

What makes a neural code convex?

Carina Curto¹, Elizabeth Gross², Jack Jeffries^{3,4}, Katherine Morrison^{1,5}, Mohamed Omar⁶,
Zvi Rosen^{1,7,8}, Anne Shiu⁹, and Nora Youngs^{6,10}

Nov 1, 2016

¹ Department of Mathematics, The Pennsylvania State University, University Park, PA 16802

² Department of Mathematics, San José State University, San José, CA 95192

³ Department of Mathematics, University of Utah, Salt Lake City, UT 84112

⁴ Department of Mathematics, University of Michigan, Ann Arbor, MI 48109

⁵ School of Mathematical Sciences, University of Northern Colorado, Greeley, CO 80639

⁶ Department of Mathematics, Harvey Mudd College, Claremont, CA 91711

⁷ Department of Mathematics, University of California, Berkeley, Berkeley, CA 94720

⁸ Department of Mathematics, University of Pennsylvania, Philadelphia, PA 19104

⁹ Department of Mathematics, Texas A&M University, College Station, TX 77843

¹⁰ Department of Mathematics, Colby College, Waterville, ME 04901

Abstract

Neural codes allow the brain to represent, process, and store information about the world. Combinatorial codes, comprised of binary patterns of neural activity, encode information via the collective behavior of populations of neurons. A code is called *convex* if its codewords correspond to regions defined by an arrangement of convex open sets in Euclidean space. Convex codes have been observed experimentally in many brain areas, including sensory cortices and the hippocampus, where neurons exhibit convex receptive fields. What makes a neural code convex? That is, how can we tell from the intrinsic structure of a code if there exists a corresponding arrangement of convex open sets? In this work, we provide a complete characterization of *local obstructions* to convexity. This motivates us to define *max intersection-complete codes*, a family guaranteed to have no local obstructions. We then show how our characterization enables one to use free resolutions of Stanley-Reisner ideals in order to detect violations of convexity. Taken together, these results provide a significant advance in understanding the intrinsic combinatorial properties of convex codes.

Keywords: neural coding, convex codes, simplicial complex, link, Nerve lemma, Hochster’s formula

Contents

1	Introduction	2
1.1	Convex codes	2
1.2	Preliminaries	3
1.3	Summary of main results	4
2	Examples	5
3	Local obstructions to convexity	7
3.1	Receptive field relationships	8
3.2	The Nerve lemma	8
3.3	Local obstructions	9
4	Characterizing local obstructions via mandatory codewords	11
4.1	Link lemmas	11
4.2	Proof of Theorem 1.3 and Lemma 1.4	12
5	Computing mandatory codewords algebraically	13
6	Supplementary Text	17

1 Introduction

Cracking the neural code is one of the central challenges of neuroscience. Typically, this has been understood as finding the relationship between the activity of neurons and the stimuli they represent. To uncover the principles of neural coding, however, it is not enough to describe the various mappings between stimulus and response. One must also understand the intrinsic structure of neural codes, independently of what is being encoded [1].

Here we focus our attention on *convex codes*, which are comprised of activity patterns for neurons with classical receptive fields. A *receptive field* U_i is the subset of stimuli that induces neuron i to respond. Often, U_i is a convex subset of some Euclidean space (see Figure 1). A collection of convex sets $U_1, \dots, U_n \subset \mathbb{R}^d$ naturally associates to each point $x \in \mathbb{R}^d$ a binary response pattern, $c_1 \cdots c_n \in \{0, 1\}^n$, where $c_i = 1$ if $x \in U_i$, and $c_i = 0$ otherwise. The set of all such response patterns is a convex code.

Convex codes have been observed experimentally in many brain areas, including sensory cortices and the hippocampus. Hubel and Wiesel’s discovery in 1959 of orientation-tuned neurons in the primary visual cortex was perhaps the first example of convex coding in the brain [2]. This was followed by O’Keefe’s discovery of hippocampal place cells in 1971 [3], showing that convex codes are also used in the brain’s representation of space. Both discoveries were groundbreaking for neuroscience, and were later recognized with Nobel Prizes in 1981 [4] and 2014 [5], respectively.

Our motivating examples of convex codes are, in fact, hippocampal place cell codes. A *place cell* is a neuron that acts as a position sensor, exhibiting a high firing rate when the animal’s location lies inside the cell’s preferred region of the environment – its *place field*. Figure 1 displays the place fields of four place cells recorded while a rat explored a two-dimensional environment. Each place field is an approximately convex subset of \mathbb{R}^2 . Taken together, the set of all neural response patterns that can arise in a population of place cells comprise a convex code for the animal’s position in space.




Figure 1: Place fields of four CA1 pyramidal neurons (place cells) in rat hippocampus, recorded while the animal explored a $1.5\text{m} \times 1.5\text{m}$ square box environment. Red areas correspond to regions of space where the corresponding place cells exhibited high firing rates, while dark blue denotes near-zero activity. Place fields were computed from data provided by the Pastalkova lab, as described in [6].

Despite their relevance to neuroscience, the mathematical theory of convex codes was initiated only recently [1, 7]. In particular, it is not clear what are the intrinsic combinatorial signatures of convex and non-convex codes. Identifying such features will enable us to infer coding properties from population recordings of neural activity, without needing *a priori* knowledge of the stimuli being encoded. Understanding the structure of convex codes is also essential to uncovering the basic principles of how neural networks are organized in order to learn, store, and process information.

1.1 Convex codes

By a *neural code*, or simply a *code*, on n neurons we mean a collection of binary strings $\mathcal{C} \subseteq \{0, 1\}^n$. The elements of a code are called *codewords*. We interpret each binary digit as the “on” or “off” state of a neuron, and consider 0/1 strings of length n and subsets of $[n] = \{1, \dots, n\}$ interchangeably. For example, 1101 and 0100 are also denoted $\{1, 2, 4\}$ and $\{2\}$, respectively. We will always assume

$00 \cdots 0 \in \mathcal{C}$ (i.e., $\emptyset \in \mathcal{C}$); this assumption simplifies notation in various places but does not alter the core results.

Let X be a topological space, and consider a collection of open sets $\mathcal{U} = \{U_1, \dots, U_n\}$, where $\bigcup_{i=1}^n U_i \subsetneq X$. Any such \mathcal{U} defines a code,

$$\mathcal{C}(\mathcal{U}) \stackrel{\text{def}}{=} \left\{ \sigma \subseteq [n] \mid U_\sigma \setminus \bigcup_{j \in [n] \setminus \sigma} U_j \neq \emptyset \right\}, \quad (1)$$

where $U_\sigma \stackrel{\text{def}}{=} \bigcap_{i \in \sigma} U_i$ for $\sigma \subseteq [n]$. In other words, each codeword σ in $\mathcal{C}(\mathcal{U})$ corresponds to the portion of U_σ that is not covered by other sets. In particular, $\mathcal{C}(\mathcal{U})$ is not the same as the *nerve* of the cover, $\mathcal{N}(\mathcal{U})$ (see Section 3.2). By convention, $U_\emptyset = X$, and so $\emptyset \in \mathcal{C}(\mathcal{U})$.

For any code \mathcal{C} , there always exists an open cover \mathcal{U} such that $\mathcal{C} = \mathcal{C}(\mathcal{U})$ [1, Lemma 2.1]. However, it may be impossible to choose the U_i s to all be convex. We thus have the following definitions, which were first introduced in [1]:

Definition 1.1. Let \mathcal{C} be a binary code on n neurons. If there exists a collection of open sets $\mathcal{U} = \{U_1, \dots, U_n\}$ such that $\mathcal{C} = \mathcal{C}(\mathcal{U})$ and the U_i s are all convex subsets of \mathbb{R}^d , then we say that \mathcal{C} is a *convex code*. The smallest d such that this is possible is called the *minimal embedding dimension*, $d(\mathcal{C})$.

