SPARSE-SKY AND HUTCH

OPEN AND CLOSED CONVEXITY IN 3-SPARSE NEURAL CODES AND BEYOND

Brianna Gambacini
Department of Mathematics
University of Connecticut
Storrs, CT 06269
brianna.gambacini@uconn.edu

Sam Macdonald
Department of Mathematics
Willamette University
Salem, OR 97301
smmacdonald@willamette.edu

July 25, 2019

ABSTRACT

Neural codes are mathematical models of neural activity. Neuroscientists have discovered neurons called place cells which fire when animals are in specific (and usually convex) regions in space. Through monitoring these place cells and recording data on when they fire, we can construct neural codes, which tell us which neurons fire together. Of particular interest to the mathematical community is identifying which codes can be represented by open or closed convex sets. In this paper, we provide counterexamples for two conjectures regarding closed convex neural codes along with several other results, including a new method for determining whether or not a code is open convex.

***K*eywords** Neural Coding · Open · Closed · Convexity · Sparsity · Local Obstruction · Place cells

1 Introduction

In recent years, mathematicians have become fascinated with neural codes corresponding to specific types of neurons called place cells. Place cells, which were discovered by John O'Keefe in 1971 and won him a joint Nobel Prize in Physiology or Medicine in 2014, are neurons in animals that fire when the animal is in a specific location. This process creates a mental map of the region in the brain. Research has shown that the locations that can be mapped in this way are almost always convex, and neural codes are derived from the intersections of these regions. This has prompted mathematicians to consider the following question: given a neural code of some finite number of neurons, is it possible for this code to correspond to a location (or set of locations) that can be mapped by the brain using place cells? Specifically, can these regions be constructed using only open convex sets? Closed convex sets?

Within the context of this paper, we consider some specific questions: How are open convex and closed convex codes related? What properties must a code have in order to be open convex but not closed convex? What about codes that are closed but not open convex? And finally, is there a faster way to show that a neural code is open convex than simply drawing the entire diagram - potentially in a high-dimensional space? This paper addresses each of these questions while proving several smaller results along the way.

In Section 2, we provide the definitions, notation, and previous results that are pertinent to this paper. In Section 3, we discuss the differences between open and closed convexity as well as disprove two conjectures regarding closed convexity. In Section 4, we examine the dimensionality of place field diagrams and their interwoven relationship with open convexity. In Section 5, we present a method for proving open convexity of a neural code given a smaller, reduced form of the code. In Section 6 we provide some additional results that were proven along the way, so that future researchers have some extra tools in their tool-belt when they investigate neural codes. We additionally provide a few conjectures in an effort to provide potential directions for future research.

2 Background

In order to set up the problems, we first provide some relevant definitions and notation.

2.1 Neural Codes and Convexity

Definition 2.1. A *codeword* σ on n neurons is a subset of the population of neurons $[n] = \{1, 2, ..., n\}$.

For example, $\sigma = \{1, 3, 4\}$ indicates that neurons 1, 3, and 4 are active, while all other neurons are silent. Throughout this paper, we will write codewords without brackets or commas for the sake of brevity. So $\sigma = 134$ will be used in place of $\sigma = \{1, 3, 4\}$.

Definition 2.2. A set of codewords is known as a *neural code* $\mathcal{C}$. For codewords $\sigma \subseteq [n]$, we have $\mathcal{C} \subseteq 2^{[n]}$.

Definition 2.3. A code $\mathcal{C}$ is n*-sparse* if no codeword $\sigma \in \mathcal{C}$ contains more than n neurons.

The problems addressed in this paper focus mainly on determining which neural codes are convex, as well as identifying differences between open and closed convex codes. The definitions of open and closed convexity are below.

Definition 2.4. A *convex set* $\mathcal{U} \subseteq \mathbb{R}^n$ is a set of points such that, given any two points $x, y \in \mathcal{U}$, the line joining them lies entirely within that set. Algebraically, a set S is convex if for all $x, y \in S$, $\lambda x + (1 - \lambda)y \in S$ for all $\lambda \in [0, 1]$.

Definition 2.5. A set $\mathcal{U}$ is *open* if for all $x \in \mathcal{U}$ there exists some neighborhood of x that lies entirely within $\mathcal{U}$. A set $\mathcal{U}$ is *closed* if it is the compliment of an open set.

Definition 2.6. For a word $\sigma \in C$, we define the set $\mathcal{U}_\sigma = \cap_{i \in \sigma} \mathcal{U}_i$.

Definition 2.7. A code $\mathcal{C} \subset 2^{[n]}$ is *open (respectively, closed) convex* if there exist open (respectively, closed) subsets $\mathcal{U}_1, \mathcal{U}_2, \ldots \mathcal{U}_n \subseteq \mathbb{R}^d$, for some d, that realize the code.

Example 2.8. Consider the neural code $\mathcal{C} = \{1, 2, 3, 12, 13, 23, 123, \emptyset\}$. This code can be expressed using both exclusively open and exclusively closed sets in the form of a traditional three circle Venn-diagram, as seen in Figure 1.

2.2 Facets and Intersection Completeness

Other important definitions refer to maximal elements of a code and their intersections.

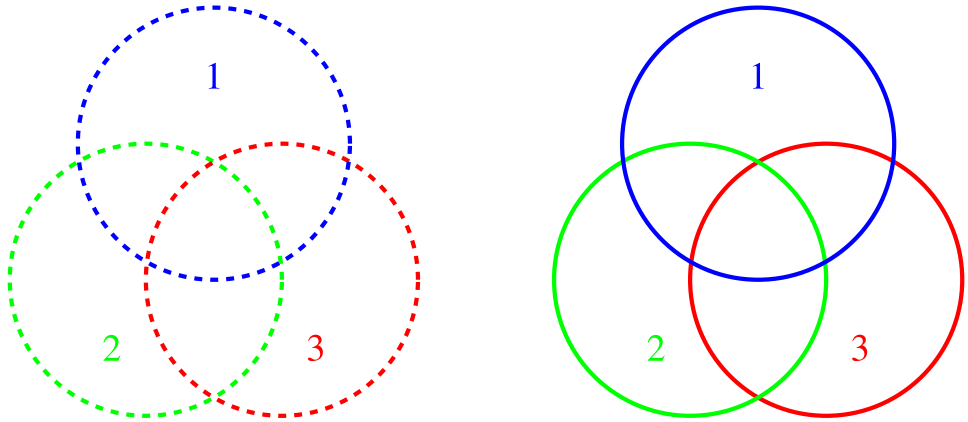

Figure 1: Open and closed convex realizations of the neural code $\mathcal{C}$, respectively.

Definition 2.9. A codeword $\sigma \in \mathcal{C}$ is a *facet* of $\mathcal{C}$ if it is a maximal element of $\mathcal{C}$ with respect to inclusion, that is, $\sigma \nsubseteq \alpha$ for all $\alpha \in \mathcal{C}$ such that $\alpha \neq \sigma$.

In this paper, whenever we list a code C, all facets of C will be bolded.

Definition 2.10. A code $\mathcal{C}$ is *max-intersection complete* if all the maximal intersections of its facets are in $\mathcal{C}$. If a code does not contain all of its facets' intersections then it is *max-intersection incomplete*.

