NON-MONOTONICITY OF CLOSED CONVEXITY IN NEURAL CODES

Brianna Gambacini
Department of Mathematics
North Carolina State University
begambac@ncsu.edu

R. Amzi Jeffs
Department of Mathematics
University of Washington
rajeffs@uw.edu

Sam Macdonald
Department of Mathematics
Willamette University
smmacdonald@willamette.edu

Anne Shiu
Department of Mathematics
Texas A&M University
annejls@math.tamu.edu

March 2021

ABSTRACT

Neural codes are lists of subsets of neurons that fire together. Of particular interest are neurons called place cells, which fire when an animal is in specific, usually convex regions in space. A fundamental question, therefore, is to determine which neural codes arise from the regions of some collection of open convex sets or closed convex sets in Euclidean space. This work focuses on how these two classes of codes – open convex and closed convex codes – are related. As a starting point, open convex codes have a desirable monotonicity property, namely, adding non-maximal codewords preserves open convexity; but here we show that this property fails to hold for closed convex codes. Additionally, while adding non-maximal codewords can only increase the open embedding dimension by 1, here we demonstrate that adding a single such codeword can increase the closed embedding dimension by an arbitrarily large amount. Finally, we disprove a conjecture of Goldrup and Phillipson, and also present an example of a code that is neither open convex nor closed convex.

Keywords Neural code · Place cell · Convex · Simplicial complex

1 Introduction

Place cells are neurons that fire (are active) when an animal is in specific locations [1]. The resulting subsets of neurons that fire together, called a neural code, can be used by the brain to form a mental map of an animal's environment. Place cells were discovered by John O'Keefe in 1971, earning him a joint (with May-Britt Moser and Edvard Moser) Nobel Prize in Physiology or Medicine in 2014.

The specific location where a place cell fires is called its *place field*, and this set is typically modeled by a convex set. Thus, neural codes arising from place cells describe the regions cut out by intersecting convex sets. This motivates the following question: Which neural codes arise from open convex sets in some Euclidean space? (Each set is required to be open to account for the fact that place fields are full-dimensional, i.e. they have nonempty interior.) Many investigations into this question have been made in recent years (for instance, [2, 3, 4, 5, 6, 7, 8, 9, 10]). Recent work such as [11] shows that this question strictly generalizes the closely related topic of *intersection patterns* of convex sets (see [12] for an overview).

In this work, we consider the above question, and also, following [2, 6], the analogous question for closed convex sets. Additionally, we ask how these two classes of codes – open convex and closed convex codes – are related. Which codes are open convex but not closed convex (or vice-versa)? Which codes are neither open convex nor closed convex?

One starting point of our work is a recent "monotonicity" result of Cruz *et al.* [2]: If two codes $\mathcal{C}$ and $\mathcal{C}'$, with $\mathcal{C} \subset \mathcal{C}'$, generate the same simplicial complex, and $\mathcal{C}$ is open convex, then so is $\mathcal{C}'$ (see Proposition 2.12). Hence, as open convexity is "inherited" from $\mathcal{C}$ to $\mathcal{C}'$, this result greatly simplifies the analysis of open convex codes. However, Cruz *et al.* did not know whether the analogous result holds for closed convexity [2], and here we show that, somewhat surprisingly, it does *not* (Theorem 3.2).

The mononotonicity result of Cruz *et al.* mentioned above can be paraphrased as follows: adding non-maximal codewords to an open convex code yields another open convex code. The open convex realization of the larger code may need to be in a Euclidean space of a higher dimension – but this dimension need only increase by 1, if at all [2]. In contrast, we show here that for closed convex codes, this increase, even if finite, can be arbitrarily large (Theorem 3.7).

We also disprove a conjecture of Goldrup and Phillipson [6] concerning the relationship between open convex and closed convex codes (Theorem 3.10). Finally, we give the first example of a code on 8 neurons that has no "local obstructions" to (open or closed) convexity, but in fact is neither open convex nor closed convex (Theorem 3.11).

The outline of our work is as follows. Section 2 provides relevant definitions and prior results. In Section 3, we prove our main results, and then we end with a discussion in Section 4.

2 Background

In this section, we recall the definitions and prior results related to convexity of neural codes (Section 2.1), simplicial complexes (Section 2.2), and sunflowers of convex sets (Section 2.3).

2.1 Neural codes and convexity

In what follows, we use the notation $[n] := \{1, 2, ..., n\}$.

Definition 2.1. A *neural code on n neurons* is a set $\mathcal{C} \subset 2^{[n]}$. Each $\sigma \in \mathcal{C}$ is a *codeword*, and σ is a *maximal codeword* of $\mathcal{C}$ if it is a maximal element of $\mathcal{C}$ with respect to inclusion.

For example, the codeword $\sigma = \{1, 3, 4\}$ indicates that neurons 1, 3, and 4 are active, while all other neurons are silent. For brevity, we will write codewords without brackets or commas; for instance, $\sigma = 134$. Also, when we list the codewords of a code, all maximal codewords will be in boldface.

Example 2.2. The following is a neural code on 6 neurons, with 12 codewords:

$$\mathcal{C} \ = \ \{\mathbf{123}, \mathbf{124}, \mathbf{135}, \mathbf{236}, 12, 13, 14, 23, 24, 1, 2, \emptyset\} \ . \tag{1}$$

The focus of this work is on open convex and closed convex codes (see Definition 2.4 below). Recall that a set $V \subset \mathbb{R}^d$ is *convex* if the line segment joining any two points in V is contained entirely within V. Also, given subsets $U_1, U_2, \ldots, U_n$ of some $\mathbb{R}^d$ and a nonempty $\sigma \subset [n]$, we use the notation $U_\sigma := \bigcap_{i \in \sigma} U_i$.

Definition 2.3. Let $\mathcal{U} = \{U_1, U_2, \ldots, U_n\}$ be a family of sets in a set $X \subset \mathbb{R}^d$ (we call X the *stimulus space*). Then $\mathrm{code}(\mathcal{U}, X)$ is the code on n neurons given by:

$$\sigma \in \mathrm{code}(\mathcal{U}, X) \iff U_\sigma \setminus \bigcup_{j \notin \sigma} U_j \neq \emptyset \ ,$$

where $U_\emptyset := X$. A code $\mathcal{C}$ on n neurons is *realized* by a family of sets $\mathcal{U} = \{U_1, U_2, \ldots, U_n\}$ in a stimulus space $X \subset \mathbb{R}^d$ if $\mathcal{C} = \mathrm{code}(\mathcal{U}, X)$. In this case, $\mathcal{U}$ is called a *realization* of $\mathcal{C}$.

