_i$,
- $f([n]) = [m]$ since $[n]$ contains all maximal codewords in $\mathcal{D}$, but not $n+1$.

Let $\mathcal{E} \subseteq 2^{[m]}$ be the collection of codewords in the first bullet above, i.e. $\mathcal{E}$ is the image of $\mathcal{C}$ under f. Since the image of an intersection complete code is again intersection complete, we see that $\mathcal{E}$ is intersection complete. Moreover, $\mathcal{E}$ contains every singleton set since $f(F_i) = \{i\}$.

The image of $\mathcal{S}_{\mathcal{C}/\mathcal{D}}$ under f therefore contains codewords in $\mathcal{E}$, codewords of the form $\{i, m+1\}$ for all $i \in [m]$, the codeword $\{m+1\}$, and $[m]$. But these are exactly the codewords of $\mathcal{S}_{\mathcal{E}/\min}$. Thus $\mathcal{S}_{\mathcal{E}/\min}$ is the image of $\mathcal{S}_{\mathcal{C}/\mathcal{D}}$ under f.

To prove the result, it remains to show that $k = \dim(\Delta(\mathcal{E})) + 1$. The codewords in $\mathcal{E}$ are of the form $f(c) = \{i \in [m] \mid F_i \subseteq c\}$. Thus a codeword in $\mathcal{E}$ corresponds to a collection of maximal codewords in $\mathcal{D}$ all of which are contained in some $c \in \mathcal{C}$. A codeword in $\mathcal{E}$ with largest weight thus corresponds to a largest possible collection of maximal codewords in $\mathcal{D}$ whose union is contained in $\Delta(\mathcal{C})$. The largest such collection has size k by definition, so any largest codeword in $\mathcal{E}$ has weight k, proving the result. □

Remark 9.13. Generalizing Remark 9.7 from the last section, we see that among all codes of the form $\mathcal{S}_{\mathcal{C}/\mathcal{D}}$ with parameters m and k as described in Proposition 9.12, the minimal elements are always of the form $\mathcal{S}_{\mathcal{E}/\min}$ where $\mathcal{E} \subseteq 2^{[m]}$ contains all singletons, and $k = \dim(\Delta(\mathcal{E})) + 1$. The following diagram shows this:

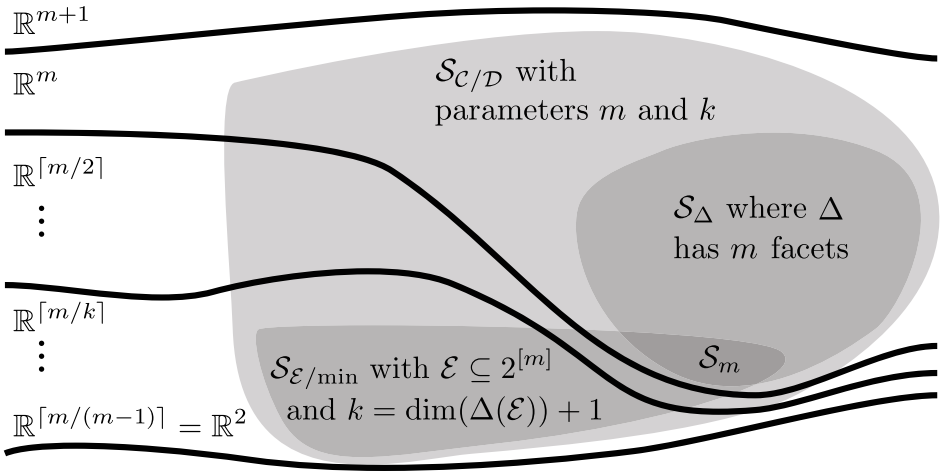

These results use Theorem 1.13 to provide a more complete picture of the open embedding dimensions of intersection complete codes. There is still much to be done, however. As one example, the bound $m \geq \operatorname{odim}(\mathcal{S}_{\mathcal{C}/\mathcal{D}}) \geq \left\lceil \frac{m}{k} \right\rceil$ of Proposition 9.12 leaves quite a large gap for $k \geq 2$. Sharpening this bound based on the combinatorial structure of $\mathcal{C}$ and $\mathcal{D}$ would be a natural task of interest.