Note that the definition of a convex code is extrinsic: a code is convex if it can be realized by an arrangement of convex open sets in some Euclidean space. How can we characterize convex codes *intrinsically*? If a code is *not* convex, how can we prove this? If a code *is* convex, what is the minimal dimension needed for the corresponding open sets?

In this work, we tackle these questions by building on mathematical ideas from [1] and [8]. In particular, we study *local obstructions* to convexity, a notion first introduced in [8]. Our main result is Theorem 1.3, which provides a complete characterization of codes with no local obstructions. In Section 2 we present a series of examples that illustrate the ideas summarized in Section 1.3. Sections 3 and 4 are devoted to additional background and technical results needed for the proof of Theorem 1.3. Finally, in Section 5 we show how tools from combinatorial commutative algebra, such as Hochster’s formula, can be used to determine that a code is not convex.

1.2 Preliminaries

Simplicial complexes. A *simplicial complex* Δ on n vertices is a nonempty collection of subsets of $[n]$ that is closed under inclusion. In other words, if $\sigma \in \Delta$ and $\tau \subset \sigma$, then $\tau \in \Delta$. The elements of a simplicial complex are called *simplices* or *faces*. The *dimension* of a face, $\sigma \in \Delta$, is defined to be $|\sigma| - 1$. The dimension of a simplicial complex Δ is equal to the dimension of its largest face: $\max_{\sigma \in \Delta} |\sigma| - 1$. If Δ consists of all 2^n subsets of $[n]$, then it is the *full simplex* of dimension $n - 1$. The *hollow simplex* contains all proper subsets of $[n]$, but not $[n]$, and thus has dimension $n - 2$.

Faces of a simplicial complex that are maximal under inclusion are referred to as *facets*. If we consider the facets together with all their intersections, we obtain the set

$$\mathcal{F}_\cap(\Delta) \stackrel{\text{def}}{=} \left\{ \bigcap_{i=1}^k \rho_i \mid \rho_i \text{ is a facet of } \Delta \text{ for each } i = 1, \dots, k \right\} \cup \{\emptyset\}.$$

We refer to the elements of $\mathcal{F}_\cap(\Delta)$ as *max intersections* of Δ . The empty set is added so that $\mathcal{F}_\cap(\Delta)$ can be regarded as a code, consistent with our convention that the all-zeros word is always included.

Restrictions and links are standard constructions from simplicial complexes. The *restriction* of Δ to σ is the simplicial complex

$$\Delta|_\sigma \stackrel{\text{def}}{=} \{\omega \in \Delta \mid \omega \subseteq \sigma\}.$$

For any $\sigma \in \Delta$, the *link* of σ inside Δ is

$$\text{Lk}_\sigma(\Delta) \stackrel{\text{def}}{=} \{\omega \in \Delta \mid \sigma \cap \omega = \emptyset \text{ and } \sigma \cup \omega \in \Delta\}.$$

Note that it is more common to write $\text{Lk}_\Delta(\sigma)$ or $\text{link}_\Delta(\sigma)$, instead of $\text{Lk}_\sigma(\Delta)$ (see, for example, [9]). However, because we will often fix σ and consider its link inside different simplicial complexes, such as $\Delta|_{\sigma \cup \tau}$, it is more convenient to put σ in the subscript.

The simplicial complex of a code. To any code \mathcal{C} , we can associate a simplicial complex $\Delta(\mathcal{C})$ by simply including all subsets of codewords:

$$\Delta(\mathcal{C}) \stackrel{\text{def}}{=} \{\sigma \subseteq [n] \mid \sigma \subseteq c \text{ for some } c \in \mathcal{C}\}.$$

$\Delta(\mathcal{C})$ is called the *simplicial complex of the code*, and is the smallest simplicial complex that contains all elements of \mathcal{C} . The facets of $\Delta(\mathcal{C})$ correspond to the codewords in \mathcal{C} that are maximal under inclusion: these are the *maximal codewords*.

Local obstructions to convexity. At first glance, it may seem that all codes should be convex, since the convex sets U_i can be chosen to reside in arbitrarily high dimensions. This is not the case, however, as non-convex codes arise for as few as $n = 3$ neurons [1]. To understand what can go wrong, consider a code with the following property: any codeword with a 1 in the first position also has a 1 in the second or third position, but no codeword has a 1 in all three positions. This implies that any corresponding cover \mathcal{U} must have $U_1 \subseteq U_2 \cup U_3$, but $U_1 \cap U_2 \cap U_3 = \emptyset$. The result is that U_1 is a disjoint union of two nonempty open sets, $U_1 \cap U_2$ and $U_1 \cap U_3$, and is hence disconnected. Since all convex sets are connected, we conclude that our code cannot be convex. The contradiction stems from a topological inconsistency that emerges if the code is assumed to be convex.

This type of topological obstruction to convexity generalizes to a family of *local obstructions*, first introduced in [8]. We define local obstructions precisely in Section 3. There we also show that a code with one or more local obstructions cannot be convex:

Lemma 1.2. If \mathcal{C} is a convex code, then \mathcal{C} has no local obstructions.

This fact was first observed in [8], using slightly different language. The converse, unfortunately, is not true. See Example 2.3 for a counterexample that first appeared in [10].

1.3 Summary of main results

To prove that a neural code is convex, it suffices to exhibit a convex realization. That is, it suffices to find a set of convex open sets $\mathcal{U} = \{U_1, \dots, U_n\}$ such that $\mathcal{C} = \mathcal{C}(\mathcal{U})$. Our strategy for proving that a code is *not* convex is to show that it has a local obstruction to convexity. Which codes have local obstructions?

Perhaps surprisingly, the question of whether or not a given code \mathcal{C} has a local obstruction can be reduced to the question of whether or not it contains a certain minimal code, $\mathcal{C}_{\min}(\Delta)$, which depends only on the simplicial complex $\Delta = \Delta(\mathcal{C})$. This is our main result:

Theorem 1.3 (Characterization of local obstructions). For each simplicial complex Δ , there is a unique minimal code $\mathcal{C}_{\min}(\Delta)$ with the following properties:

- (i) the simplicial complex of $\mathcal{C}_{\min}(\Delta)$ is Δ , and
- (ii) for any code \mathcal{C} with simplicial complex Δ , \mathcal{C} has no local obstructions if and only if $\mathcal{C} \supseteq \mathcal{C}_{\min}(\Delta)$.

Moreover, $\mathcal{C}_{\min}(\Delta)$ depends only on the topology of the links of Δ :

$$\mathcal{C}_{\min}(\Delta) = \{\sigma \in \Delta \mid \text{Lk}_\sigma(\Delta) \text{ is non-contractible}\} \cup \{\emptyset\}.$$

We will regard the elements of $\mathcal{C}_{\min}(\Delta)$ as “mandatory” codewords with respect to convexity, because they must all be included in order for a code \mathcal{C} with $\Delta(\mathcal{C}) = \Delta$ to be convex. From the above description of $\mathcal{C}_{\min}(\Delta)$, we can prove the following lemma.

Lemma 1.4. $\mathcal{C}_{\min}(\Delta) \subseteq \mathcal{F}_{\cap}(\Delta)$. That is, each nonempty element of $\mathcal{C}_{\min}(\Delta)$ is an intersection of facets of Δ .