Definition 2.4. A code $\mathcal{C}$ on n neurons is *open convex* (respectively, *closed convex*) if there exists a stimulus space $X \subset \mathbb{R}^d$ (for some d) and a family of open (respectively, closed) convex sets $\mathcal{U} = \{U_1, U_2, \ldots, U_n\}$ such that (1) each U_i is a subset of X, and (2) $\mathcal{C} = \mathrm{code}(\mathcal{U}, X)$. The minimum such value of d is the *open embedding dimension* (respectively, *closed embedding dimension*) of $\mathcal{C}$.

Remark 2.5. For the codes in this work, we always take the stimulus space X to be $\mathbb{R}^d$ (cf. [13, Remark 2.19]).

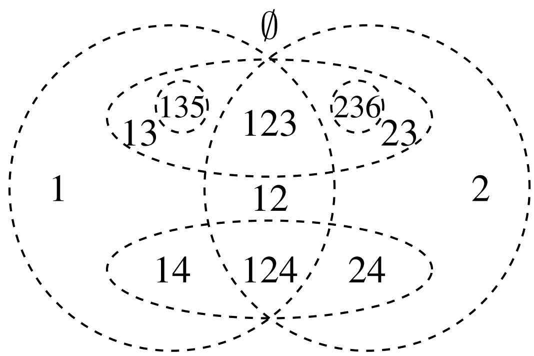

Figure 1: Open-convex realization of the code in Example 2.2.

Remark 2.6. The open embedding dimension is also called the "minimal embedding dimension" [3].

Example 2.7 (Example 2.2 continued). Consider again the code $\mathcal{C}$ in (1). First, $\mathcal{C}$ is open convex: an open-convex realization is shown in Figure 1 (more precisely, each set U_i is the interior of the union of all closures of regions labeled by some codeword containing i). Also, $\mathcal{C}$ is closed convex. Indeed, by replacing each U_i in Figure 1 by its closure, we obtain a closed-convex realization of $\mathcal{C}$.

We end this subsection with two more useful definitions.

Definition 2.8. A code $\mathcal{C}$ is *max-intersection complete* if every intersection of two or more maximal codewords is in $\mathcal{C}$. Otherwise, $\mathcal{C}$ is *max-intersection incomplete*.

If a code is max-intersection complete, then it is both open convex and closed convex [2]. The converse, however, is not true. For instance, the code $\mathcal{C}$ in (1) is open convex and closed convex (see Example 2.7), but not max-intersection complete ($135 \cap 236 = 3$ is not in $\mathcal{C}$).

Definition 2.9. Let $\mathcal{C}$ be a code on n neurons, and let $\tau \subset [n]$. The *code obtained from $\mathcal{C}$ by restricting to τ* is the neural code $\{\sigma \cap \tau \mid \sigma \in \mathcal{C}\}$.

Restricting to a set of neurons may be interpreted geometrically: if $\mathcal{U} = \{U_1, U_2, \ldots, U_n\}$ is a realization of a code $\mathcal{C}$, then $\{U_i \mid i \in \tau\}$ is a realization of the code obtained from $\mathcal{C}$ by restricting to τ.

2.2 Simplicial complexes and mandatory codewords

An (abstract) *simplicial complex* on $[n]$ is a subset of $2^{[n]}$ that is closed under taking subsets.

Definition 2.10. For a neural code $\mathcal{C}$ on n neurons, the *simplicial complex of $\mathcal{C}$* is the smallest simplicial complex containing $\mathcal{C}$:

$$\Delta(\mathcal{C}) := \{\sigma \subset [n] : \sigma \subset \alpha \text{ for some } \alpha \in \mathcal{C}\} .$$

Example 2.11 (Example 2.7 continued). The simplicial complex $\Delta(\mathcal{C})$ of the code $\mathcal{C}$ in (1) has maximal faces 123, 124, 135, and 236. The geometric realization of $\Delta(\mathcal{C})$ is shown in Figure 2.

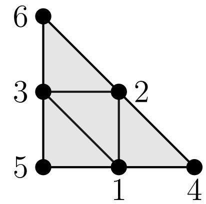

Figure 2: The simplicial complex of the code in Example 2.2.

The following result, due to Cruz *et al.* [2], states that for codes having the same simplicial complex, open-convexity is a monotone property with respect to inclusion:

Proposition 2.12 (Monotonicity property for open convex codes [2]). *Let $\mathcal{C}$ and $\mathcal{C}'$ be codes with $\mathcal{C} \subset \mathcal{C}'$ and $\Delta(\mathcal{C}) = \Delta(\mathcal{C}')$. If $\mathcal{C}$ is open convex, then $\mathcal{C}'$ is also open convex and, additionally, the open embedding dimension of $\mathcal{C}'$ is at most 1 more than that of $\mathcal{C}$.*

Definition 2.13. Let Δ be a simplicial complex on $[n]$ and let $\sigma \in \Delta$. The *link* of σ in Δ is:

$$\text{Lk}_\sigma(\Delta) := \{\tau \subset [n]\backslash\sigma : \sigma \cup \tau \in \Delta\}.$$

Recall that a *contractible* set, by definition, is homotopy-equivalent to a single point.

Definition 2.14. Let Δ be a simplicial complex. A nonempty face $\sigma \in \Delta(\mathcal{C})$ is a *mandatory codeword of* Δ if (the geometric realization of) $\text{Lk}_\sigma(\Delta)$ is non-contractible. Otherwise, σ is *non-mandatory*.

The following definition, pertaining to codes without certain "local obstructions" to convexity, is equivalent to the original definition [4].

Definition 2.15. A code $\mathcal{C}$ is *locally good* if it contains every mandatory codeword of $\Delta(\mathcal{C})$.

If a code is open convex or closed convex, then it is locally good [2, 5].

2.3 Sunflowers

A *sunflower* is a collection of sets whose pairwise intersections are all equal and nonempty. We will be interested in sunflowers that consist of convex sets, as introduced in [14]. We define sunflowers using codes as follows.

Definition 2.16. A collection $\mathcal{U} = \{U_1, U_2, \ldots, U_n\}$ of convex sets is a *sunflower* if $\text{code}(\mathcal{U}, \mathbb{R}^d)$ contains the codeword $[n]$, and all other codewords have size at most one. When $\mathcal{U}$ is a sunflower, we refer to the U_i as *petals*.