10. Conclusion

We have seen a number of phenomena arise in the closed and open embedding dimensions of intersection complete codes. Some of these, like Theorems 1.6, 1.8, and 1.9, gave us improved control over the embedding dimensions. Others, like Theorem 1.11, showed that embedding dimension may be difficult to control. With Theorem 1.13, we developed new

tools to understand open embedding dimension using k-flexible sunflowers, but the picture is still far from complete.

One direction for future work would be to search for analogous phenomena among codes that are not intersection complete. One could start with the following.

Question 10.1. *Does there exist a code $\mathcal{C}$ with* $\operatorname{odim}(\mathcal{C}) < \operatorname{cdim}(\mathcal{C}) < \infty$*?*

Theorem 1.8 tells us that such a code cannot be intersection complete. There are examples due to [10, 3] of codes with $\operatorname{odim}(\mathcal{C}) < \operatorname{cdim}(\mathcal{C}) = \infty$. These examples rely on a compactness argument to construct contradictory line segments of minimum distance in a hypothetical closed realization, proving $\operatorname{cdim}(\mathcal{C}) = \infty$. A similar approach, paired with a classic convexity theorem that depends on dimension such as Radon's theorem, could yield a positive answer to Question 10.1, and also possibly Question 10.2 below.

Question 10.2. *Little is known about whether closed embedding dimension can be large relative to the number of neurons, n. A few open areas to investigate are the following, in increasing order of difficulty:*

- *Does there exist a code $\mathcal{C} \subseteq 2^{[n]}$ for which* $\operatorname{cdim}(\mathcal{C})$ *is finite, but larger than $n-1$?*
- *Does there exist a family of codes $\{\mathcal{C}_n \subseteq 2^{[n]} \mid n \geq 1\}$ such that* $\operatorname{cdim}(\mathcal{C}_n)$ *grows faster than any linear function of n?*
- *Does there exist a family of codes $\{\mathcal{C}_n \subseteq 2^{[n]} \mid n \geq 1\}$ such that* $\operatorname{cdim}(\mathcal{C}_n)$ *grows faster than any polynomial function of n?*

Note that Theorem 1.9 tells us that if such codes exist, they cannot be intersection complete. We have provided an affirmative answers to the odim *versions of the above questions via the family $\mathcal{S}_\Delta$, and in particular Corollary 6.3.*

Regarding the tangled sunflower codes $\mathcal{T}_n$ of Section 8, there is much to be done. A good first step would be to improve the embedding dimension bounds that we currently have, or, more ambitiously, find an exact characterization of the embedding dimension.

Question 10.3. *Does there exist an explicit characterization of the open embedding dimensions t_n described in Definition 8.1? Can we improve the bounds of $\lceil n/2 \rceil \leq t_n \leq n$?*

One might also consider codes that describe more than two sunflowers whose petals are "tangled" (i.e. incident) in some way. This would be a significantly more complicated problem, but perhaps of some interest. Another generalization would be to consider a notion of tangled flexible sunflowers. This would be even more challenging to investigate, but would perhaps be more relevant to applications in experimental data.