It was proved in [1] that if a code $\mathcal{C}$ is max-intersection complete, then it is both open and closed convex.

2.3 Simplicial Complexes and Local Obstructions

Definition 2.11. A *simplicial complex* is a subset of $2^{[n]}$ that is closed under taking subsets, where $[n] := \{1, 2, ..., n\}$ is the population of neurons.

Definition 2.12. More specifically, for a neural code $\mathcal{C}$ on n neurons, the *simplicial complex* $\Delta(\mathcal{C})$ is defined as:

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

Definition 2.13. Let Δ be a simplicial complex on n neurons and $\sigma \in \Delta$. Then the *link* of σ in Δ is:

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

Definition 2.14. A set is *contractible* if it can be continuously deformed to a single point within the set.

Definition 2.15. A word $\sigma \in \Delta(\mathcal{C})$ is *mandatory* if $\text{Lk}_\sigma(\Delta(\mathcal{C}))$ is not contractible. If $\text{Lk}_\sigma(\Delta(\mathcal{C}))$ is contractible, then σ is said to be *non-mandatory*.

Definition 2.16. A code $\mathcal{C}$ with $\sigma \in \Delta(\mathcal{C})$ has a *local obstruction* at σ if σ is mandatory and not in $\mathcal{C}$. If $\mathcal{C}$ contains all of its mandatory codewords, then $\mathcal{C}$ is *locally good*.

It was proved in [2] and [3] that if $\mathcal{C}$ is convex, then it is locally good.

We introduce the following definition.

Definition 2.17. For a 3-sparse neural code, the *reduced* sub-code of $\mathcal{C}$, denoted $\mathcal{C}_{red}$, is the code containing all size-three codewords of $\mathcal{C}$ and their subsets that are also in $\mathcal{C}$.

3 Results on Conjectures Regarding Closed Convexity

3.1 The Goldrup and Phillipson Conjecture

Recently, [4] posed the following conjecture as an attempt to distinguish codes that are open convex but not closed convex.

Conjecture 3.1. *[4] Let $\mathcal{C}$ be an open convex, max-intersection incomplete code with at least two non-mandatory codewords. Suppose $\mathcal{C}$ has at least 3 facets 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 show the conjecture is false by providing a counter-example.

Theorem 3.2. *The neural code $\mathcal{C} = \{\emptyset, 1, 2, 12, 13, 14, 23, 24, \mathbf{123}, \mathbf{124}, \mathbf{135}, \mathbf{236}\}$ fulfills the hypotheses of Conjecture 3.1 and is closed convex.*

Proof. We begin by checking that $\mathcal{C}$ satisfies the hypotheses of Conjecture 3.1. First, we must show that $\mathcal{C}$ is open convex. An open convex realization of $\mathcal{C}$ is provided in Figure 2.

Additionally, Conjecture 3.1 states that $\mathcal{C}$ must be max-intersection incomplete. Note that the facets of $\mathcal{C}$ are $123, 124, 135$, and 236. The code $\mathcal{C}$ is max-intersection incomplete because $135 \cap 236 = 3 \notin \mathcal{C}$.

We must also show that $\mathcal{C}$ has at least 2 non-mandatory codewords. The simplicial complex of $\mathcal{C}$ is $\Delta(\mathcal{C}) = \{\mathbf{123}, \mathbf{124}, \mathbf{135}, \mathbf{236}, 12, 13, 14, 15, 23, 24, 26, 35, 36, 1, 2, 3, 4, 5, 6, \emptyset\}$. From this set, we get $\text{Lk}_{\{3\}}(\Delta(\mathcal{C})) = \{12, 15, 26, 1, 2, 5, 6, \emptyset\}$, which is contractible (see Figure 3). Additionally, $\text{Lk}_{\{4\}}(\Delta(\mathcal{C})) = \{12, 1, 2, \emptyset\}$ is contractible (see Figure 4). Therefore $\mathcal{C}$ has at least two non-mandatory codewords, namely 3 and 4.

Finally, we must show that $\mathcal{C}$ has at least 3 facets M_1, M_2, M_3 such that there is a $\sigma \in \mathcal{C}$ with the property $\sigma \subset M_1$ 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 it is true that $\sigma \subset M_1$, because $13 \subset 123$. It is also true that $\sigma \cap M_2 \notin \mathcal{C}$, as $13 \cap 236 = 3 \notin \mathcal{C}$.

Therefore, our code $\mathcal{C} = \{\emptyset, 1, 2, 12, 13, 14, 23, 24, \mathbf{123}, \mathbf{124}, \mathbf{135}, \mathbf{236}\}$ is open convex, max-intersection-incomplete, has at least 2 non-mandatory codewords, and has 3 facets M_1, M_2, M_3 such that there is a $\sigma \in \mathcal{C}$ with the property $\sigma \subset M_1$ and $\sigma \cap M_2 \notin \mathcal{C}$. However, a closed convex realization of $\mathcal{C}$ can be seen in Figure 5, which was obtained by including the boundaries of the place fields given in the open convex realization in Figure 2. This proves that $\mathcal{C}$ is a counterexample to Conjecture 3.1. □

3.2 Open but not Closed Convex Codes

Conjecture 3.3. *Let $\mathcal{C}$ be a locally good, 3-sparse code. Then $\mathcal{C}$ must be closed convex.*

Conjecture 3.4. *Let $\mathcal{C}$ be a locally good, 3-sparse code. Then $\mathcal{C}$ must be open convex.*

It has already been proven by [2] and [3] that if a neural code is convex, it is also locally good. It is natural to wonder if the converse is also true. It was also stated in [2] that 2-sparse, locally good codes are always both open and closed convex. Additionally, it has been proven in [5] that codes which are 4-sparse or higher are not necessarily convex, even when they are locally good. So all that is left is the 3-sparse case. In this section we provide a counterexample to Conjecture 3.3.

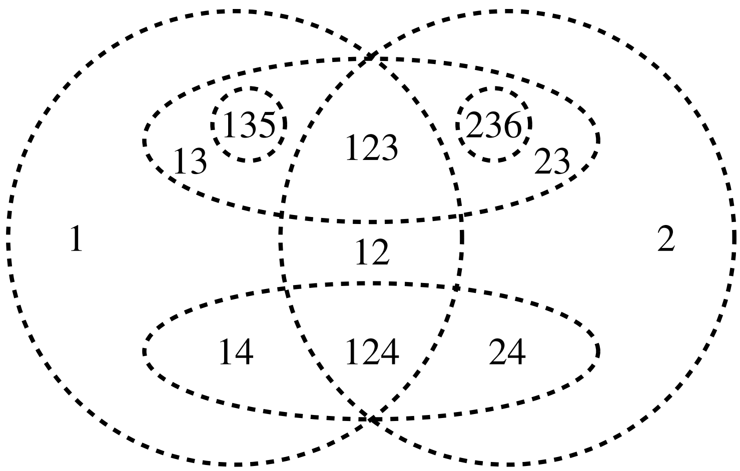

Figure 2: Open convex realization of counterexample to Goldrup and Phillipson conjecture

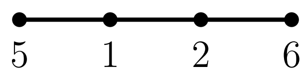

Figure 3: $\mathrm{Lk}_{\{3\}}(\Delta(\mathcal{C})) = \{12, 15, 26, 1, 2, 5, 6, \emptyset\}$ is contractible