A 3-petal sunflower is shown in Figure 3.

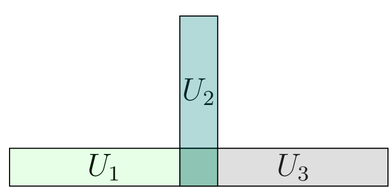

Figure 3: A sunflower $\mathcal{U} = \{U_1, U_2, U_3\}$, with $\text{code}(\mathcal{U}, \mathbb{R}^2) = \{123, 1, 2, 3, \emptyset\}$.

For our work, we will require the following theorem which constrains how the sets in a sunflower consisting of convex open sets may be arranged.

Theorem 2.17 (Sunflower Theorem [14]). *Let $\mathcal{U} = \{U_1, U_2, \ldots, U_n\}$ be a sunflower of convex open sets in $\mathbb{R}^d$, and assume that $n > d$. Then every hyperplanein $\mathbb{R}^d$ that has nonempty intersection with every U_i also has nonempty intersection with $U_{[n]}$.*

3 Results

Our main results are as follows. First, closed convex codes do *not* possess the same monotonicity property that open convex codes have (Theorem 3.2). Next, adding non-maximal codewords can increase the embedding dimension of closed convex codes by arbitrarily large, finite amounts (Theorem 3.7). We also disprove a conjecture on the relationship between open convexity and closed convexity (Theorem 3.10). Finally, we give an example of code on 8 neurons that is locally good, but neither open convex nor closed convex (Theorem 3.11), and then conjecture that there are no such codes on fewer neurons.

3.1 Closed convexity is non-monotone

Recall that open convex codes have a monontonicity property (Proposition 2.12). It is natural to ask whether the same is true for closed convexity (indeed, Cruz *et al.* did not know the answer [2, §3]):

Question 3.1. *Let $\mathcal{C}$ and $\mathcal{C}'$ be codes with $\mathcal{C} \subset \mathcal{C}'$ and $\Delta(\mathcal{C}) = \Delta(\mathcal{C}')$.*

(a) If $\mathcal{C}$ is closed convex, does it follow that $\mathcal{C}'$ is also closed convex?

(b) If $\mathcal{C}$ and $\mathcal{C}'$ are closed convex, does it follow that the closed embedding dimension of $\mathcal{C}'$ is at most 1 more than that of $\mathcal{C}$?

In a special case, Question 3.1(a) has an affirmative answer. Specifically, this is true for closed convex codes that have a realization in which the region of each codeword is top-dimensional (including max-intersection-complete codes); this result follows from results of Cruz *et al.* [2, Theorem 1.3, Lemma 2.11, and Theorem 2.12]. In general, however, Question 3.1(a) and (b) have a negative answer. We show this perhaps surprising result in the following theorem and Theorem 3.7.

Theorem 3.2 (Closed convexity is non-monotone)**.** *Consider the code*

$$\mathcal{C} = \{\mathbf{12378}, \mathbf{1457}, \mathbf{2456}, \mathbf{3468}, 17, 38, 45, 46, 2, \emptyset\}.$$

This code has a closed convex realization in $\mathbb{R}^2$, but $\mathcal{C} \cup \{278\}$ is not closed convex (in any dimension).

We first require a lemma regarding a closely related code, $\mathcal{C}_0$, which is (up to permutation of neurons) the minimally non-open-convex code of [15, Theorem 5.10] (see also [14, Theorem 4.2]). A convex set $Y \subset \mathbb{R}^d$ is *full-dimensional* if its affine hull is $\mathbb{R}^d$. Note that for a convex set, being full-dimensional is equivalent to being top-dimensional. Moreover, a convex set is full-dimensional if and only if it has nonempty interior.

Lemma 3.3. *The code $\mathcal{C}_0 = \{\mathbf{2456}, \mathbf{123}, \mathbf{145}, \mathbf{346}, 45, 46, 1, 2, 3, \emptyset\}$ is closed convex in $\mathbb{R}^2$, and every closed convex realization $\{V_1, V_2, \ldots, V_6\}$ in $\mathbb{R}^2$ is such that V_{123} is not full-dimensional.*

Proof. A closed convex realization of $\mathcal{C}_0$ in $\mathbb{R}^2$ is shown in Figure 4.

To prove the rest of the lemma, let $\{V_1, V_2, \ldots, V_6\}$ be a closed convex realization of $\mathcal{C}_0$ in $\mathbb{R}^2$. By intersecting each of the V_i's by a single sufficiently large closed ball, we may assume that each V_i is compact (cf. [13, Remark 2.19]). We will show that V_{123} is not full-dimensional. Below, we let U_i denote the interior of V_i (for $1 \leq i \leq 6$).

Suppose for contradiction that V_{123} is full-dimensional. Then V_1, V_2, and V_3 are full-dimensional. We claim that $\{U_1, U_2, U_3\}$ forms a sunflower. Indeed, 123 is the only codeword containing more than one neuron from the set $\{1, 2, 3\}$, so $V_i \cap V_j = V_{123} \neq \emptyset$ for $1 \leq i < j \leq 3$. The same relationship holds for $\{U_1, U_2, U_3\}$ because V_{123} is full-dimensional (and so has nonempty interior) and the intersection of the interiors of two sets is the interior of their intersection.

Next, no codeword contains 1234, so V_4 is disjoint from V_{123}. It follows that there exists a line L properly separating the two (compact and convex) sets. We now claim that L intersects U_i, for

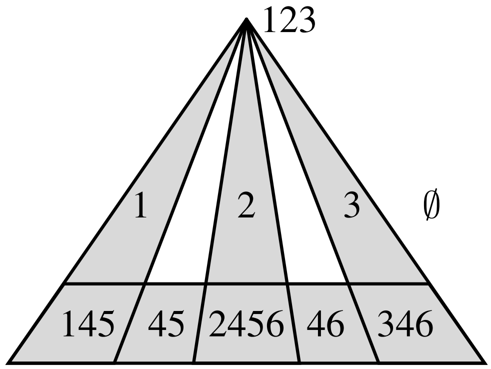

Figure 4: A closed convex realization of $\mathcal{C}_0$ in $\mathbb{R}^2$.

$1 \leq i \leq 3$. This claim follows from the fact that one side of L properly contains a point from V_{123} and the other side properly contains a point from the region corresponding to the codeword 145 (or, respectively, 2456 or 346) which is contained in V_4.