Our proofs of Theorem 1.3 and Lemma 1.4 are given in Section 4.2. Unfortunately, finding all elements of $\mathcal{C}_{\min}(\Delta)$ for arbitrary Δ is, in general, undecidable (see Section 5). Nevertheless, we can algorithmically compute a subset, $\mathcal{M}_H(\Delta)$, of “homologically-detectable” mandatory codewords. In Section 5 we show how to compute $\mathcal{M}_H(\Delta)$ using machinery from combinatorial commutative algebra. Lemma 1.4 also tells us that every element of $\mathcal{C}_{\min}(\Delta)$ must be an intersection of facets of Δ – that is, an element of $\mathcal{F}_{\cap}(\Delta)$. We thus have the inclusions:

$$\mathcal{M}_H(\Delta) \subseteq \mathcal{C}_{\min}(\Delta) \subseteq \mathcal{F}_{\cap}(\Delta), \quad (2)$$

where both $\mathcal{M}_H(\Delta)$ and $\mathcal{F}_{\cap}(\Delta)$ are straightforwardly computable. Note that if $\mathcal{M}_H(\Delta) = \mathcal{F}_{\cap}(\Delta)$, then we can conclude that $\mathcal{C}_{\min}(\Delta) = \mathcal{F}_{\cap}(\Delta)$. Moreover, if $\mathcal{C} \supseteq \mathcal{F}_{\cap}(\Delta)$, then $\mathcal{C} \supseteq \mathcal{C}_{\min}(\Delta)$, and thus \mathcal{C} has no local obstructions (and is potentially convex). This motivates the following definition:

Definition 1.5. A neural code \mathcal{C} is *max \cap -complete* (or *max intersection-complete*) if $\mathcal{C} \supseteq \mathcal{F}_{\cap}(\Delta(\mathcal{C}))$.

We therefore have a simple combinatorial condition for a code that guarantees it has no local obstructions:

Corollary 1.6. If a neural code \mathcal{C} is max \cap -complete, then \mathcal{C} has no local obstructions.

For $n \leq 4$, the convex codes are precisely those codes that are max \cap -complete.

Proposition 1.7. Let \mathcal{C} be a code on $n \leq 4$ neurons. Then \mathcal{C} is convex if and only if \mathcal{C} is max \cap -complete.

This is shown in Supplementary Text S1, where we provide a complete classification of convex codes on $n = 4$ neurons. Proposition 1.7 does not extend to $n > 4$, however, since beginning in $n = 5$ there are convex codes that are not max \cap -complete (see Example 2.2). This raises the question: are there max \cap -complete codes that are not convex? In a previous version of this paper, we conjectured that all max \cap -complete codes are convex [11]. This conjecture has recently been proven [12], using ideas similar to what we illustrate in Example 2.4.

Proposition 1.8 ([12, Theorem 4.4]). If \mathcal{C} is max \cap -complete, then \mathcal{C} is convex.

Finally, in Supplementary Text S3 we present some straightforward bounds on the minimal embedding dimension $d(\mathcal{C})$, obtained using results about d -representability of the associated simplicial complex $\Delta(\mathcal{C})$. In particular, we find bounds from Helly’s theorem and the Fractional Helly theorem. Unfortunately, these results all stem from $\Delta(\mathcal{C})$. In our classification of convex codes for $n \leq 4$, however, it is clear that the presence or absence of specific codewords can affect $d(\mathcal{C})$, even if $\Delta(\mathcal{C})$ remains unchanged (see Table 2 in Supplementary Text S1). The problem of how to use this additional information about a code in order to improve the bounds on $d(\mathcal{C})$ remains wide open.

2 Examples

Our first example depicts a convex code with minimal embedding dimension $d(\mathcal{C}) = 2$.




Figure 2: **(a)** An arrangement $\mathcal{U} = \{U_1, U_2, U_3, U_4\}$ of convex open sets. Black dots mark regions corresponding to distinct codewords in $\mathcal{C} = \mathcal{C}(\mathcal{U})$. From left to right, the codewords are 0000, 1000, 1100, 1010, 1110, 0110, 0010, 0111, 0011, and 0001. **(b)** The simplicial complex $\Delta(\mathcal{C})$ for the code \mathcal{C} defined in (a). The two facets, 123 and 234, correspond to the two maximal codewords, 1110 and 0111, respectively. This simplicial complex is also equal to the nerve of the cover $\mathcal{N}(\mathcal{U})$ (see Section 3.2).

Example 2.1. Consider the open cover \mathcal{U} illustrated in Figure 2a. The corresponding code, $\mathcal{C} = \mathcal{C}(\mathcal{U})$, has 10 codewords. \mathcal{C} is a convex code, by construction, and it is easy to check that $d(\mathcal{C}) = 2$. The simplicial complex $\Delta(\mathcal{C})$ (see Figure 2b) loses some of the information about the cover \mathcal{U} that is present in \mathcal{C} . In particular, $U_2 \subseteq U_1 \cup U_3$ and $U_2 \cap U_4 \subseteq U_3$ is reflected in \mathcal{C} , but not in $\Delta(\mathcal{C})$. Note that we can infer $U_2 \subseteq U_1 \cup U_3$ directly from the code, because any codeword with neuron 2 “on” also has neuron 1 or 3 “on.”

Note that the convex code in Example 2.1 is also max \cap -complete, as guaranteed by Proposition 1.7. The next example shows that this proposition does not hold for $n \geq 5$.




Figure 3: **(a)** A simplicial complex Δ on $n = 5$ vertices. The vertex 1 is an intersection of facets, but is not contained in the code \mathcal{C} of Example 2.2. **(b)** The link $\text{Lk}_1(\Delta)$ (see Section 3.3). **(c)** A convex realization of the code \mathcal{C} . The set U_1 corresponding to neuron 1 (shaded) is completely covered by the other sets U_2, \dots, U_5 , consistent with the fact that $1 \notin \mathcal{C}$.

Example 2.2. The simplicial complex Δ shown in Figure 3a has facets 123, 134, and 145. Their intersections yield the faces 1, 13, and 14, so that $\mathcal{F}_\cap(\Delta) = \{123, 134, 145, 13, 14, 1, \emptyset\}$. For this Δ , we can compute the minimal code with no local obstructions, $\mathcal{C}_{\min}(\Delta) = \{123, 134, 145, 13, 14, \emptyset\}$, as in Theorem 1.3. Note that the element $1 \in \mathcal{F}_\cap(\Delta)$ is not present in $\mathcal{C}_{\min}(\Delta)$. Now consider the code $\mathcal{C} = \Delta \setminus \{1\}$. Clearly, this code has simplicial complex $\Delta(\mathcal{C}) = \Delta$: it has a codeword for each face of Δ , except the vertex 1 (see Figure 3a). By Theorem 1.3, \mathcal{C} has no local obstructions because $\mathcal{C} \supseteq \mathcal{C}_{\min}(\Delta)$. However, \mathcal{C} is not max \cap -complete because $\mathcal{F}_\cap(\Delta) \not\subseteq \mathcal{C}$. Nevertheless, \mathcal{C} is convex. A convex realization is shown in Figure 3c.

The absence of local obstructions is a necessary condition for convexity. Unfortunately, it is not sufficient: the following example shows a code with no local obstructions that is *not* convex.

Example 2.3 ([10]). Consider the code $\mathcal{C} = \{2345, 123, 134, 145, 13, 14, 23, 34, 45, 3, 4, \emptyset\}$. The simplicial complex of this code, $\Delta = \Delta(\mathcal{C})$, has facets $\{2345, 123, 134, 145\}$. It is straightforward to show that $\mathcal{C} = \mathcal{C}_{\min}(\Delta)$, and thus \mathcal{C} has no local obstructions. Despite this, it was shown in [10] using geometric arguments that \mathcal{C} is *not* convex. Note that this code is not max \cap -complete (the max intersection $123 \cap 134 \cap 145 = 1$ is not in \mathcal{C}).

The next example illustrates how a code with a single maximal codeword can be embedded in \mathbb{R}^2 . This basic construction is used repeatedly in our proof of Proposition 1.7 (see Supplementary Text S1), and inspired aspects of the proof of Proposition 1.8, given in [12].