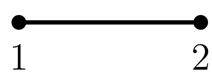

Figure 4: $\mathrm{Lk}_{\{4\}}(\Delta(\mathcal{C})) = \{12, 1, 2, \emptyset\}$ is contractible

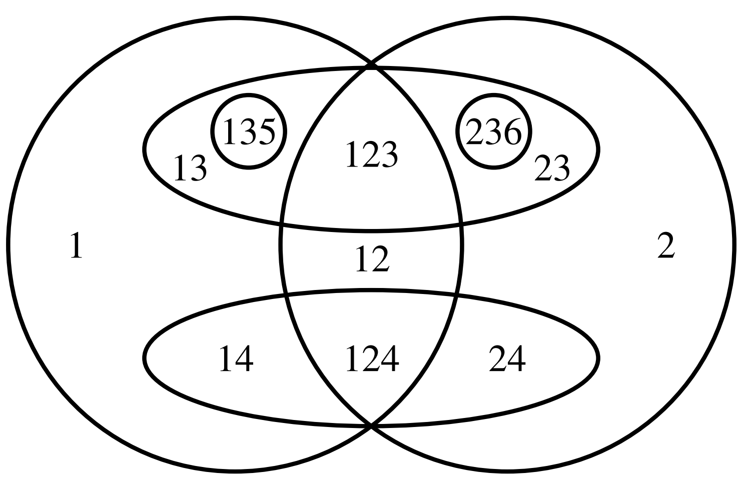

Figure 5: Closed convex realization of counterexample to Goldrup and Phillipson conjecture

In Theorem 3.5, we give a neural code that is not closed convex. In Corollary 3.6, we show that this code in fact is also 3-sparse and locally good, making it a valid counterexample to Conjecture 3.3.

Theorem 3.5. *The neural code* $\mathcal{C} = \{\emptyset, 4, 5, 12, 14, 23, 35, \mathbf{45}, \mathbf{123}, \mathbf{124}, \mathbf{235}\}$ *is open convex and not closed convex.*

Proof. An open convex realization of this neural code is provided in Figure 6. Additionally, this code can be obtained by adding the codewords $\{14, 25\}$ to the code labeled $\mathcal{C}_3$ in [4] and then permuting the indexing of the neurons. We will utilize the method outlined in [1] to exhibit that $\mathcal{C}$ is not closed convex.

Assume for contradiction that there exists some closed convex realization $\mathcal{U} = \{\mathcal{U}_i\}_{i=1}^5$ for $\mathcal{C}$ in $\mathbb{R}^d$. As $\mathcal{U}_{14} \cap \mathcal{U}_{35} = \emptyset$, we can pick distinct points $x_{14} \in \mathcal{U}_{14}$ and $x_{35} \in \mathcal{U}_{35}$. Let $L_1 = \overline{x_{14}, x_{35}}$. As $\mathcal{U}_{123}$ is closed, we can choose x_{123} such that $d(x_{123}, L_1) \leq d(y_{123}, L_1)$ for all $y_{123} \in \mathcal{U}_{123}$.

Let $L_2 = \overline{x_{14}, x_{123}}$. Both x_{14} and x_{123} lie in $\mathcal{U}_1$. As $\mathcal{U}_1$ is convex, $L_2 \subset \mathcal{U}_1$. Note that $\mathcal{U}_1 \subset \mathcal{U}_2 \cup \mathcal{U}_4$, as whenever the neuron 1 appears in the code it is always accompanied by either a 2 or a 4 (or both). Thus we can say that $L_2 \subset \mathcal{U}_2 \cup \mathcal{U}_4$. Both $L_2 \cap \mathcal{U}_2$ and $L_2 \cap \mathcal{U}_4$ are closed and non-empty. As L_2 is a line segment it is connected, and therefore $L_2 \cap \mathcal{U}_2 \cap \mathcal{U}_4 \subset \mathcal{U}_{124}$ must be nonempty.So there exists some $x_{124} \in \mathcal{U}_{124}$ such that $x_{124} \in L_2$.

Let $L_3 = \overline{x_{35}, x_{123}}$. Both x_{35} and x_{123} lie in $\mathcal{U}_3$. As $\mathcal{U}_3$ is convex, $L_3 \subset \mathcal{U}_3$. Note that $\mathcal{U}_3 \subset \mathcal{U}_2 \cup \mathcal{U}_5$, as whenever the neuron 3 appears in the code it is always accompanied by either a 2 or a 5 (or both). Thus we can say that $L_3 \subset \mathcal{U}_2 \cup \mathcal{U}_5$. Both $L_3 \cap \mathcal{U}_2$ and $L_3 \cap \mathcal{U}_5$ are closed and non-empty. As L_3 is a line segment it is connected, and therefore $L_3 \cap \mathcal{U}_2 \cap \mathcal{U}_5 \subset \mathcal{U}_{124}$ must be nonempty. So there exists some $x_{235} \in \mathcal{U}_{235}$ such that $x_{235} \in L_3$. Let $K = \overline{x_{124}, x_{235}}$. Both x_{124} and x_{235} lie in $\mathcal{U}_2$. As $\mathcal{U}_2$ is convex, $K \subset \mathcal{U}_2$. Note that $\mathcal{U}_2 \subset \mathcal{U}_1 \cup \mathcal{U}_3$, as whenever the neuron 2 appears in the code it is always accompanied by either a 1 or a 3 (or both). Thus we can say that $K \subset \mathcal{U}_1 \cup \mathcal{U}_3$. Both $K \cap \mathcal{U}_1$ and $K \cap \mathcal{U}_3$ are closed and non-empty. As K is a line segment it is connected, and therefore $K \cap \mathcal{U}_1 \cap \mathcal{U}_3 \subset \mathcal{U}_{123}$ must be nonempty. So there exists some $y_{123} \in \mathcal{U}_{123}$ such that $y_{123} \in K$.

Thus we see that y_{123} lies in the interior of the triangle $\triangle(x_{123}, x_{14}, x_{35})$. However, this would mean that $d(y_{123}, L_1) < d(x_{123}, L_1)$, which would contradict the fact that we said x_{123} was the closest point to L_1 in $\mathcal{U}_{123}$. Thus we have reached a contradiction, implying that $\mathcal{C}$ cannot be realized as a collection of closed convex sets. ☐

We now disprove Conjecture 3.3 by showing that $\mathcal{C}$ from Theorem 3.5 is a counterexample.