In summary, $\{U_1, U_2, U_3\}$ forms a sunflower of open, convex sets in $\mathbb{R}^2$; and the line L passes through U_1, U_2, U_3, but does not pass through their intersection (as this intersection is contained in V_{123}). These facts directly contradict Theorem 2.17, and so V_{123} is not full-dimensional. $\square$

Proof of Theorem 3.2. Notice that $\mathcal{C}$ is the same as $\mathcal{C}_0$ from Lemma 3.3, except that we have added neurons 7 and 8 which duplicate neurons 1 and 3, respectively. Thus the realization of $\mathcal{C}_0$ given in Figure 4 provides a closed realization of $\mathcal{C}$ in $\mathbb{R}^2$ by setting $V_7 = V_1$ and $V_8 = V_3$. Also, $\mathcal{C}_0$ is the restriction of $\mathcal{C}$ to the neurons $\{1, 2, \ldots, 6\}$.

Now suppose for contradiction that there is a closed convex realization $\{V_1, V_2, \ldots, V_8\}$ of $\mathcal{C} \cup \{278\}$ in $\mathbb{R}^d$. It is straightforward to check that $d = 1$ is impossible. So, assume that $d \geq 2$.

Let $p_1 \in V_{1457}$, $p_2 \in V_{278} \setminus V_{12378}$, and $p_3 \in V_{3468}$ (so, $p_i \in V_i$ and the three points are distinct). Let A be a 2-dimensional affine subspace of $\mathbb{R}^d$ containing p_1, p_2, and p_3; and let $W_i = V_i \cap A$. We claim that $\{W_1, W_2, \ldots, W_8\}$ is a realization of $\mathcal{C} \cup \{278\}$ in A (i.e. in $\mathbb{R}^2$), and moreover that W_{123} is full-dimensional in this realization.

Clearly the code of $\{W_1, W_2, \ldots, W_8\}$ is contained in $\mathcal{C} \cup \{278\}$ since the V_i's realize that code. So, we must show that every codeword from $\mathcal{C} \cup \{278\}$ arises inside A. By choice of p_1, p_2, and p_3, A contains points that realize the codewords 1457, 3468, and 278.

Consider the line segment L_1 from p_2 to p_3. This line segment is contained entirely in W_8, and so the codewords that appear along it must come from the set $\{278, 12378, 38, 3468\}$. In fact, each of these codewords must appear, and in exactly this order, since the code along the line segment must be a 1-dimensional code (see the arguments in [16]). A symmetric argument shows that the line segment L_2 from p_2 to p_1 has the codewords $\{278, 12378, 17, 1457\}$ along it in that order. Finally, a similar argument shows that the line segment L_3 from p_1 to p_3 has along it the codewords $\{1457, 45, 2456, 46, 3468\}$ in that order.

Thus, only the codewords 2 and $\emptyset$ need to be shown to arise in A. The codeword 2 can be recovered by examining a line segment from p_2 to a point in W_{2456}, and $\emptyset$ can be obtained by assuming that the W_i are bounded (cf. [13, Remark 2.19]).

To see that W_{123} is full-dimensional in A, we again consider the line segments L_1, L_2, and L_3. The points p_1, p_2, and p_3 must be in general position: the codeword 1457 that p_1 gives rise to does not appear on the line segment L_1 between p_2 and p_3, the codeword 3468 corresponding to p_3 does not appear on L_2, and the codeword 278 corresponding to p_2 does not appear on L_3. Thus p_1, p_2, p_3 define a triangle in $\mathbb{R}^2$ with edges L_1, L_2, L_3; see Figure 5.

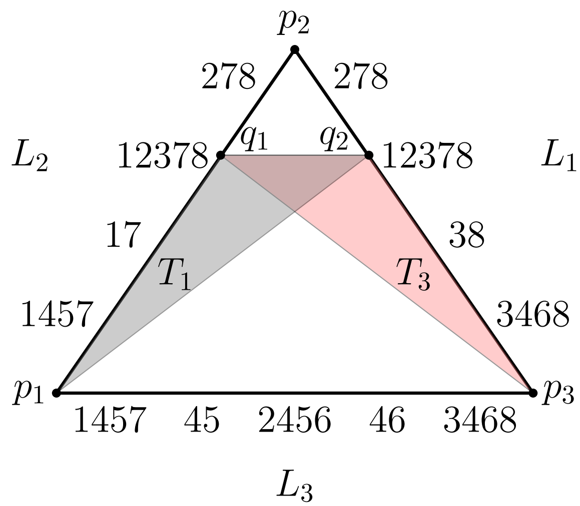

Figure 5: The triangle with vertices p_1, p_2, and p_3. Also depicted are the codewords appearing along each edge. The points q_1, q_2 and triangles T_1, T_2 are as in the proof of Theorem 3.2.

Next, L_1 and L_2 both pass through W_{12378} and intersect only at p_2, so we now choose distinct points q_1 and q_2 in $L_1 \cap W_{12378}$ and $L_2 \cap W_{12378}$, respectively. Now consider the triangles T_1 and T_3 with respective vertex sets $\{q_1, q_2, p_1\}$ and $\{q_1, q_2, p_3\}$ (see the figure). The vertices of T_1 are contained in W_1, so $T_1 \subset W_1$. Similarly, $T_3 \subset W_3$. Hence, $T_1 \cap T_3 \subset W_1 \cap W_3$. The intersection $T_1 \cap T_2$ is full-dimensional (the doubly shaded region in Figure 5), and therefore so is $W_1 \cap W_3$.

However, $W_1 \cap W_3 = W_{123} = W_{12378}$ (because only the codeword 12378 contains both neurons 1 and 3). So, by deleting the sets W_7 and W_8, we obtain a closed convex realization $\{W_1, W_2, \ldots, W_6\}$ of the code $\mathcal{C}_0$ in $A \cong \mathbb{R}^2$ with W_{123} full-dimensional in A. This contradicts Lemma 3.3, and so the proof is complete. □

Remark 3.4. Theorem 3.2 answered Question 3.1 in the negative using a code on 8 neurons and codewords of size up to 5. We do not know whether such a result is possible using 7 or fewer neurons and/or codewords of size at most 4.

Remark 3.5. Previous works such as [2, Lemma 2.9] and [6, Theorem 4.1] have used minimum-distance arguments to prove that certain codes are not closed convex. Our proof of Theorem 3.2 took a different approach, effectively reducing the argument to the case of open sets. In the future, we would like a general set of criteria that preclude closed convexity, and which prove, as special cases, that the code $\mathcal{C}$ of Theorem 3.2 and the relevant codes in [2, 6] are not closed convex.