Question 10.4. *In Section 9 we contextualized our results via a partial order on codes, denoted by $\mathbf{P}_{\mathbf{Code}}$. In this partial order, both* odim *and* cdim *are monotone functions. In* [14] *we showed that a code is intersection complete if and only if it lies below a simplicial complex in $\mathbf{P}_{\mathbf{Code}}$. An interesting question is thus the following: do the simplicial complexes lying above an intersection complete code $\mathcal{C}$ in $\mathbf{P}_{\mathbf{Code}}$ determine* $\operatorname{odim}(\mathcal{C})$*? That is, among the simplicial complexes lying above $\mathcal{C}$ in $\mathbf{P}_{\mathbf{Code}}$, does one have minimal embedding dimension equal to* $\operatorname{odim}(\mathcal{C})$*?*

A positive answer to the above question would reduce the problem of determining open embedding dimension for intersection complete codes to the problem of determining open embedding dimension for simplicial complexes, which is very closely tied to the well-studied problem of determining when a complex is d-representable, as described in [25, Section 1.2].

Note that the answer to Question 10.4 cannot be positive when we replace odim with cdim. Open and closed embedding dimension for simplicial complexes are always the same (recall Theorem 1.6), but the code $\mathcal{S}_3$ already shows that closed dimension and open dimension are different for intersection complete codes. In particular, $\operatorname{cdim}(\mathcal{S}_3) = 2$ while $\operatorname{odim}(\mathcal{S}_3) = 3$. Thus the simplicial complexes that $\mathcal{S}_3$ is a minor of do not determine its closed embedding dimension, at least not simply as the minimum of their closed embedding dimensions, which must be at least 3.

In the proof of Theorem 1.11 (see Section 6), we showed that $\operatorname{odim}(\mathcal{S}_\Delta)$ was equal to the number of facets in Δ by showing that any realization of $\mathcal{S}_\Delta$ gives rise to a realization of the code $\mathcal{S}_m$ described in Section 5. Later, in Proposition 9.6, we translated this argument to the context of minors, demonstrating a surjective morphism from $\mathcal{S}_\Delta$ to $\mathcal{S}_m$. This technique could be generalized to analyze arbitrary codes as follows. Given a code $\mathcal{C}$, look for the largest m so that there is a surjective morphism $\mathcal{C} \to \mathcal{S}_m$ (i.e. the largest m so that $\mathcal{S}_m$ is a minor of $\mathcal{C}$). This largest m then provides a lower bound on the open embedding dimension of $\mathcal{C}$.

Importantly, this method is distinct from existing techniques for providing lower bounds on $\operatorname{odim}(\mathcal{C})$, which rely on homological information obtained from $\Delta(\mathcal{C})$ (see for example [7]). In contrast, the sunflower approach is completely agnostic to homology of $\Delta(\mathcal{C})$. Whether sunflowers of convex open sets could be useful in analyzing experimental data may be an interesting open question. As a start, it would be useful to determine whether searching for specific minors of a given code can be done in a computationally feasible manner.

Question 10.5. *Given a code $\mathcal{C} \subseteq 2^{[n]}$, is there an efficient algorithm to determine the largest m so that $\mathcal{S}_m$ is a minor of $\mathcal{C}$? More generally, for what pairs of codes $\mathcal{C} \subseteq 2^{[n]}$ and $\mathcal{D} \subseteq 2^{[m]}$ can one efficiently recognize whether or not $\mathcal{D}$ is a minor of $\mathcal{C}$?*

Acknowledgements

We would like to thank Florian Frick for raising the question of whether there exist open convex codes $\mathcal{C} \subseteq 2^{[n]}$ with $\operatorname{odim}(\mathcal{C}) > n-1$, and for interesting discussions on this question. Isabella Novik provided detailed feedback on initial drafts of this paper, as well as helpful discussions on its mathematical content. Anne Shiu also provided thorough feedback on an initial draft of the paper. Finally, we would like to thank the anonymous referees for their comments, which greatly improved the presentation of our results.

References

[1] Nina Amenta, Jesús A. De Loera, and Pablo Soberón. Helly's theorem: new variations and applications. In *Algebraic and geometric methods in discrete mathematics*, volume 685 of *Contemp. Math.*, pages 55–95. Amer. Math. Soc., Providence, RI, 2017.

[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 Journal on Applied Algebra and Geometry*, 1(1):222–238, 2017.