Corollary 3.6. *The neural code* $\mathcal{C} = \{\emptyset, 4, 5, 12, 14, 23, 35, \mathbf{45}, \mathbf{123}, \mathbf{124}, \mathbf{235}\}$ *is a 3-sparse locally good code that is not closed convex.*

Proof. From Theorem 3.5 we know that $\mathcal{C}$ is not closed convex. A quick inspection of the code verifies that $\mathcal{C}$ is 3-sparse. Thus all that remains is to show that $\mathcal{C}$ is locally good. Thankfully, [2] and [3] have shown that if a neural code is open convex it must be locally good. Thus, as $\mathcal{C}$ is open convex by Theorem 3.5, it must be locally good. ☐

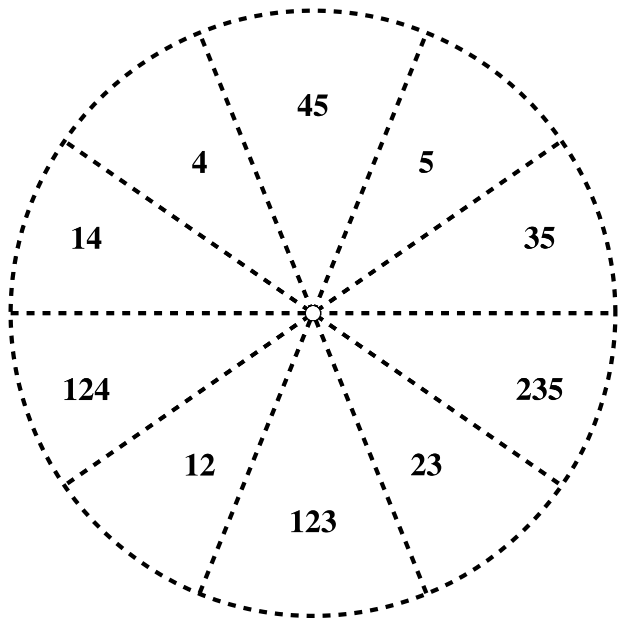

Figure 6: An open convex realization of the code in Theorem 3.4 and Corollary 3.5

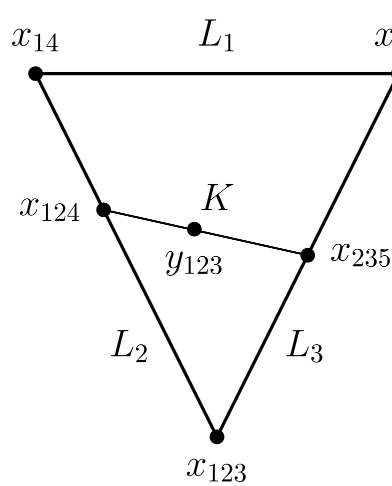

Figure 7: Visual proof of Theorem 3.5

4 Open and Closed Convexity in $\mathbb{R}^d$

In this section, we provide several results establishing connections between closed and open convex codes. Specifically, we provide methods of taking closed convex realizations of neural codes and using them to construct open convex realizations.

As much of this section works with dimensionality of codes, we provide the following definition:

Definition 4.1. A set $\mathcal{U}_\sigma$ is *full-dimensional* in $\mathbb{R}^d$ if it contains an open ball in $\mathbb{R}^d$. Similarly, a realization $\mathcal{U} = \{\mathcal{U}_i\}_{i=1}^n$ of a code C is *full-dimensional* if each $\mathcal{U}_i$ is full-dimensional.

Lemma 4.2. *Let $\mathcal{C}$ be a closed convex neural code on n neurons. Then there exists a closed convex realization of $\mathcal{C}$, $\mathcal{U} = \{\mathcal{U}_i\}_{i=1}^n$ in $\mathbb{R}^d$, for some $d \in \mathbb{N}$, such that each $\mathcal{U}_i$ is compact.*

Proof. Let $\mathcal{C}$ be a closed convex neural code on n neurons. Then there exists a closed convex realization of $\mathcal{C}$, $\mathcal{U} = \{\mathcal{U}_i\}_{i=1}^n$ in $\mathbb{R}^d$. If $\mathcal{U}_\sigma$ is compact for each $\sigma \subset [n]$, then we are done. So assume that there exists some subset of $[n]$, S, such that for each $j \in S$, it is the case that $\mathcal{U}_j$ is not compact. Draw a closed ball $\overline{B}$ such that, for all $i \in [n]$, $\overline{B}$ contains a point in $\mathcal{U}_i$. Redraw each $\mathcal{U}_i$ as $\mathcal{U}_i \cap \overline{B}$. As the intersection of closed sets is closed and the intersection of convex sets is convex, this process preserves both. Thus we are left with a compact, closed convex realization of $\mathcal{C}$. $\square$

For the remainder of this paper we will assume for brevity that all closed convex realizations are compact due to Lemma 4.2.

Lemma 4.3. *Let $\mathcal{C}$ be a neural code on n neurons with a closed convex realization $\mathcal{U} = \{\mathcal{U}_i\}_{i=1}^n$ in $\mathbb{R}^d$. If every set in $\mathcal{U}$ is full-dimensional, then removing the boundaries of every set $\mathcal{U}$ will not add any codewords.*

Proof. Let $\mathcal{C}$ be a neural code on n neurons with a closed convex realization $\mathcal{U} = \{\mathcal{U}_i\}_{i=1}^n$ in $\mathbb{R}^d$. Suppose that every set in in $\mathcal{U}$ is full-dimensional. Take some $x \in \mathbb{R}^d$. If x does not lie on the boundary of some set in $\mathcal{U}$, then it will remain in whatever sets it lies in if the boundaries of all sets in $\mathcal{U}$ were removed, and thus will not be included in any new sets.

Suppose without loss of generality that for $1 \leq i \leq j \leq n$:

- $x \in \partial\mathcal{U}_1, \ldots, \partial\mathcal{U}_i$
- $x \in \text{int}(\mathcal{U}_{i+1}), \ldots, \text{int}(\mathcal{U}_j)$
- $x \notin \mathcal{U}_{j+1}, \ldots, \mathcal{U}_n$

Thus, by removing the boundary of every set, $x \in \mathcal{U}_{i+1} \cap ... \cap \mathcal{U}_j \cap \mathcal{U}_1^C \cap \cdots \cap \mathcal{U}_i^C \cap \mathcal{U}_{j+1}^C \cap \cdots \cap \mathcal{U}_n^C$.

We will show that before removing the boundaries of every set, the set $\mathcal{U}_{i+1} \cap ... \cap \mathcal{U}_j \cap \mathcal{U}_1^C \cap \cdots \cap \mathcal{U}_i^C \cap \mathcal{U}_{j+1}^C \cap \cdots \cap \mathcal{U}_n^C$ is non-empty.

As $x \in \text{int}(\mathcal{U}_{i+1}), \ldots, \text{int}(\mathcal{U}_j)$, this means that $\text{int}(\mathcal{U}_{i+1}) \cap \cdots \cap \text{int}(\mathcal{U}_j) \neq \emptyset$. As we are intersecting the interiors of these sets, which are in full-dimension, their intersection must too be in full-dimension. Ergo there must exist an open ball in $\mathbb{R}^d$ that is contained in this intersection. Let $\tilde{x} \neq x$ be some point in this open ball. Thus $\tilde{x} \in \mathcal{U}_{i+1} \cap ... \cap \mathcal{U}_j \cap \mathcal{U}_1^C \cap \cdots \cap \mathcal{U}_i^C \cap \mathcal{U}_{j+1}^C \cap \cdots \cap \mathcal{U}_n^C$. As it resides within the interior of this set, it would not be removed with the boundary, showing that the set is therefore non-empty.