3.2 Arbitrarily large increases in closed embedding dimension

In Theorem 3.2, we saw that adding a non-maximal codeword to a closed convex code may yield a non-closed-convex code. Nevertheless, when the resulting code is closed convex, one might hope that its closed embedding dimension has not greatly increased, in line with the fact that open embedding dimension increases by at most 1 when a non-maximal codeword is added (recall Proposition 2.12).

However, this is not the case. In fact, adding a non-maximal codeword may increase the closed embedding dimension by any amount, as we show in the next theorem.

Definition 3.6. For $n \geq 2$, let $\mathcal{A}_n$ be the code whose neurons are $\{1, 2, \ldots, n+1\}$ and $\{\overline{1}, \overline{2}, \ldots, \overline{n}\}$ and which consists of the following $2n+3$ codewords:

(i) The following three codewords: $\{1, 2, \ldots, n, \overline{1}, \overline{2}, \ldots, \overline{n}\}$, $\{n+1\}$, and the empty set,

(ii) The codeword $\{i, \bar{i}, n+1\}$ for all $i \in [n]$, and

(iii) The codeword $\{i, \bar{i}\}$ for all $i \in [n]$.

For $i \in [n]$, note that the neuron i appears in a codeword of $\mathcal{A}_n$ if and only if $\bar{i}$ appears. Thus the receptive fields of neurons i and $\bar{i}$ will be identical in every realization of $\mathcal{A}_n$ (i.e., $U_i = U_{\bar{i}}$).

Theorem 3.7 (Large increase in closed embedding dimension). *For $n \geq 2$, the code $\mathcal{A}_n$ has a closed convex realization in $\mathbb{R}^2$, and the code $\mathcal{A}_n \cup \{\{\bar{1}, \bar{2}, \ldots, \bar{n}\}\}$ is closed convex with closed embedding dimension equal to n.*

To prove Theorem 3.7 we first require a lemma similar to Lemma 3.3.

Lemma 3.8. *Assume $n \geq 2$. Let $\mathcal{S}_n$ be the code obtained from $\mathcal{A}_n$ by restricting to the neurons $\{1, 2, \ldots, n+1\}$. Let $\{V_1, V_2, \ldots, V_{n+1}\}$ be a closed convex realization of $\mathcal{S}_n$ in $\mathbb{R}^d$. If $d < n$, then the region $V_{[n]}$ is not full-dimensional.*

Proof. By intersecting each of the V_i's by a single sufficiently large closed ball, we may assume that each V_i is compact (cf. [13, Remark 2.19]). Suppose for contradiction that $V_{[n]}$ is full-dimensional. Then the sets $V_1, V_2, \ldots, V_n$ are also full-dimensional.

Next, from Definition 3.6, we see that $\mathcal{S}_n = \{[n], \{1, n+1\}, \{2, n+1\}, \ldots, \{n, n+1\}, \{1\}, \{2\}, \ldots, \{n+1\}, \emptyset\}$. So, $\{V_1, V_2, \ldots, V_n\}$ forms a sunflower, that is, $V_i \cap V_j = V_{[n]} \neq \emptyset$ for all $1 \leq i < j \leq n$. We claim that the interiors $\{U_1, U_2, \ldots, U_n\}$ of the V_i also form a sunflower. Indeed, for all $1 \leq i < j \leq n$, we see that $U_i \cap U_j$ is the interior of $V_i \cap V_j = V_{[n]}$, which is nonempty because $V_{[n]}$ is full-dimensional.

Next, observe that $V_{[n]}$ is disjoint from V_{n+1}, because $[n+1]$ is not a codeword of $\mathcal{S}_n$. As both $V_{[n]}$ and V_{n+1} are compact and convex, there exists a hyperplane H properly separating the two sets. For all $i \in [n]$, the set V_i intersects V_{n+1} (because $\{i, n+1\} \in \mathcal{S}_n$) and so each side of H properly contains a point in V_i. Thus, H passes through the interior of V_i for all $i \in [n]$, but does not pass through their common intersection $U_{[n]}$. When $d < n$, this contradicts Theorem 2.17. □

We are now ready to prove the main result of this subsection.

Proof of Theorem 3.7. A closed convex realization $\{W_1, W_2, \ldots, W_{n+1}, W_{\bar{1}}, W_{\bar{2}}, \ldots, W_{\bar{n}}, \}$ of $\mathcal{A}_n$ in $\mathbb{R}^2$ is shown in Figure 6: n line segments meet at a common point, and an additional line segment crosses all the line segments away from the common point.

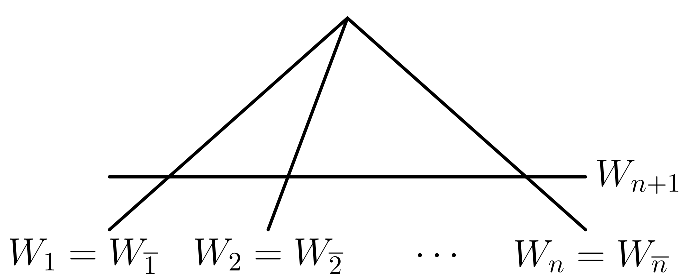

Figure 6: A closed convex realization of $\mathcal{A}_n$ in $\mathbb{R}^2$.

Our next task is to construct a closed convex realization of $\mathcal{C}_n := \mathcal{A}_n \cup \{\{\bar{1}, \bar{2}, \ldots, \bar{n}\}\}$ in $\mathbb{R}^n$. An informal description of this construction is as follows. Starting from the realization of $\mathcal{A}_n$ in Figure 6, rotate each $W_i = W_{\bar{i}}$ so it lies along the i-th coordinate axis in $\mathbb{R}^n$ and then fatten it into an n-dimensional rectangular prism, so that the common intersection is a unit n-cube at the origin. Next, W_{n+1} becomes a thickened hyperplane that meets each of the $W_i = W_{\bar{i}}$'s. So far, we have a realization of $\mathcal{A}_n$, and now we obtain the new codeword $\{\overline{1}, \overline{2}, \ldots, \overline{n}\}$ by "slicing off" a corner of the n-cube from each of $W_1, W_2, \ldots, W_n$. This construction is shown for $n = 2$ in Figure 7.