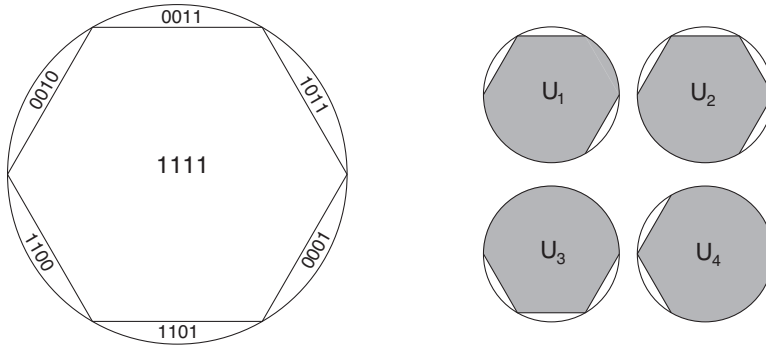


Figure 4: A convex realization in \mathbb{R}^2 of the code in Example 2.4. (Left) Each non-maximal codeword is assigned a region outside the polygon, but inside the disk. (Right) For each neuron i , the convex set U_i is the union of all regions corresponding to codewords with a 1 in the i^{th} position.

Example 2.4. Consider the code $\mathcal{C} = \{1111, 1011, 1101, 1100, 0011, 0010, 0001, 0000\}$, with unique maximal codeword 1111. Figure 4 depicts the construction of a convex realization in \mathbb{R}^2 . All regions corresponding to codewords are subsets of a disk in \mathbb{R}^2 . For each $i = 1, \dots, 4$, the convex set U_i is the union of all regions whose corresponding codewords have a 1 in the i^{th} position. For example, U_1 is the union of the four regions corresponding to codewords 1111, 1011, 1101, and 1100.

The above construction can be generalized to any code with a unique maximal codeword.

Lemma 2.5. Let \mathcal{C} be a code with a unique maximal codeword. Then \mathcal{C} is convex, and $d(\mathcal{C}) \leq 2$.

Proof. Let $\rho \in \mathcal{C}$ be the unique maximal codeword, and let $m = |\mathcal{C}| - 2$ be the number of non-maximal codewords, excluding the all-zeros word. Inscribe a regular open m -gon P in an open disk, so that there are m sectors surrounding P , as in Figure 4. (If $m < 3$, let P be an open triangle.) Assign each non-maximal codeword (excluding $00 \dots 0$) to a distinct sector inside the disk but outside of P , and assign the maximal codeword ρ to P . Next, for each $i \in \rho$ let U_i be the union of P and all sectors whose corresponding codewords have a 1 in the i^{th} position, together with their common boundaries with P . For $j \in [n] \setminus \rho$, set $U_j = \emptyset$. Note that each U_i is open and convex, and $\mathcal{C} = \mathcal{C}(\{U_i\})$. \square

Lemma 2.5 can easily be generalized to any code whose maximal codewords are non-overlapping (that is, having disjoint supports). In this case, each nonzero codeword is contained in a unique facet of $\Delta(\mathcal{C})$, and the facets thus yield a partition of the code. We can repeat the above construction in parallel for each part, obtaining the same dimension bound.

Proposition 2.6. Let \mathcal{C} be a code with non-overlapping maximal codewords (i.e., the facets of $\Delta(\mathcal{C})$ are disjoint). Then \mathcal{C} is convex and $d(\mathcal{C}) \leq 2$.

3 Local obstructions to convexity

For any simplicial complex Δ , there exists a convex cover \mathcal{U} in a high-enough dimensional space \mathbb{R}^d such that Δ can be realized as $\Delta(\mathcal{C}(\mathcal{U}))$ [13]. For this reason, the simplicial complex $\Delta(\mathcal{C})$ alone is not sufficient to determine whether or not \mathcal{C} is convex. Obstructions to convexity must emerge from information in the code that goes beyond what is reflected in $\Delta(\mathcal{C})$. As was shown in [1], this additional information is precisely the receptive field relationships, which we turn to now.

3.1 Receptive field relationships

For a code \mathcal{C} on n neurons, let $\mathcal{U} = \{U_1, \dots, U_n\}$ be any collection of open sets such that $\mathcal{C} = \mathcal{C}(\mathcal{U})$, and recall that $U_\sigma = \bigcap_{i \in \sigma} U_i$.

Definition 3.1. A *receptive field relationship* (RF relationship) of \mathcal{C} is a pair (σ, τ) corresponding to the set containment

$$U_\sigma \subseteq \bigcup_{i \in \tau} U_i,$$

where $\sigma \neq \emptyset$, $\sigma \cap \tau = \emptyset$, and $U_\sigma \cap U_i \neq \emptyset$ for all $i \in \tau$.

If $\tau = \emptyset$, then the relationship (σ, \emptyset) simply states that $U_\sigma = \emptyset$. Note that relationships of the form (σ, \emptyset) reproduce the information in $\Delta(\mathcal{C})$, while those of the form (σ, τ) for $\tau \neq \emptyset$ reflect additional structure in \mathcal{C} that goes beyond the simplicial complex. A *minimal* RF relationship is one such that no single neuron can be removed from σ or τ without destroying the containment.

It is important to note that RF relationships are independent of the choice of open sets \mathcal{U} (see Lemma 4.2 of [1]). Hence we denote the set of all RF relationships $\{(\sigma, \tau)\}$ for a given code \mathcal{C} as simply $\text{RF}(\mathcal{C})$. In [1], it was shown that one can compute $\text{RF}(\mathcal{C})$ algebraically, using an associated ideal $I_{\mathcal{C}}$.

Example 3.2 (Example 2.1 continued). The code $\mathcal{C} = \mathcal{C}(\mathcal{U})$ from Example 2.1 has RF relationships $\text{RF}(\mathcal{C}) = \{(\{1, 4\}, \emptyset), (\{1, 2, 4\}, \emptyset), (\{1, 3, 4\}, \emptyset), (\{1, 2, 3, 4\}, \emptyset), (\{2\}, \{1, 3\}), (\{2\}, \{1, 3, 4\}), (\{2, 4\}, \{3\})\}$. Of these, the pairs $(\{1, 4\}, \emptyset)$, $(\{2\}, \{1, 3\})$, and $(\{2, 4\}, \{3\})$, corresponding to $U_1 \cap U_4 = \emptyset$, $U_2 \subseteq U_1 \cup U_3$, and $U_2 \cap U_4 \subseteq U_3$ respectively, are the minimal RF relationships.

The following lemma illustrates a simple case where RF relationships can be used to show that a code cannot have a convex realization. (This is a special case of Lemma 3.6 below.)

Lemma 3.3. Let $\mathcal{C} = \mathcal{C}(\mathcal{U})$. If \mathcal{C} has RF relationships $U_\sigma \subseteq U_i \cup U_j$ and $U_\sigma \cap U_i \cap U_j = \emptyset$ for some $\sigma \subseteq [n]$ and distinct $i, j \notin \sigma$, then \mathcal{C} is not a convex code.

Proof. By assumption, $\{(\sigma, \{i, j\}), (\sigma \cup \{i, j\}, \emptyset)\} \subseteq \text{RF}(\mathcal{C})$. It follows that the sets $V_i = U_\sigma \cap U_i \neq \emptyset$ and $V_j = U_\sigma \cap U_j \neq \emptyset$ are disjoint open sets that each intersect U_σ , and $U_\sigma \subseteq V_i \cup V_j$. We can thus conclude that U_σ is disconnected in any open cover \mathcal{U} such that $\mathcal{C} = \mathcal{C}(\mathcal{U})$. This implies that \mathcal{C} cannot have a convex realization, because if the U_i s were all convex then U_σ would be convex, contradicting the fact that it is disconnected. \square

The above proof relies on the observation that U_σ must be convex in any convex realization \mathcal{U} , but the properties of the code imply that U_σ is covered by a collection of open sets whose topology does not match that of a convex set. This topological inconsistency between a set and its cover is, at its core, a contradiction arising from the Nerve lemma, which we discuss next.