Thus every boundary point of every set in $\mathcal{U}$ exists in the interior of some other set, which would remain after the boundaries were removed. Thus no new codewords could be added by removing the boundaries of all sets in $\mathcal{U}$. $\square$

Theorem 4.4. *Let $\mathcal{C}$ be a neural code on n neurons with a closed convex realization $\mathcal{U} = \{\mathcal{U}_i\}_{i=1}^n$ in $\mathbb{R}^d$. If every region corresponding to a codeword in $\mathcal{C}$ is full-dimensional, then $\mathcal{C}$ is open convex.*

Proof. Let $\mathcal{C}$ be a neural code on n neurons with a closed convex realization $\mathcal{U} = \{\mathcal{U}_i\}_{i=1}^n$ in $\mathbb{R}^d$. Suppose that every set in $\mathcal{U}$ can be expressed fully in $\mathbb{R}^d$. As $\mathcal{C}$ is closed convex, $\mathcal{U}_i$ is closed for all $i \in [n]$. Thus $\mathcal{U}_i = \partial\mathcal{U}_i \cup \text{int}(\mathcal{U}_i)$. For each $i \in [n]$, remove the boundary of $\mathcal{U}_i$, leaving only the interior of each set, which is open. By definition, boundaries of convex sets in $\mathbb{R}^d$ are not full-dimensional in $\mathbb{R}^d$. As we supposed that every set in $\mathcal{U}$ can be expressed as full-dimensional in $\mathbb{R}^d$, removing these boundaries does not remove any codewords from the code. From Lemma 4.3 we know that no new codewords could be added. Thus we are left with an open convex realization of $\mathcal{C}$. $\square$

Lemma 4.5. *Let $\mathcal{C}$ be a neural code on n neurons with a closed convex realization $\mathcal{U} = \{\mathcal{U}_i\}_{i=1}^n$ in $\mathbb{R}^d$. Suppose every region corresponding to a codeword in C is full-dimensional in $\mathbb{R}^d$ except for $\mathcal{U}_{\alpha_1}, \mathcal{U}_{\alpha_2}, \dots, \mathcal{U}_{\alpha_m}$. If, for each $1 \le i \le m$, the sets $\mathcal{U}_{\alpha_i}$ are disjoint from the boundary of $\mathcal{U}_\sigma$ for all $\sigma \in \mathcal{C}$, then $\mathcal{C}$ is open convex.*

See Figure 8 for a visual example of this lemma.

Proof. Let $\mathcal{C}$ be a neural code on n neurons with a closed convex realization $\mathcal{U} = \{\mathcal{U}_i\}_{i=1}^n$ in $\mathbb{R}^d$. Suppose $\mathcal{U}_{\alpha_1}, \mathcal{U}_{\alpha_2}, \dots, \mathcal{U}_{\alpha_m}$ are less than d-dimensional and none intersect the boundary of $\mathcal{U}_\sigma$ for any $\sigma \in \mathcal{C}$. We begin with $\mathcal{U}_{\alpha_1}$. As all sets in $\mathcal{U}$ are closed and compact, we can choose $\epsilon > 0$ such that $\epsilon < d(\mathcal{U}_{\alpha_1}, \mathcal{U}_\beta)$ for all $\beta \in \mathcal{C}$ such that $\mathcal{U}_{\alpha_1} \cap \mathcal{U}_\beta = \emptyset$. If $\mathcal{U}_{\alpha_1} \subset \mathcal{U}_\lambda$ for some $\lambda \in \mathcal{C}$, then re-choose ϵ such that $\epsilon < d(\mathcal{U}_{\alpha_1}, \partial\mathcal{U}_\lambda)$ as well. Then draw a closed ϵ-neighborhood around $\mathcal{U}_{\alpha_1}$. Redefine $\mathcal{U}_{\alpha_1}$ as this larger closed set. Since the ϵ-neighborhood of a convex set is also convex, the closed convexity of $\mathcal{U}_{\alpha_1}$ is preserved.

As this new realization of $\mathcal{U}_{\alpha_1}$ does not intersect any sets that it was not already overlapping (or fire by itself if it was the subset of some $\mathcal{U}_\lambda$), the code remains unchanged. Repeat this process for each $\mathcal{U}_{\alpha_i}$. Then we have obtained a closed convex realization of $\mathcal{C}$ with the property that every set in $\mathcal{U}$ can be expressed fully in $\mathbb{R}^d$. Thus by Theorem 4.4, $\mathcal{C}$ is open convex. $\square$

These results may also be applied to Lemma 4.2. If by drawing a closed ball $\overline{B}$ containing at least one point from each $\mathcal{U}_i$, we end up with any set that is not full-dimensional, we can use the process outlined in Lemma 4.5 to expand this set into d dimensions.

Lemma 4.6. *Let $\mathcal{C}$ be a neural code on n neurons with a closed convex realization $\mathcal{U} = \{\mathcal{U}_i\}_{i=1}^n$ in $\mathbb{R}^d$. Suppose every set in $\mathcal{U}$ can be expressed fully in $\mathbb{R}^d$ except for $\mathcal{U}_{\alpha_1}, \mathcal{U}_{\alpha_2}, \dots, \mathcal{U}_{\alpha_m}$. If, for each $1 \le i \le m$, the set $\mathcal{U}_{\alpha_i} \cap \partial\mathcal{U}_\gamma \ne \emptyset$ for exactly one $\gamma \in \mathcal{C}$ such that $\mathcal{U}_\gamma$ is in full-dimension, then $\mathcal{C}$ is open convex.*

Lemma 4.6 states that we can expand any closed convex set to be full-dimensional in $\mathbb{R}^d$ without changing the original code. The following proof utilizes closed ϵ-neighborhoods to expand each such set, and then removes their boundaries, leaving open convex sets.

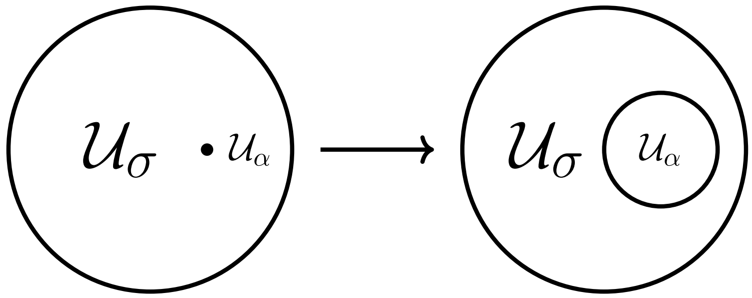

Figure 8: Visualization of Proof of Lemma 4.5 in $\mathbb{R}^2$

Proof. Let $\mathcal{C}$ be a neural code on n neurons with a closed convex realization $\mathcal{U} = \{\mathcal{U}_i\}_{i=1}^n$ in $\mathbb{R}^d$. Suppose $\mathcal{U}_{\alpha_1}, \mathcal{U}_{\alpha_2}, \ldots, \mathcal{U}_{\alpha_n}$ are not full-dimensional in $\mathbb{R}^d$, and for each $1 \leq i \leq n$ the set $\mathcal{U}_{\alpha_i} \cap \partial\mathcal{U}_\gamma \neq \emptyset$ for exactly one $\gamma \in \mathcal{C}$.

We begin with $\mathcal{U}_{\alpha_1}$. There are two cases: either $\mathcal{U}_{\alpha_1} \subset \mathcal{U}_\gamma$ or $\mathcal{U}_{\alpha_1} \not\subset \mathcal{U}_\gamma$.