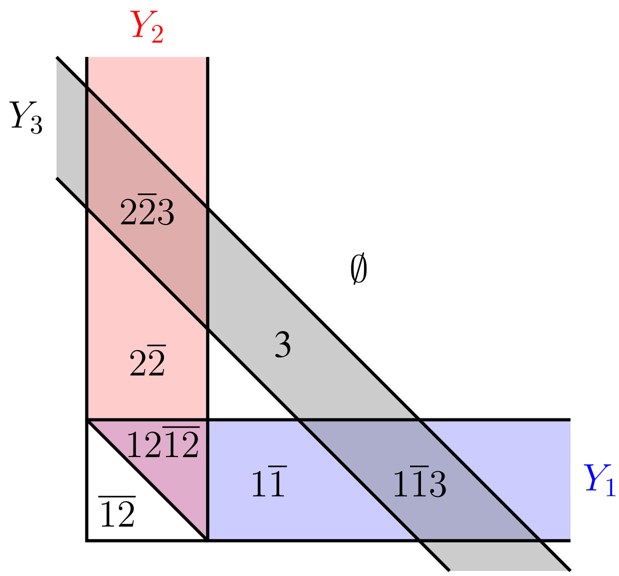

Figure 7: A closed convex realization of $\mathcal{C}_2$ in $\mathbb{R}^2$. The receptive fields Y_1, Y_2, and Y_3 are labeled; and $Y_{\overline{1}}$ and $Y_{\overline{2}}$ are rectangles (as indicated by the codewords).

More precisely, the realization $\{Y_1, Y_2, \ldots, Y_{n+1},\ Y_{\overline{1}}, Y_{\overline{2}}, \ldots, Y_{\overline{n}}\}$ of $\mathcal{C}_n$ in $\mathbb{R}^n$ is given by, for $i \in [n]$, $Y_{\overline{i}} := \{x \in \mathbb{R}^n_{\geq 0} \mid 0 \leq x_j \leq 1 \text{ for all } j \neq i\}$ and $Y_i := Y_{\overline{i}} \cap \{x \in \mathbb{R}^n \mid x_1 + x_2 + \ldots x_n \geq 1\}$; and $Y_{n+1} := \{x \in \mathbb{R}^n \mid 2n \leq x_1 + x_2 + \cdots + x_n \leq 2n + 1\}$. In this realization, the codeword $\{\overline{1}, \overline{2}, \ldots, \overline{n}\}$ arises in the region of the unit n-cube where the sum of all coordinates is less than 1. In the remainder of the realization, we see that $i \in [n]$ appears if and only if $\overline{i}$ appears. Moreover, the receptive fields Y_i, for $i \in [n]$, form a sunflower whose petals meet in the subset of the n-cube where the sum of coordinates is greater than or equal to 1, and so the only codewords arising involving i are $\{i, \overline{i}\}$, $\{1, 2, \ldots, n, \overline{1}, \overline{2}, \ldots, \overline{n}\}$, and the codeword $\{i, \overline{i}, n+1\}$, the latter arising where the thickened hyperplane Y_{n+1} meets the receptive field of i. Finally, the codeword $\{n+1\}$ arises anywhere in the thickened hyperplane Y_{n+1} where it does not meet any of the Y_i's, for example, at a point where one of the coordinates is negative.

Now it remains only to show that there is no closed convex realization of $\mathcal{C}_n$ in $\mathbb{R}^{n-1}$. Suppose for contradiction that we have such a realization $\{V_1, V_2, \ldots, V_{n+1},\ V_{\overline{1}}, V_{\overline{2}}, \ldots, V_{\overline{n}}\}$. Choose a point p^* in the region that gives rise to the codeword $\{\overline{1}, \overline{2}, \ldots, \overline{n}\}$, and for $i \in [n]$ choose a point p_i in $V_i \cap V_{n+1}$ (i.e., p_i lies in the region that gives rise to the codeword $\{i, \overline{i}, n+1\}$).

For $i \in [n]$, let L_i denote the line segment from p^* to p_i. Observe that the containment $L_i \subset V_{\overline{i}}$ holds (as both endpoints of L are in $V_{\overline{i}}$). Also, the codewords $\{\overline{1}, \overline{2}, \ldots, \overline{n}\}$, $\{1, 2, \ldots, n, \overline{1}, \overline{2}, \ldots, \overline{n}\}$, $\{i, \overline{i}\}$, and $\{i, \overline{i}, n+1\}$ appear along L_i in precisely that order. Also, the codeword $\{n+1\}$ must arise along any line segment between distinct p_i. Thus, all codewords of $\mathcal{C}_n$ arise inside the affine hull of $\{p^*, p_1, \ldots, p_n\}$. We denote this affine hull by A.

It follows that by replacing our receptive fields $\{V_1, V_2, \ldots, V_{n+1},\ V_{\overline{1}}, V_{\overline{2}}, \ldots, V_{\overline{n}}\}$ by their intersections with A, we obtain a closed convex realization of $\mathcal{C}_n$ inside $A \cong \mathbb{R}^d$, for some $d \leq n-1$, such that the convex hull of the points $\{p^*,\ p_1, p_2, \ldots, p_n\}$ is full-dimensional in A (by construction). Observe that $d \geq 2$, as $\mathcal{C}_n$ is not convex in $\mathbb{R}^1$; one reason for this is that $\Delta(\mathcal{C}_n)$ has a 1-dimensional hole.

The code $\mathcal{C}_n$ is invariant under permutations of $[n]$ provided we also apply the permutation to $\{\overline{1}, \overline{2}, \ldots, \overline{n}\}$, and so we may assume without loss of generality that $\{p^*,\ p_1, p_2, \ldots, p_d\}$ form the vertices of a d-simplex Δ in A.

For $i \in [d]$, each L_i is a distinct edge of Δ, and we let q_i be a point on L_i where the codeword $\{1, 2, \ldots, n,\ \overline{1}, \overline{2}, \ldots, \overline{n}\}$ arises. In particular, $p_i \neq q_i \in V_{[n]}$. Since the q_i lie on distinct edges of Δ, the affine hull H of $\{q_1, q_2, \ldots, q_d\}$ has dimension $d-1$ and so is a hyperplane $H \subset A$. We orient H so that its negative side contains p^* and hence its positive side contains $\{p_1, p_2, \ldots, p_d\}$.