3.2 The Nerve lemma

The *nerve* of an open cover $\mathcal{U} = \{U_1, \dots, U_n\}$ is the simplicial complex

$$\mathcal{N}(\mathcal{U}) \stackrel{\text{def}}{=} \{\sigma \subseteq [n] \mid U_\sigma \neq \emptyset\}.$$

In fact, $\mathcal{N}(\mathcal{U}) = \Delta(\mathcal{C}(\mathcal{U}))$, so the nerve can be recovered directly from the code $\mathcal{C}(\mathcal{U})$. The Nerve lemma tells us that $\mathcal{N}(\mathcal{U})$ carries a surprising amount of topological information about the underlying space covered by \mathcal{U} , provided \mathcal{U} is a good cover. Recall that a *good cover* is a collection of open sets $\{U_i\}$ where every non-empty intersection, $U_\sigma = \bigcap_{i \in \sigma} U_i$, is contractible.¹

¹A set is *contractible* if it is homotopy-equivalent to a point [14].

Lemma 3.4 (Nerve lemma). If \mathcal{U} is a good cover, then $\bigcup_{i=1}^n U_i$ is homotopy-equivalent to $\mathcal{N}(\mathcal{U})$. In particular, $\bigcup_{i=1}^n U_i$ and $\mathcal{N}(\mathcal{U})$ have exactly the same homology groups.

This result is well known, and can be obtained as a direct consequence of [14, Corollary 4G.3].

Now observe that an open cover comprised of convex sets is always a good cover, because the intersection of convex sets is convex, and hence contractible. For example, if $\mathcal{C} = \mathcal{C}(\mathcal{U})$ for a convex cover \mathcal{U} , then $\Delta(\mathcal{C})$ must match the homotopy type of $\bigcup_{i=1}^n U_i$. This fact was previously exploited to extract topological information about the represented space from hippocampal place cell activity [15].

The Nerve lemma is also key to our notion of local obstructions, which we turn to next.

3.3 Local obstructions

Local obstructions arise when a code contains a RF relationship (σ, τ) , so that $U_\sigma \subseteq \bigcup_{i \in \tau} U_i$, but the nerve of the corresponding cover of U_σ by the restricted sets $\{U_i \cap U_\sigma\}_{i \in \tau}$ is not contractible. By the Nerve lemma, if the U_i s are all convex then $\mathcal{N}(\{U_\sigma \cap U_i\}_{i \in \tau})$ must have the same homotopy type as U_σ , which is contractible. If $\mathcal{N}(\{U_\sigma \cap U_i\}_{i \in \tau})$ fails to be contractible, we can conclude that the U_i s can not all be convex.

Now, observe that the nerve of the restricted cover $\mathcal{N}(\{U_\sigma \cap U_i\}_{i \in \tau})$ is related to the nerve of the original cover $\mathcal{N}(\mathcal{U})$ as follows:

$$\mathcal{N}(\{U_\sigma \cap U_i\}_{i \in \tau}) = \{\omega \in \mathcal{N}(\mathcal{U}) \mid \sigma \cap \omega = \emptyset, \sigma \cup \omega \in \mathcal{N}(\mathcal{U}), \text{ and } \omega \subseteq \tau\}.$$

In fact, letting $\Delta = \mathcal{N}(\mathcal{U})$ and considering the restricted complex $\Delta|_{\sigma \cup \tau}$, we recognize that the right-hand side above is precisely the link,

$$\mathcal{N}(\{U_\sigma \cap U_i\}_{i \in \tau}) = \text{Lk}_\sigma(\Delta|_{\sigma \cup \tau}).$$

We can now define a *local obstruction* to convexity.

Definition 3.5. Let $(\sigma, \tau) \in \text{RF}(\mathcal{C})$, and let $\Delta = \Delta(\mathcal{C})$. We say that (σ, τ) is a *local obstruction* of \mathcal{C} if $\tau \neq \emptyset$ and $\text{Lk}_\sigma(\Delta|_{\sigma \cup \tau})$ is *not* contractible.

Local obstructions are thus detected via non-contractible links of the form $\text{Lk}_\sigma(\Delta|_{\sigma \cup \tau})$, where $(\sigma, \tau) \in \text{RF}(\mathcal{C})$. Figure 5 displays all possible links that can arise for $|\tau| \leq 4$. Non-contractible links are highlighted in red. Note also that $\tau \neq \emptyset$ implies $\sigma \notin \mathcal{C}$ and $U_\sigma \neq \emptyset$, as the definition of a RF relationship requires that $U_\sigma \cap U_i \neq \emptyset$ for all $i \in \tau$. Any local obstruction (σ, τ) must therefore have $\sigma \in \Delta(\mathcal{C}) \setminus \mathcal{C}$ and $\text{Lk}_\sigma(\Delta|_{\sigma \cup \tau})$ nonempty.

The arguments leading up to the definition of local obstruction imply the following simple consequence of the Nerve lemma, which was previously observed in [8].

Lemma 3.6 (Lemma 1.2). If \mathcal{C} has a local obstruction, then \mathcal{C} is not a convex code.

In general, the question of whether or not a given simplicial complex is contractible is undecidable [16]; however, in some cases it is easy to see that all relevant links will be contractible. This yields a simple condition on RF relationships that guarantees that a code has no local obstructions.

Lemma 3.7. Let $\mathcal{C} = \mathcal{C}(\mathcal{U})$. If for each $(\sigma, \tau) \in \text{RF}(\mathcal{C})$ we have $U_\sigma \cap U_\tau \neq \emptyset$, then \mathcal{C} has no local obstructions.

Proof. Let $\Delta = \Delta(\mathcal{C})$. $U_\sigma \cap U_\tau \neq \emptyset$ implies $\text{Lk}_\sigma(\Delta|_{\sigma \cup \tau})$ is the full simplex on the vertex set τ , which is contractible. If this is true for every RF relationship, then none can give rise to a local obstruction. \square

For example, if $11 \cdots 1 \in \mathcal{C}$, then $U_\sigma \cap U_\tau \neq \emptyset$ for any pair $\sigma, \tau \subset [n]$, so \mathcal{C} has no local obstructions.




Figure 5: All simplicial complexes on up to $n = 4$ vertices, up to permutation equivalence. These can each arise as links of the form $Lk_{\sigma}(\Delta|_{\sigma \cup \tau})$ for $|\tau| \leq 4$. Red labels correspond to non-contractible complexes. Note that L13 is the only simplicial complex on $n \leq 4$ vertices that is contractible but *not* a cone.

4 Characterizing local obstructions via mandatory codewords

From the definition of local obstruction, it seems that in order to show that a code has no local obstructions one would need to check the contractibility of *all* links of the form $\text{Lk}_\sigma(\Delta|_{\sigma\cup\tau})$ corresponding to all pairs $(\sigma, \tau) \in \text{RF}(\mathcal{C})$. We shall see in this section that in fact we only need to check for contractibility of links inside the full complex Δ – that is, links of the form $\text{Lk}_\sigma(\Delta)$. This is key to obtaining a list of mandatory codewords, $\mathcal{C}_{\min}(\Delta)$, that depends only on Δ , and not on any further details of the code.

In Section 4.1 we prove some important lemmas about links, and then use them in Section 4.2 to prove Theorem 1.3.

4.1 Link lemmas

In what follows, the notation

$$\text{cone}_v(\Delta) \stackrel{\text{def}}{=} \{\{v\} \cup \omega \mid \omega \in \Delta\} \cup \Delta$$

denotes the *cone* of v over Δ , where v is a new vertex not contained in Δ . Any simplicial complex that is a cone over a sub-complex, so that $\Delta = \text{cone}_v(\Delta')$, is automatically contractible. In Figure 5, the only contractible link that is *not* a cone is L13. This is the same link that appeared in Figure 3b of Example 2.2.