First, we suppose $\mathcal{U}_{\alpha_1} \subset \mathcal{U}_\gamma$. As all sets in $\mathcal{U}$ are closed, we can choose $\epsilon > 0$ such that $\epsilon < d(\mathcal{U}_{\alpha_1}, \mathcal{U}_\beta)$ for all $\beta \in \mathcal{C}$ such that $\mathcal{U}_{\alpha_1} \cap \mathcal{U}_\beta = \emptyset$. If $\mathcal{U}_{\alpha_1} \subset \mathcal{U}_\lambda$ for some $\lambda \neq \gamma \in \mathcal{C}$, then re-choose ϵ such that $\epsilon < d(\mathcal{U}_{\alpha_1}, \partial\mathcal{U}_\lambda)$ as well. Draw a closed ϵ-neighborhood around $\mathcal{U}_{\alpha_1}$ and call this $\tilde{\mathcal{U}}_{\alpha_1}$. Define $\tilde{\mathcal{U}}_{\alpha_1}$ as the union of the original set and its closed ϵ-neighborhood. Redefine $\mathcal{U}_{\alpha_1}$ as the set $\mathcal{U}_\gamma \cap \tilde{\mathcal{U}}_{\alpha_1}$. As both $\mathcal{U}_\gamma$ and $\tilde{\mathcal{U}}_{\alpha_1}$ are convex, so is $\mathcal{U}_{\alpha_1}$. By construction, this new realization of $\mathcal{U}_{\alpha_1}$ does not intersect any sets that it was not already overlapping (or fire by itself if it was the subset of some $\mathcal{U}_\lambda$), so no codewords appear. Repeat this process for each $\mathcal{U}_{\alpha_i}$. Since every set in $\mathcal{U}$ is now in full-dimensional, by Theorem 4.4 we can remove the boundary of each set, leaving an open convex realization of $\mathcal{C}$.

Now suppose that $\mathcal{U}_{\alpha_1} \not\subset \mathcal{U}_\gamma$. As all sets in $\mathcal{U}$ are closed and compact, we can define z to be the point furthest away from $\mathcal{U}_{\alpha_1}$ that is in $\mathcal{U}_\gamma$. Next, we choose $\epsilon > 0$ such that $\epsilon < \min\{d(\mathcal{U}_{\alpha_1}, \mathcal{U}_\beta), z\}$ for all $\beta \in \mathcal{C}$ such that $\mathcal{U}_{\alpha_1} \cap \mathcal{U}_\beta = \emptyset$. If $\alpha_1 \in \mathcal{C}$, then this is a sufficient ϵ. If $\alpha_1 \notin \mathcal{C}$, then there must exist $\mathcal{U}_\lambda \subset \mathcal{U}$ such that $\mathcal{U}_{\alpha_1} \subset \mathcal{U}_\lambda$. If this is the case, re-choose ϵ such that $\epsilon < d(\mathcal{U}_{\alpha_1}, \partial\mathcal{U}_\lambda)$ as well. Then draw a closed ϵ-neighborhood around $\mathcal{U}_{\alpha_1}$. Define $\tilde{\mathcal{U}}_{\alpha_1}$ as the union of the original set and its closed ϵ-neighborhood.

As $\tilde{\mathcal{U}}_{\alpha_1}$ does not intersect any sets that it was not already overlapping (or fire by itself if it was the subset of some $\mathcal{U}_\lambda$), the code remains unchanged. Repeat this process for each $\mathcal{U}_{\alpha_i}$. Since every set in $\mathcal{U}$ is now full-dimensional, by Theorem 4.4 we can remove the boundary of each set, leaving an open convex realization of $\mathcal{C}$. □

Theorem 4.7. *Let $\mathcal{C}$ be a neural code on n neurons with a closed convex realization $\mathcal{U} = \{\mathcal{U}_i\}_{i=1}^n$ in $\mathbb{R}^d$. Suppose every set in $\mathcal{U}$ is full-dimensional in $\mathbb{R}^d$ except for $\mathcal{U}_{\alpha_1}, \mathcal{U}_{\alpha_2}, \ldots, \mathcal{U}_{\alpha_m}$, which are $(d-1)$-dimensional. Then $\mathcal{C}$ is open convex.*

Proof. Let $\mathcal{C}$ be a neural code on n neurons with a closed convex realization $\mathcal{U} = \{\mathcal{U}_i\}_{i=1}^n$ in $\mathbb{R}^d$. Suppose for contradiction that $\mathcal{C}$ is not open convex. Thus the closed convex realization must be constructed in such a way that removing the boundary of each $\mathcal{U}_i$ would remove some codeword from or add some codeword to the neural code. Moreover, there could easily be more than one such codeword. For $\alpha_i \in \mathcal{C}$, let $\mathcal{U}_\alpha = \mathcal{U}_{\alpha_1} \cup \mathcal{U}_{\alpha_2} \cup \cdots \cup \mathcal{U}_{\alpha_n}$, where $\mathcal{U}_{\alpha_1}, \mathcal{U}_{\alpha_2}, \ldots, \mathcal{U}_{\alpha_n}$ are all of the sets that would be removed in this way. This means that each $\mathcal{U}_{\alpha_i}$ must have existed on the boundary of

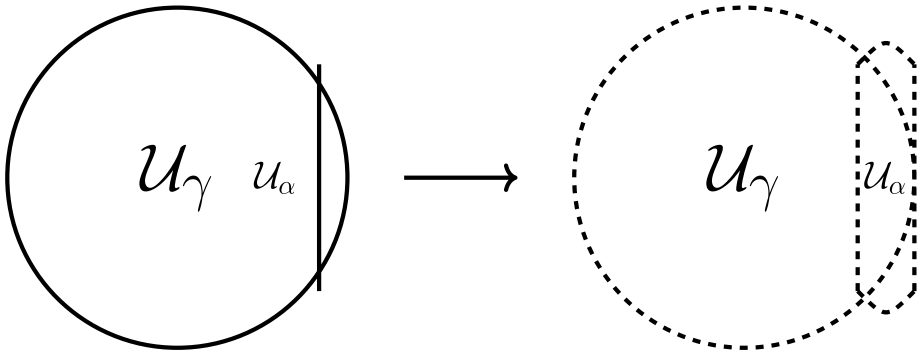

Figure 9: Visualization of Proof of Lemma 4.6 in $\mathbb{R}^2$

some other set. Boundaries of sets are not full-dimensional by definition. Thus, by Lemmas 4.5 and 4.6, we know that each $\mathcal{U}_{\alpha_i}$ is completely contained in the intersection of at least two sets. As the intersection of convex sets is convex, this intersection must either be full-dimensional in $\mathbb{R}^{d-1}$ or full-dimensional in some lower dimensional, affine subspace. It will now be shown that $\mathcal{U}_\alpha$ cannot be full-dimensional in $\mathbb{R}^{d-1}$.

Suppose for contradiction that each $\mathcal{U}_{\alpha_i}$ can be drawn as full-dimensional in $\mathbb{R}^{d-1}$ and is contained in the intersection of at least two sets. For each $\sigma \in \mathcal{C}$, choose some point $x_\sigma \in \mathcal{U}_\sigma$. Let z_σ equal the smallest distance from x_σ to the boundary of any other set in $\mathcal{C}$. Let $z = \min\{z_\sigma | \sigma \in C\}$.