For $i \in [d]$, let Δ_i be the d-simplex with vertices $\{q_1, q_2, \ldots, q_d,\ p_i\}$, and observe that $\Delta_i \subset V_i$. Since all Δ_i lie on the nonnegative side of H and share the common face whose vertices are $\{q_1, q_2, \ldots, q_d\}$, we may choose a point q^* that lies in the interior of all Δ_i, and hence in $V_{[d]}$. Since $d \geq 2$ and $\{V_1, V_2, \ldots, V_n\}$ is a sunflower, we have $V_{[d]} = V_{[n]}$. Thus, q^* lies in $V_{[n]}$. The point q^* lies strictly on the positive side of H, and so the convex hull of $\{q^*,\ q_1, q_2, \ldots, q_d\}$ is a d-simplex contained in $V_{[n]}$. Therefore, $V_{[n]}$ is full-dimensional in A. Since $d < n$, this contradicts Lemma 3.8. □

3.3 A counterexample to a conjecture of Goldrup and Phillipson

Recall that Theorem 3.2 contrasts open convex codes with closed convex codes by showing that the latter does not possess the same monotonicity property as the former. In the same spirit, there is much interest in comparing and relating open convexity to closed convexity. In particular, we would like to know properties that cause codes to be solely open convex or solely closed convex. In an attempt to distinguish codes that are open convex but not closed convex, Goldrup and Phillipson posed the following conjecture [6, Conjecture 4.3].

Conjecture 3.9. *Let $\mathcal{C}$ be a max-intersection incomplete open convex code, where $\Delta(\mathcal{C})$ has at least two non-mandatory codewords not contained in $\mathcal{C}$. Suppose $\mathcal{C}$ has at least three maximal codewords M_1, M_2, M_3, and there is $\sigma \subset M_1$ with $\sigma \in \mathcal{C}$ such that $\sigma \cap M_2 \notin \mathcal{C}$. Then $\mathcal{C}$ is not closed convex.*

We disprove Conjecture 3.9 through a counterexample, namely, the code from Example 2.2.

Theorem 3.10. *The neural code $\mathcal{C} = \{\mathbf{123}, \mathbf{124}, \mathbf{135}, \mathbf{236}, 12, 13, 14, 23, 24, 1, 2, \emptyset\}$ fulfills the hypotheses of Conjecture 3.9 and is closed convex.*

Proof. We begin by checking that $\mathcal{C}$ satisfies the hypotheses of Conjecture 3.1. First, we saw that $\mathcal{C}$ is open convex (Example 2.2). Next, $\mathcal{C}$ is max-intersection incomplete, as the intersection of maximal codewords $135 \cap 236 = 3$ is not in $\mathcal{C}$.

We must also show that $\Delta(\mathcal{C})$ has at least two non-mandatory codewords that are not in $\mathcal{C}$. It is straightforward to check that the links $\mathrm{Lk}_{\{3\}}(\Delta(\mathcal{C}))$ and $\mathrm{Lk}_{\{4\}}(\Delta(\mathcal{C}))$ are the following contractible simplicial complexes (respectively):

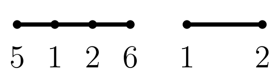

Therefore, 3 and 4 are non-mandatory codewords. Also, neither 3 nor 4 is in $\mathcal{C}$.

Next, we must show that $\mathcal{C}$ has three maximal codewords M_1, M_2, M_3 and a codeword $\sigma \in \mathcal{C}$ such that $\sigma \subset M_1$ and $\sigma \cap M_2 \notin \mathcal{C}$. Let $M_1 = 123, M_2 = 236, M_3 = 124$, and let $\sigma = 13 \in \mathcal{C}$. Then $13 = \sigma \subset M_1 = 123$. Also, $13 \cap 236 = \sigma \cap M_2 = 3 \notin \mathcal{C}$.

Finally, we already saw (in Example 2.2) that $\mathcal{C}$ is closed convex. □

Although Conjecture 3.9 is false, we nevertheless still wish to discover conditions guaranteeing that a code is open convex but not closed convex – or vice-versa.

3.4 A locally good code that is neither open convex nor closed convex

Next, we consider codes that are neither open convex nor closed convex. Based on the examples in this work and in other articles, one might wonder whether every locally good codes is open convex or closed convex. However, here we present a code on 8 neurons that, despite being locally good, is neither open convex nor closed convex (Theorem 3.11). As seen in the proof, this code is built by combining two locally good codes, one that is not closed convex and the other not open convex.

Theorem 3.11. *The following code is locally good, but neither open convex nor closed convex:*

$$\mathcal{C} = \{\mathbf{2345}, \mathbf{123}, \mathbf{124}, \mathbf{145}, 12, 14, 23, 24, 45, 2, 4, \emptyset\} \cup \{\mathbf{237}, \mathbf{238}, \mathbf{367}, \mathbf{678}, 26, 37, 67, 6, 8\}.$$

Proof. By [4, Theorem 1.3 and Lemma 1.4], being locally good is equivalent to the following: If $\emptyset \neq \sigma \notin \mathcal{C}$ and σ is the intersection of two or more maximal codewords of $\mathcal{C}$, then $\mathrm{Lk}_\sigma(\Delta(\mathcal{C}))$ is contractible. It is straightforward to check that only 1, 3, and 7 are nonempty intersections of maximal codewords and not in $\mathcal{C}$. Their links in $\Delta(\mathcal{C})$, respectively, are shown below:

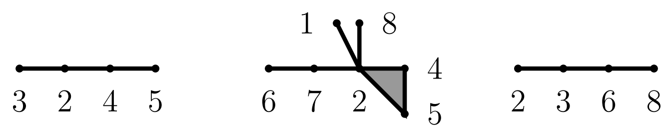

Each link is contractible, and so $\mathcal{C}$ is locally good.

Next, we show that $\mathcal{C}$ is neither open convex nor closed convex. Let $\mathcal{U} = \{U_1, U_2, \ldots, U_8\}$ be a realization of $\mathcal{C}$ in some $\mathbb{R}^d$. We must show that some U_i is not open, and some U_j is not closed.

First, $\{U_1, U_2, U_3, U_4, U_5\}$ is a realization of the restriction of $\mathcal{C}$ to the neurons $\{1, 2, 3, 4, 5\}$, which is the code called "$C4$" in [6, Table 1] and is non-open-convex [9]. Thus, at least one of U_1, U_2, U_3, U_4, U_5 is not open.