Lemma 4.1. Let Δ be a simplicial complex on $[n]$, $\sigma \in \Delta$, and $v \in [n]$ such that $v \notin \sigma$ and $\sigma \cup \{v\} \in \Delta$. Then $\text{Lk}_{\sigma \cup \{v\}}(\Delta) \subseteq \text{Lk}_\sigma(\Delta|_{[n] \setminus \{v\}})$, and

$$\text{Lk}_\sigma(\Delta) = \text{Lk}_\sigma(\Delta|_{[n] \setminus \{v\}}) \cup \text{cone}_v(\text{Lk}_{\sigma \cup \{v\}}(\Delta)).$$

Proof. The proof follows from the definition of the link. First, observe that

$$\begin{aligned} \text{Lk}_{\sigma \cup \{v\}}(\Delta) &= \{\omega \subset [n] \mid v \notin \omega, \omega \cap \sigma = \emptyset, \text{ and } \omega \cup \sigma \cup \{v\} \in \Delta\} \\ &= \{\omega \subset [n] \setminus \{v\} \mid \omega \cap \sigma = \emptyset, \text{ and } \omega \cup \sigma \cup \{v\} \in \Delta\} \\ &\subseteq \{\omega \subset [n] \setminus \{v\} \mid \omega \cap \sigma = \emptyset, \text{ and } \omega \cup \sigma \in \Delta|_{[n] \setminus \{v\}}\} \\ &= \text{Lk}_\sigma(\Delta|_{[n] \setminus \{v\}}), \end{aligned}$$

which establishes that $\text{Lk}_{\sigma \cup \{v\}}(\Delta) \subseteq \text{Lk}_\sigma(\Delta|_{[n] \setminus \{v\}})$. Next, observe that

$$\begin{aligned} \text{cone}_v(\text{Lk}_{\sigma \cup \{v\}}(\Delta)) \setminus \text{Lk}_{\sigma \cup \{v\}}(\Delta) &= \{\{v\} \cup \omega \mid \omega \in \text{Lk}_{\sigma \cup \{v\}}(\Delta)\} \\ &= \{\{v\} \cup \omega \mid \omega \subset [n], v \notin \omega, \omega \cap \sigma = \emptyset, \text{ and } \omega \cup \{v\} \cup \sigma \in \Delta\} \\ &= \{\tau \subset [n] \mid v \in \tau, \tau \cap \sigma = \emptyset, \text{ and } \tau \cup \sigma \in \Delta\} \\ &= \{\omega \in \text{Lk}_\sigma(\Delta) \mid v \in \omega\}. \end{aligned}$$

Finally,

$$\text{Lk}_\sigma(\Delta|_{[n] \setminus \{v\}}) = \{\omega \in \text{Lk}_\sigma(\Delta) \mid v \notin \omega\}.$$

From here the second statement is clear. \square

Corollary 4.2. Let Δ be a simplicial complex on $[n]$, $\sigma \in \Delta$, and $v \in [n]$ such that $v \notin \sigma$ and $\sigma \cup \{v\} \in \Delta$. If $\text{Lk}_{\sigma \cup \{v\}}(\Delta)$ is contractible, then $\text{Lk}_\sigma(\Delta)$ and $\text{Lk}_\sigma(\Delta|_{[n] \setminus \{v\}})$ are homotopy-equivalent.

Proof. Lemma 4.1 states that $\text{Lk}_{\sigma \cup \{v\}}(\Delta) \subseteq \text{Lk}_\sigma(\Delta|_{[n] \setminus \{v\}})$, and that $\text{Lk}_\sigma(\Delta)$ can be obtained from $\text{Lk}_\sigma(\Delta|_{[n] \setminus \{v\}})$ by coning off the subcomplex $\text{Lk}_{\sigma \cup \{v\}}(\Delta)$ – that is, by including $\text{cone}_v(\text{Lk}_{\sigma \cup \{v\}}(\Delta))$. If this subcomplex is itself contractible, then the homotopy type is preserved. \square

Another useful corollary follows from the one above by simply setting $\Delta = \Delta|_{\sigma \cup \tau \cup \{v\}}$ and $[n] = \sigma \cup \tau \cup \{v\}$. We immediately see that if both $\text{Lk}_\sigma(\Delta|_{\sigma \cup \tau \cup \{v\}})$ and $\text{Lk}_{\sigma \cup \{v\}}(\Delta|_{\sigma \cup \tau \cup \{v\}})$ are contractible, then $\text{Lk}_\sigma(\Delta|_{\sigma \cup \tau})$ is contractible. Stated another way:

Corollary 4.3. Assume $v \notin \sigma$ and $\sigma \cap \tau = \emptyset$. If $\text{Lk}_\sigma(\Delta|_{\sigma \cup \tau})$ is *not* contractible, then either (i) $\text{Lk}_\sigma(\Delta|_{\sigma \cup \tau \cup \{v\}})$ is not contractible, and/or (ii) $\sigma \cup \{v\} \in \Delta$ and $\text{Lk}_{\sigma \cup \{v\}}(\Delta|_{\sigma \cup \tau \cup \{v\}})$ is not contractible.

This corollary can be extended to show that for every non-contractible link $\text{Lk}_\sigma(\Delta|_{\sigma \cup \tau})$, there exists a non-contractible “big” link $\text{Lk}_{\sigma'}(\Delta)$ for some $\sigma' \supseteq \sigma$. This is because vertices outside of $\sigma \cup \tau$ can be added one by one to either σ or its complement, preserving the non-contractibility of the new link at each step. (Note that if $\sigma \cup \{v\} \notin \Delta$, we can always add v to the complement. In this case, $\text{Lk}_\sigma(\Delta|_{\sigma \cup \tau}) = \text{Lk}_\sigma(\Delta|_{\sigma \cup \tau \cup \{v\}})$, so we are in case (i) of Corollary 4.3.) In other words, we have the following lemma:

Lemma 4.4. Let $\sigma, \tau \in \Delta$. Suppose $\sigma \cap \tau = \emptyset$, and $\text{Lk}_\sigma(\Delta|_{\sigma \cup \tau})$ is *not* contractible. Then there exists $\sigma' \in \Delta$ such that $\sigma' \supseteq \sigma$, $\sigma' \cap \tau = \emptyset$, and $\text{Lk}_{\sigma'}(\Delta)$ is not contractible.

The next results show that only intersections of facets (maximal faces under inclusion) can possibly yield non-contractible links. For any $\sigma \in \Delta$, we denote by f_σ the intersection of all facets of Δ containing σ . In particular, $\sigma = f_\sigma$ if and only if σ is an intersection of facets of Δ . It is also useful to observe that a simplicial complex is a cone if and only if the common intersection of all its facets is non-empty. (Any element of that intersection can serve as a cone point, and a cone point is necessarily contained in all facets.)

Lemma 4.5. Let $\sigma \in \Delta$. Then $\sigma = f_\sigma \Leftrightarrow \text{Lk}_\sigma(\Delta)$ is *not* a cone.

Proof. Recall that $\text{Lk}_\sigma(\Delta)$ is a cone if and only if all facets of $\text{Lk}_\sigma(\Delta)$ have a non-empty common intersection ν . This can happen if and only if $\sigma \cup \nu \subseteq f_\sigma$. Note that since $\nu \in \text{Lk}_\sigma(\Delta)$, we must have $\nu \cap \sigma = \emptyset$ and hence $\text{Lk}_\sigma(\Delta)$ is a cone if and only if $\sigma \neq f_\sigma$. \square

Furthermore, it is easy to see that every simplicial complex that is not a cone can in fact arise as the link of an intersection of facets. For any Δ that is not a cone, simply consider $\tilde{\Delta} = \text{cone}_v(\Delta)$; v is an intersection of facets of $\tilde{\Delta}$, and $\text{Lk}_v(\tilde{\Delta}) = \Delta$.

The above lemma immediately implies the following corollary:

Corollary 4.6. Let $\sigma \in \Delta$ be nonempty. If $\sigma \neq f_\sigma$, then $\text{Lk}_\sigma(\Delta)$ is a cone and hence contractible. In particular, if $\text{Lk}_\sigma(\Delta)$ is *not* contractible, then σ must be an intersection of facets of Δ (i.e., $\sigma \in \mathcal{F}_\cap(\Delta)$).