Choose $\epsilon > 0$ such that $\epsilon < \min\{d(\mathcal{U}_i, \mathcal{U}_j), z\}$ for each $\mathcal{U}_i, \mathcal{U}_j \subset \mathcal{U}$ such that $\mathcal{U}_i \cap \mathcal{U}_j = \emptyset$. For each $1 \leq i \leq n$, draw a closed neighborhood of radius $\frac{\epsilon}{n}$. Redefine each $\mathcal{U}_i$ as this closed neighborhood. As our realization was full-dimensional, by Lemma 4.3, no new codewords were added to the code. Because of our choice of epsilon, no codewords were removed either. Thus the code remains unchanged, and every $\mathcal{U}_{\alpha_i}$ has now been expressed fully in $\mathbb{R}^d$. Thus by Theorem 4.4, $\mathcal{C}$ is open convex. □

Conjecture 4.8. *Let $\mathcal{C}$ be a neural code of n neurons with a closed convex realization $\mathcal{U} = \{\mathcal{U}_i\}_{i=1}^n$ in $\mathbb{R}^d$. If for every $\sigma \in \mathcal{C}$ such that the region $\mathcal{U}_\sigma$ must have dimension $(d-2)$ or less, if $\mathcal{U}_\sigma$ is equal to the intersection of three or fewer sets in $\mathcal{U}$, then $\mathcal{C}$ is open convex.*

Thus far, each locally good neural code that is not open convex has been closed convex. If Conjecture 4.8 can be proved, it would show that for a code to be closed convex but not open convex, an intersection of four or more sets would be required. This would mean that there does not exist a 3-sparse, locally good, closed convex code that is not open convex.

Proposition 4.9. *If Conjecture 4.8 is true, then any 3-sparse, locally good, closed convex code is also open convex.*

5 A Faster Method to Prove Open Convexity

Theorem 5.1. *Let $\mathcal{C}$ be a 3-sparse neural code on n neurons. If there exists a closed convex realization $\mathcal{U} = \{\mathcal{U}_i\}_{i=1}^n$ in $\mathbb{R}^d$ of $\mathcal{C}_{red}$ such that each $\mathcal{U}_i$ is $(d-1)$ or d-dimensional, then $\mathcal{C}$ is open convex.*

Proof. Let $\mathcal{C}$ be a 3-sparse locally good neural code on n neurons. Suppose that there exists some full-dimensional closed realization of $\mathcal{C}_{red}$, denoted $\mathcal{U} = \{\mathcal{U}_i\}_{i=1}^n$ in $\mathbb{R}^d$. We will construct an open realization of $\mathcal{C}$ using $\mathcal{U}$.

We begin with intersections of neurons in $\mathcal{C}_{red}$. Let $\lambda_1 \subseteq \sigma_1$ and $\lambda_2 \subseteq \sigma_2$ for some distinct $\sigma_1, \sigma_2 \in C_{red}$ such that $\lambda_1\lambda_2$ is a facet of $\mathcal{C}$. As both $\mathcal{U}_{\lambda_1}$ and $\mathcal{U}_{\lambda_2}$ are involved in separate triple-wise intersections, to avoid local obstructions it must be the case that $\lambda_1, \lambda_2 \in \mathcal{C}$. Extend both $\mathcal{U}_{\lambda_1}$ and $\mathcal{U}_{\lambda_2}$ out into $\mathbb{R}^{d+1}$ some distance such that they overlap in full dimension. This is possible because no other sets exist in $\mathbb{R}^{d+1}$. We will write x to denote the distance of this intersection from the rest of the realization.

Without loss of generality, it is now the case that the intersection between $\mathcal{U}_{\lambda_1}$ and any other set is still in $\mathbb{R}^d$. However, by using the same epsilonic expansion technique used in the proof of Theorem 4.7, we can re-express our realization as full-dimensional in $\mathbb{R}^{d+1}$. However, unlike in Theorem 4.7, we are refraining from removing the boundaries of our closed sets just yet.

Repeat this process for any other intersections of neurons that are already in $\mathcal{C}_{red}$. However, for each successive intersection, extend the sets into the dimension one above the previous. For example, the next sets would be extended into $\mathbb{R}^{d+2}$. This prevents any unwanted overlap between sets.

Thus all that remains is to incorporate neurons in $\mathcal{C}$ that are missing from $\mathcal{C}_{red}$. The only neurons missing from $\mathcal{U}$ are the ones not involved in any triple-wise intersection. Let $A = \{a_1, a_2, \ldots, a_m\} \subset \mathcal{C}$ denote the set of these neurons.

Begin with some a_i such that a_i fires with some neuron $b_1 \in C_{red}$. If it fires with a second neuron b_2, begin this process with b_1 and repeat again for b_2. As $\mathcal{C}$ is locally good, it follows that $b_1 \in \mathcal{C}$ to avoid a local obstruction. Extend b_1 some distance x into a new dimension as we did previously.

Either $a_i \in \mathcal{C}$ or it isn't. If not, then a_i cannot fire with any other neuron in $\mathcal{C}$. Thus $\mathcal{U}_{a_i} \subset \mathcal{U}_b$, and we can draw $\mathcal{U}_{a_i}$ inside $\mathcal{U}_b$ such that $\mathcal{U}_{a_i}$ does not come into contact with any other set. If $a_i \in \mathcal{C}$, then draw $\mathcal{U}_{a_i} \in \mathbb{R}^{d+1}$ such that it only overlaps with $\mathcal{U}_b$ but is not a subset of $\mathcal{U}_b$.

It is now the case that the intersection between $\mathcal{U}_b$ and any other set that was extended in the opposite direction is still in $\mathbb{R}^d$. However, once again by Theorem 4.7 we can re-express our realization fully once again.

Repeat this process for each $1 \leq i \leq m$, each time selecting a codeword that contains a neuron already existing in the realization. This provides us with a full-dimensional closed realization of $\mathcal{C}$. Thus, by Theorem 4.7, $\mathcal{C}$ is open convex. $\square$

Example 5.2. Consider the neural code

$$\mathcal{C} = \{123, 134, 145, 13, 14, 26, 27, 29, 36, 37, 38, 46, 48, 49, 58, 67, 79, 89, 2, 3, 4, 5, 6, 7, 8, 9, \emptyset\}.$$

Not the easiest code to visualize. Unfortunately, $\mathcal{C}$ is not max-intersection complete as $123 \cap 145 = 1$ but $1 \notin \mathcal{C}$. However, the link of 1 is contractible, and as $\mathcal{C}$ contains all of its mandatory codewords, it is locally good. So how are we to know if this code is open convex? We could attempt to draw a realization of the entire code, but this would be tiresome and should be avoided if possible.

Thankfully, by Theorem 5.1, we know that if $\mathcal{C}_{red} = \{123, 134, 145, 13, 14, \emptyset\}$ is closed convex such that every set can be realized in $\mathbb{R}^{d-1}$ or higher, then $\mathcal{C}$ must be open. This is a much simpler code to think about, and closed convex realization of $\mathcal{C}_{red}$ meeting these requirements is given below. Thus $\mathcal{C}$ is open convex.