Similarly, $\{U_2, U_3, U_6, U_7, U_8\}$ realizes $\mathcal{C}$ restricted to $\{2, 3, 6, 7, 8\}$. After relabeling neurons $2, 3, 6, 7, 8$ by $1, 3, 2, 4, 5$, respectively, this restricted code is the code "$C10$" in [6, Table 1], which is non-closed-convex [6, Theorem 4.1]. Hence, at least one of U_2, U_3, U_6, U_7, U_8 is not closed. □

Remark 3.12. Another locally good code on 8 neurons that is neither open convex nor closed convex, is the code $\mathcal{C} \cup \{278\}$ from Theorem 3.2. Non-closed-convexity is shown in that theorem. As non-open-convexity, restricting the code to $\{1, 2, 3, 4, 5, 6\}$ yields (up to permuting neurons) the minimally non-open-convex code in [15, Theorem 5.10] (this is the code in Lemma 3.3), and restriction preserves convexity.

The code in Theorem 3.11 is on 8 neurons. We want to know whether there is a code with the same properties but on fewer neurons. (For instance, to our knowledge, the codes that Jeffs and Novik show are locally good – in fact, "locally perfect" – but neither open convex nor closed convex, require at least 8 neurons [17, §9].)

Conjecture 3.13. *Every locally good code on at most 7 neurons is open convex or closed convex.*

For codes on up to 4 neurons, Conjecture 3.13 is true, as such locally good codes are open convex [4]. For codes on 5 neurons, most of the work toward resolving the conjecture was done by Goldrup and Phillipson [6]. The only codes left to analyze are those with the same simplicial complex as the code in [9, Theorem 3.1] (this is the code "C4" in the proof of Theorem 3.11).

Another approach to resolving Conjecture 3.13 comes from the fact that the set of all neural codes forms a partially ordered set (poset), which arises from surjective morphisms as defined in [15]. In this poset, the open convex codes form a down-set, that is, all codes lying below an open convex code are also open convex. Also forming down-sets are closed-convex codes [11, Proposition 9.3] and locally good codes [18, Corollary 4.2]. Therefore, it would be interesting to check whether any codes lying below the 8-neuron code in Theorem 3.11 are neither open convex nor closed convex (i.e., whether or not this code is minimal among codes that are neither open nor closed convex). If trying this approach, however, one should beware that it is possible for the number of neurons in a code to *increase* while moving downwards in this poset.

4 Discussion

Open convex and closed convex codes share several important properties. For instance, both classes of codes are locally good (and, in fact, "locally perfect" [19]). Also, max-intersection complete codes are both open convex and closed convex [2]. However, while open convex codes possess a monotonicity property, which greatly simplifies the analysis of all codes with a given simplicial complex, here we showed that this property fails for closed convex codes (Theorem 3.2). Also, even when monotonicity holds, the embedding dimension can greatly increase (Theorem 3.7).

Additional results in our work also address fundamental questions pertaining to open convex and closed convex codes. For instance, there is a locally good code on 8 neurons that is neither open convex nor closed convex (Theorem 3.11).

Our results lead to several open questions. First, is there an instance of non-monotonicity in codes on up to 7 neurons and/or codewords with size up to 5 (Remark 3.4)? Next, is there a locally good code on 7 neurons that is neither open convex nor closed convex (Conjecture 3.13)? Also, are there general criteria (beyond local obstructions) for ruling out closed convexity (Remark 3.5)? This is an important future direction, as existing approaches are somewhat ad-hoc, and progress here will therefore aid in classifying closed convex codes.

Answers to these questions, together with the results we already have on convex codes, will clarify the theories of open convex and closed convex codes. Specifically, we will better understand when the convexity properties are the same (for instance, for nondegenerate codes [2]) and when they differ (as seen in this work). In turn, this knowledge contributes to answering the questions from neuroscience that originally motivated our work. Specifically, we will better understand what types of neural codes allow the brain to represent structured environments.

Acknowledgements

The authors thank Aaron Chen, Chad Giusti, Kaitlyn Phillipson, and Thomas Yahl for insightful comments and suggestions, and thank Nida Obatake for editorial suggestions. The authors are also grateful to four referees for detailed comments which improved this work. BG, SM, and AS are grateful to Ramona Heuing and Martina Juhnke-Kubitzke for finding a crucial error in a prior version of this work. BG and SM initiated this research in the 2019 REU in the Department of Mathematics at Texas A&M University, supported by the NSF (DMS-1757872). AS was supported by the NSF (DMS-1752672). RAJ was supported by the NSF (DGE-1761124).

References

[1] John O'Keefe and Jonathan Dostrovsky. The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res.*, 34(1):171–175, 1971.

[2] Joshua Cruz, Chad Giusti, Vladimir Itskov, and Bill Kronholm. On open and closed convex codes. *Discrete Comput. Geom.*, 61(2):247–270, 2019.

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

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

[5] Chad Giusti and Vladimir Itskov. A no-go theorem for one-layer feedforward networks. *Neural Comput.*, 26(11):2527–2540, 2014.

[6] Sarah Ayman Goldrup and Kaitlyn Phillipson. Classification of open and closed convex codes on five neurons. *Adv. Appl. Math.*, 112:101948, January 2020.

[7] Megan Franke and Samuel Muthiah. Every binary code can be realized by convex sets. *Adv. Appl. Math.*, 99:83–93, 2018.

[8] Vladimir Itskov, Alexander Kunin, and Zvi Rosen. Hyperplane neural codes and the polar complex. In Nils A. Baas, Gunnar E. Carlsson, Gereon Quick, Markus Szymik, and Marius Thaule, editors, *Topological Data Analysis*, pages 343–369. Springer International Publishing, 2020.

[9] Caitlin Lienkaemper, Anne Shiu, and Zev Woodstock. Obstructions to convexity in neural codes. *Adv. Appl. Math.*, 85:31–59, 2017.

[10] Robert Williams. Strongly maximal intersection-complete neural codes on grids are convex. *Appl. Math. Comput.*, 336:162–175, 2018.

[11] R. Amzi Jeffs. Embedding Dimension Phenomena in Intersection Complete Codes. Available at `arXiv:1909.13406`, 2019.

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

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

[14] R. Amzi Jeffs. Sunflowers of convex open sets. *Adv. Appl. Math.*, 111:101935, 2019.

[15] R. Amzi Jeffs. Morphisms of neural codes. *SIAM Journal on Applied Algebra and Geometry*, 4:99–122, 2020.

[16] Zvi Rosen and Yan X Zhang. Convex neural codes in dimension 1. Available at `arXiv:1702.06907`, 2017.

[17] R. Amzi Jeffs and Isabella Novik. Convex union representability and convex codes. *Int. Math. Res. Notices*, Apr 2019.

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

[19] 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, 2019.