Finally, we note that all pairwise intersections of facets that are not also higher-order intersections give rise to non-contractible links.

Lemma 4.7. Let Δ be a simplicial complex. If $\sigma = \tau_1 \cap \tau_2$, where τ_1, τ_2 are distinct facets of Δ , and σ is not contained in any other facet of Δ , then $\text{Lk}_\sigma(\Delta)$ is not contractible.

Proof. Observe that $\text{Lk}_\sigma(\Delta)$ consists of all subsets of $\omega_1 = \tau_1 \setminus \sigma$ and $\omega_2 = \tau_2 \setminus \sigma$, but ω_1 and ω_2 are disjoint because τ_1 and τ_2 do not overlap outside of σ . This means $\text{Lk}_\sigma(\Delta)$ has two connected components, and is thus not contractible. \square

Note that if σ is a pairwise intersection of facets that is also contained in another facet, then $\text{Lk}_\sigma(\Delta)$ could be contractible. For example, the vertex 1 in Figure 3a can be expressed as a pairwise intersection of facets 123 and 145, but is also contained in 134. As shown in Figure 3b, the corresponding link $\text{Lk}_1(\Delta)$ is contractible.

4.2 Proof of Theorem 1.3 and Lemma 1.4

Using the above facts about links, we can now prove Theorem 1.3 and Lemma 1.4. First, we need the following key proposition.

Proposition 4.8. A code \mathcal{C} has no local obstructions if and only if $\sigma \in \mathcal{C}$ for every $\sigma \in \Delta(\mathcal{C})$ such that $\text{Lk}_\sigma(\Delta(\mathcal{C}))$ is non-contractible.

Proof. Let $\Delta = \Delta(\mathcal{C})$, and let $\mathcal{U} = \{U_i\}$ be any collection of open sets such that $\mathcal{C} = \mathcal{C}(\mathcal{U})$. (\Rightarrow) We prove the contrapositive. Suppose there exists $\sigma \in \Delta(\mathcal{C}) \setminus \mathcal{C}$ such that $\text{Lk}_\sigma(\Delta)$ is non-contractible. Then U_σ must be covered by the other sets $\{U_i\}_{i \notin \sigma}$, and since $\text{Lk}_\sigma(\Delta)$ is not contractible, the RF relationship $(\sigma, \bar{\sigma})$ is a local obstruction. (\Leftarrow) We again prove the contrapositive. Suppose \mathcal{C} has a local obstruction (σ, τ) . This means that $\sigma \cap \tau = \emptyset$, $U_\sigma \subseteq \bigcup_{i \in \tau} U_i$, and $\text{Lk}_\sigma(\Delta|_{\sigma \cup \tau})$ is not contractible. By Lemma 4.4, there exists $\sigma' \supseteq \sigma$ such that $\sigma' \cap \tau = \emptyset$ and $\text{Lk}_{\sigma'}(\Delta)$ is not contractible. Moreover, $U_{\sigma'} \subseteq U_\sigma \subseteq \bigcup_{i \in \tau} U_i$ with $\sigma' \cap \tau = \emptyset$, which implies $\sigma' \notin \mathcal{C}$. \square

Theorem 1.3 now follows as a corollary of Proposition 4.8. To see this, let

$$\mathcal{C}_{\min}(\Delta) = \{\sigma \in \Delta \mid \text{Lk}_\sigma(\Delta) \text{ is non-contractible}\} \cup \{\emptyset\},$$

and note that $\mathcal{C}_{\min}(\Delta)$ has simplicial complex Δ . This is because for any facet $\rho \in \Delta$, $\text{Lk}_\rho(\Delta) = \emptyset$, which is non-contractible, and thus $\mathcal{C}_{\min}(\Delta)$ contains all the facets of Δ . By Proposition 4.8, any code \mathcal{C} with simplicial complex Δ has no local obstructions precisely when $\mathcal{C} \supseteq \mathcal{C}_{\min}(\Delta)$. Thus, $\mathcal{C}_{\min}(\Delta)$ is the unique code satisfying the required properties in Theorem 1.3.

Finally, it is easy to see that Lemma 1.4 follows directly from Corollary 4.6.

5 Computing mandatory codewords algebraically

Computing $\mathcal{C}_{\min}(\Delta)$ is certainly simpler than finding *all* local obstructions. However, it is still difficult in general because determining whether or not a simplicial complex is contractible is undecidable [16]. For this reason, we now consider the subset of $\mathcal{C}_{\min}(\Delta)$ corresponding to non-contractible links that can be detected via homology:

$$\mathcal{M}_H(\Delta) \stackrel{\text{def}}{=} \{\sigma \in \Delta \mid \dim \tilde{H}_i(\text{Lk}_\sigma(\Delta), \mathbf{k}) > 0 \text{ for some } i\}, \quad (3)$$

where the $\tilde{H}_i(\cdot)$ are reduced simplicial homology groups, and \mathbf{k} is a field. Homology groups are topological invariants that can be easily computed for any simplicial complex, and *reduced* homology groups simply add a shift in the dimension of $\tilde{H}_0(\cdot)$. This shift is designed so that for any contractible space X , $\dim \tilde{H}_i(X, \mathbf{k}) = 0$ for all integers i . Clearly, $\mathcal{M}_H(\Delta) \subseteq \mathcal{C}_{\min}(\Delta)$, and $\mathcal{M}_H(\Delta)$ is thus a subset of the mandatory codewords that must be included in any convex code \mathcal{C} with $\Delta(\mathcal{C}) = \Delta$.² On the other hand, $\mathcal{M}_H(\Delta) \subseteq \mathcal{C}$ does not guarantee that \mathcal{C} has no local obstructions, as a homologically trivial simplicial complex may be non-contractible.³

It turns out that the entire set $\mathcal{M}_H(\Delta)$ can be computed algebraically, via a minimal free resolution of an ideal built from Δ . Specifically,

$$\mathcal{M}_H(\Delta) = \{\sigma \in \Delta \mid \beta_{i, \bar{\sigma}}(S/I_{\Delta^*}) > 0 \text{ for some } i > 0\}. \quad (4)$$

where $S = \mathbf{k}[x_1, \dots, x_n]$, the ideal I_{Δ^*} is the Stanley-Reisner ideal of the Alexander dual Δ^* , and $\beta_{i, \bar{\sigma}}(S/I_{\Delta^*})$ are the Betti numbers of a minimal free resolution of the ring S/I_{Δ^*} . This is a direct consequence of Hochster's formula:

$$\dim \tilde{H}_i(\text{Lk}_\sigma(\Delta), \mathbf{k}) = \beta_{i+2, \bar{\sigma}}(S/I_{\Delta^*}). \quad (5)$$

A derivation of the above dual version of Hochster's formula, which follows easily from alternative formulations in [9], is given in Supplementary Text S2 (see Lemma 6.5).

Moreover, the subset of mandatory codewords $\mathcal{M}_H(\Delta)$ can be easily computed using existing computational algebra software, such as `Macaulay2` [18]. We now describe this via an explicit example.

²Note that while $\mathcal{M}_H(\Delta)$ depends on the choice of field \mathbf{k} , $\mathcal{M}_H(\Delta) \subseteq \mathcal{C}_{\min}(\Delta)$ for any \mathbf{k} .

³For example, consider a triangulation of the punctured Poincaré homology sphere: this simplicial complex has all-vanishing reduced homology groups, but is non-contractible [17].