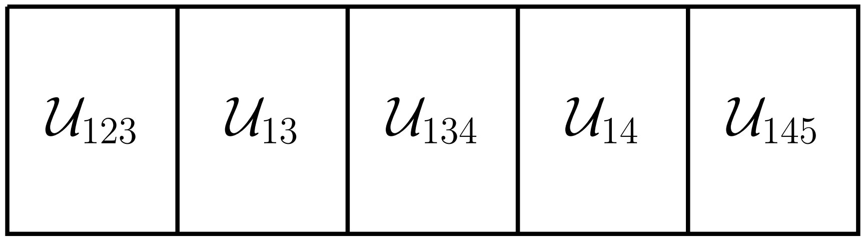

Figure 10: Open convex visualization of $\mathcal{C}_{red}$ in $\mathbb{R}^2$

6 Additional Results and Future Research

Mathematics never has been and never will be a linear pursuit. Progress occurs with many backtracks and even more dead ends. This paper was certainly no exception. However, one mathematician's dead end may be the key to another's most important proof. In this section we provide a number of results we proved that were not essential to our primary goals. We hope that these results, though not extraordinary in their own right, may at some point down the road prove themselves useful. Or at least avoid the case in which some other mathematician spends an afternoon of proving some intermediary step that we could have provided here.

Theorem 6.1. *Let $\mathcal{C}$ be a locally good neural code on n neurons. If each intersection of two or more facets consists of at most one neuron, then $\mathcal{C}$ is max-intersection complete.*

Proof. Let $\mathcal{C}$ be a locally good neural code on n neurons with distinct facets $M_1, M_2, \ldots, M_k$ such that no intersection of two or more facets contains more than one neuron. Let σ equal the intersection of some facets $M_{i_1}, \ldots M_{i_j}$. From our suppositions we see $M_{i_a} \cap M_{i_b} = \sigma$ for $1 \leq a \neq b \leq j$. Thus $Lk_\sigma \Delta(\mathcal{C})$ is a collection of j disjoint points. Thus the link is not contractible, and to remain locally good σ must be in the code. Thus $\mathcal{C}$ is max-intersection complete. □

Lemma 6.2. *Let $\mathcal{C}$ be a 3-sparse code on n neurons. If $\mathcal{C}$ has an intersection of facets, σ, of size two, then σ is a mandatory codeword.*

Proof. Let $\mathcal{C}$ be a 3-sparse code on n neurons. Suppose that $\mathcal{C}$ has an intersection of distinct facets, σ, of size two. As $\mathcal{C}$ is 3-sparse, by definition $Lk_\sigma(\Delta)$ can only consist of isolated points. The existence of any connected points would require a facet of size four. Thus $Lk_\sigma(\Delta)$ is non-contractible, and σ is a mandatory codeword of $\mathcal{C}$. □

Theorem 6.3. *Let $\mathcal{C}$ be a 3-sparse locally good max-intersection incomplete code. Then there must be at least three codewords of size three or more.*

Proof. Let $\mathcal{C}$ be a 3-sparse locally good max-intersection incomplete code with distinct facets $M_1, M_2, \ldots, M_n$. Then there must exist some codeword $\sigma \notin \mathcal{C}$ and $M_i, M_j \in \mathcal{C}$ such that $M_i \cap M_j = \sigma$. From Lemma 6.2 we know that σ must be of size one. The facets M_i and M_j must be of at least size two to remain distinct, given their shared σ. However, if M_i and M_j were of size two, then σ would be a mandatory codeword. Therefore M_i and M_j must be of size three. From Lemma 6.2 we see that σ cannot be of size two, and is thus a single neuron.

However, this could not be the entire code, as this would make σ a mandatory codeword. Thus, there must exist some other facet $M_k \in \mathcal{C}$ such that $M_i \cap M_j \cap M_k = \sigma$. This would maintain σ as the maximal intersection of facets.

To remain distinct from M_i and M_j, M_k must contain some neuron $\tau \notin M_i, M_j$. However, if both $M_i \cap M_k = \sigma$ and $M_j \cap M_k = \sigma$, then there would exist a local obstruction at σ, as τ would be disjoint from all other neurons in $Lk_\sigma(\Delta)$. Thus, without loss of generality, there must exist some α such that $M_i \cap M_k = \sigma\beta$. Therefore, as $\mathcal{C}$ is 3-sparse we know that $M_k = \sigma\beta\tau$, and is thus of size three, completing the proof. □

Theorem 6.4. *The neural code*

$$\mathcal{C} = \{\emptyset, 4, 5, 12, 14, 23, 35, 45, 46, 47, 56, 58, 123, 124, 235, 456, 467, 568, 4578\}$$

is a locally good code that is neither open convex nor closed convex.

Proof. What we have done here is taken a code that has been proved in [5] to not be open convex and placed it inside the code provided in Theorem 3.5. Thus the proofs for each of those codes still apply here. It follows that this code can be neither open convex nor closed convex. Ergo all that is left is to make sure that the code is locally good. The only intersection of facets that are not contained in the code are 2 and 5. It can be verified that neither of these codewords are mandatory. Thus $\mathcal{C}$ is a locally good code that is neither open convex nor closed convex. □

Conjecture 6.5. *Let $\mathcal{C}$ be a locally good neural code on n neurons. If $n \leq 7$, then $\mathcal{C}$ must be either open or closed convex.*

All of the results and codes found to date suggest that for a code to be closed convex but not open convex (or vice-versa) its visualization must contain (or obtain by adding in boundaries) a set that exists in two dimensions below that of the ambient space. The code in Theorem 6.4 was the smallest code that we could find that was neither open convex nor closed convex. Looking at the code in Theorem 6.4 it is difficult to imagine a smaller code with the same properties, but such is the case with all codes until they are discovered.

Conjecture 6.6. *Let $\mathcal{C}$ be a locally good neural code on n neurons. If $\mathcal{C}$ is not open convex, then any convex realization of $\mathcal{C}$ in $\mathbb{R}^d$ must contain a set that can only be realized in $\mathbb{R}^{d-2}$ or below.*

It was shown in [6] that every neural code can be expressed convexly if we do not specify that all sets must be either open or closed. However, this conjecture would introduce some situations in which there are specifications on what these convex realizations would look like.

Conjecture 6.7. *Let $\mathcal{C}$ be a closed convex neural code on n neurons. Let $\mathcal{U} = \{\mathcal{U}_i\}_{i=1}^n$ in $\mathbb{R}^d$ be an arbitrary open convex realization of $\mathcal{C}$. If filling in the boundary of each $\mathcal{U}_i \in \mathcal{U}$ creates a region corresponding to a codeword that is $(d-2)$-dimensional or less, then $\mathcal{C}$ is not open convex.*

7 Acknowledgements

We would like to sincerely thank Dr. Anne Shiu for her invaluable assistance and mentorship. This paper would not have been possible without her trust and dedication. Additionally, we would like to thank Thomas Yahl for staring at the board with us and Nida Obatake for emotionally supporting us through the trials and tribulations of research.

This research was completed in the 2019 REU in the Department of Mathematics at Texas A&M University, supported by NSF grant DMS-1757872.

References

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

[2] 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.

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

[4] Sarah Ayman Goldrup and Kaitlyn Phillipson. Classification of open and closed convex codes on five neurons. *Preprint*, 2019.

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

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