[5] Carina Curto, Elizabeth Gross, Jack Jeffries, Katherine Morrison, Zvi Rosen, Anne Shiu, and Nora Youngs. Algebraic signatures of convex and non-convex codes. *Journal of Pure and Applied Algebra*, 223(9):3919–3940, 2019.

[6] 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.
[7] Carina Curto and Ramón Vera. The Leray dimension of a convex code. *arXiv e-prints: 1612.07797*, 2016.
[8] Megan K. Franke and Samuel Muthiah. Every binary code can be realized by convex sets. *Advances in Applied Mathematics*, 99:83–93, 2017.
[9] Rebecca Garcia, Luis Garcia-Puente, Ryan Kruse, Jessica Liu, Dane Miyata, Ethan Petersen, Kaitlyn Phillipson, and Anne Shiu. Gröbner bases of neural ideals. *International Journal of Algebra and Computation*, 28(4):553–571, 2018.
[10] Sarah Ayman Goldrup and Kaitlyn Phillipson. Classification of open and closed convex codes on five neurons. *Advances in Applied Mathematics*, 112:101948, 2020.
[11] Sema Gunturkun, Jack Jeffries, and Jeffrey Sun. Polarization of neural rings. *Journal of Algebra and Its Applications*, 19(8), 2019.
[12] Vladimir Itskov, Alexander Kunin, and Zvi Rosen. Hyperplane neural codes and the polar complex. In Nils A. Baas, Gereon Quick, Markus Szymik, Marius Thaule, and Gunnar E. Carlsson, editors, *Topological Data Analysis - The Abel Symposium, 2018*, Abel Symposia, pages 343–369. Springer, 2020.
[13] R. Amzi Jeffs. Sunflowers of convex open sets. *Advances in Applied Mathematics*, 111:101935, 2019.
[14] R. Amzi Jeffs. Morphisms of neural codes. *SIAM Journal on Applied Algebra and Geometry*, 4:99–122, 2020.
[15] R. Amzi Jeffs. *Morphisms, Minors, and Minimal Obstructions to Convexity of Neural Codes*. PhD thesis, University of Washington, Seattle, 2021. Available online at `http://hdl.handle.net/1773/48062`.
[16] R. Amzi Jeffs and Isabella Novik. Convex union representability and convex codes. *International Mathematics Research Notices*, 2019.
[17] R. Amzi Jeffs, Mohamed Omar, Natchanon Suaysom, Aleina Wachtel, and Nora Youngs. Sparse neural codes and convexity. *Involve, a Journal of Mathematics*, 12(5):737–754, 2015.
[18] Alexander Kunin, Caitlin Lienkaemper, and Zvi Rosen. Oriented matroids and combinatorial neural codes. *arXiv e-prints: 2002.03542*, page arXiv:2002.03542, 2020.
[19] Caitlin Lienkaemper, Anne Shiu, and Zev Woodstock. Obstructions to convexity in neural codes. *Advances in Applied Mathematics*, 85:31–59, 2017.
[20] Jiří Matoušek. *Lectures on discrete geometry*, volume 212 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 2002.
[21] Raffaella Mulas and Ngoc M Tran. Minimal embedding dimensions of connected neural codes. *Algebraic Statistics*, 11(1):99–106, 2020.
[22] Alexander Ruys de Perez, Laura Felicia Matusevich, and Anne Shiu. Neural codes and the factor complex. *Advances in Applied Mathematics*, 114:101977, 2020.
[23] Alexander Schrijver. *Theory of integer and linear programming*. Discrete Mathematics and Optimization. Wiley Interscience, 1986.
[24] Martin Tancer. d-representability of simplicial complexes of fixed dimension. *Journal of Computational Geometry*, 2(1):183–188, 2011.
[25] 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.
[26] Günter M. Ziegler. *Lectures on polytopes*, volume 152 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1995.