Example 5.1. Let Δ be the simplicial complex L25 in Figure 5. The Stanley-Reisner ideal is given by $I_\Delta = \langle x_1x_2x_4, x_2x_3x_4 \rangle$, and its Alexander dual is $I_{\Delta^*} = \langle x_1, x_2, x_4 \rangle \cap \langle x_2, x_3, x_4 \rangle = \langle x_1x_3, x_2, x_4 \rangle$. A minimal free resolution of S/I_{Δ^*} is:

$$0 \leftarrow S/I_{\Delta^*} \xleftarrow{[x_1x_3 \quad x_2 \quad x_4]} S(-2) \oplus S(-1)^2 \xleftarrow{\begin{bmatrix} x_2 & x_4 & 0 \\ -x_1x_3 & 0 & x_4 \\ 0 & -x_1x_3 & -x_2 \end{bmatrix}} S(-3)^2 \oplus S(-2) \xleftarrow{\begin{bmatrix} x_4 \\ -x_2 \\ x_1x_3 \end{bmatrix}} S(-4) \leftarrow 0$$

The Betti number $\beta_{i,\sigma}(S/I_{\Delta^*})$ is the dimension of the module in multidegree σ at step i of the resolution, where S/I_{Δ^*} is step 0 and the steps increase as we move from left to right. At step 0, the total degree is always 0. For the above resolution, the multidegrees at $S(-2) \oplus S(-1)^2$ (step 1) are 1010, 0100, and 0001; at $S(-3)^2 \oplus S(-2)$ (step 2), we have 1110, 1011, and 0101; and at $S(-4)$ (step 4) the multidegree is 1111. This immediately gives us the nonzero Betti numbers:

$$\begin{aligned} \beta_{0,0000}(S/I_{\Delta^*}) &= 1, & \beta_{1,1010}(S/I_{\Delta^*}) &= 1, & \beta_{1,0100}(S/I_{\Delta^*}) &= 1, & \beta_{1,0001}(S/I_{\Delta^*}) &= 1, \\ \beta_{2,1110}(S/I_{\Delta^*}) &= 1, & \beta_{2,1011}(S/I_{\Delta^*}) &= 1, & \beta_{2,0101}(S/I_{\Delta^*}) &= 1, & \beta_{3,1111}(S/I_{\Delta^*}) &= 1. \end{aligned}$$

Recalling from equation (4) that the multidegrees correspond to complements $\bar{\sigma}$ of faces in Δ , we can now immediately read off the elements of $\mathcal{M}_H(\Delta)$ from the above $\beta_{i,\bar{\sigma}}$ for $i > 0$ as:

$$\mathcal{M}_H(\Delta) = \{0101, 1011, 1110, 0001, 0100, 1010, 0000\} = \{24, 134, 123, 4, 2, 13, \emptyset\}.$$

Note that the first three elements of $\mathcal{M}_H(\Delta)$ above, obtained from the Betti numbers $\beta_{1,*}$ in step 1 of the resolution, are precisely the facets of Δ . The next three elements, 0001, 0100, and 1010, are mandatory codewords: they must be included for a code with simplicial complex Δ to be convex. These all correspond to pairwise intersections of facets, and are obtained from the Betti numbers $\beta_{2,*}$ at step 2 of the resolution; this is consistent with the fact that the corresponding links are all disconnected, resulting in non-trivial $\tilde{H}_0(\text{Lk}_\sigma(\Delta), \mathbf{k})$. The last element, 0000, reflects the fact that $\text{Lk}_\emptyset(\Delta) = \Delta$, and $\dim \tilde{H}_1(\Delta, \mathbf{k}) = 1$ for $\Delta = \text{L25}$. By convention, however, we always include the all-zeros codeword in our codes (see Section 1.2).

Using Macaulay2 [18], the Betti numbers for the simplicial complex Δ above can be computed through the following sequence of commands (choosing $\mathbf{k} = \mathbb{Z}_2$, and suppressing outputs except for the Betti tally at the end):

```
i1 : kk = ZZ/2;
i2 : S = kk[x1,x2,x3,x4, Degrees => {{1,0,0,0},{0,1,0,0},{0,0,1,0},{0,0,0,1}}];
i3 : I = monomialIdeal(x1*x2*x4,x2*x3*x4);
i4 : Istar = dual(I);
i5 : M = S^1/Istar;
i6 : Mres = res M;          [comment: this step computes the minimal free resolution]
i7 : peek betti Mres
o7 = BettiTally{(0, {0, 0, 0, 0}, 0) => 1}
      (1, {0, 0, 0, 1}, 1) => 1
      (1, {0, 1, 0, 0}, 1) => 1
      (1, {1, 0, 1, 0}, 2) => 1
      (2, {0, 1, 0, 1}, 2) => 1
      (2, {1, 0, 1, 1}, 3) => 1
      (2, {1, 1, 1, 0}, 3) => 1
      (3, {1, 1, 1, 1}, 4) => 1
```

Each line of the BettiTally displays $(i, \{\sigma\}, |\sigma|) \Rightarrow \beta_{i,\sigma}$. This yields (in order):

$$\beta_{0,0000} = 1, \beta_{1,0001} = 1, \beta_{1,0100} = 1, \beta_{1,1010} = 1, \beta_{2,0101} = 1, \beta_{2,1011} = 1, \beta_{2,1110} = 1, \beta_{3,1111} = 1,$$

which is the same set of nonzero Betti numbers we previously obtained. Recalling again that the multidegrees correspond to complements $\bar{\sigma}$ in (4), and we care only about $i > 0$, this output immediately gives us $\mathcal{M}_H(\Delta)$ – exactly as before.

The above example illustrates how computational algebra can help us to determine whether a code has local obstructions. However, as noted in Section 2, even codes without local obstructions may fail to be convex. Though we have made significant progress via Theorem 1.3, finding a complete combinatorial characterization of convex codes is still an open problem.

Acknowledgments

This work began at a 2014 AMS Mathematics Research Community, “Algebraic and Geometric Methods in Applied Discrete Mathematics,” which was supported by NSF DMS-1321794. CC was supported by NSF DMS-1225666/1537228, NSF DMS-1516881, and an Alfred P. Sloan Research Fellowship; EG was supported by NSF DMS-1304167 and NSF DMS-1620109; JJ was supported by NSF DMS-1606353; and AS was supported by NSF DMS-1004380 and NSF DMS-1312473/1513364. We thank Joshua Cruz, Chad Giusti, Vladimir Itskov, Carly Klivans, William Kronholm, Keivan Monfared, and Yan X Zhang for numerous discussions, and Eva Pastalkova for providing the data used to create Figure 1.

References

- [1] C. Curto, V. Itskov, A. Veliz-Cuba, and N. Youngs. The neural ring: an algebraic tool for analyzing the intrinsic structure of neural codes. *Bull. Math. Biol.*, 75(9):1571–1611, 2013.
- [2] D. H. Hubel and T. N. Wiesel. Receptive fields of single neurons in the cat’s striate cortex. *J. Physiol.*, 148(3):574–591, Oct 1959.
- [3] J. O’Keefe and J. Dostrovsky. The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res.*, 34(1):171–175, 1971.
- [4] Physiology or Medicine 1981-Press Release, 2014. *Nobelprize.org* Nobel Media AB. http://www.nobelprize.org/nobel_prizes/medicine/laureates/1981/press.html.
- [5] N. Burgess. The 2014 Nobel Prize in Physiology or Medicine: A Spatial Model for Cognitive Neuroscience. *Neuron*, 84(6):1120–1125, Dec 2014.
- [6] C. Giusti, E. Pastalkova, C. Curto, and V. Itskov. Clique topology reveals intrinsic geometric structure in neural correlations. *Proc. Natl. Acad. Sci.*, 112(44):13455–13460, 2015.
- [7] C. Curto, V. Itskov, K. Morrison, Z. Roth, and J. L. Walker. Combinatorial neural codes from a mathematical coding theory perspective. *Neural Comput.*, 25(7):1891–1925, 2013.
- [8] C. Giusti and V. Itskov. A no-go theorem for one-layer feedforward networks. *Neural Comput.*, 26(11):2527–2540, 2014.
- [9] E. Miller and B. Sturmfels. *Combinatorial commutative algebra*, volume 227 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 2005.
- [10] C. Lienkaemper, A. Shiu, and Z. Woodstock. Obstructions to convexity in neural codes. Available online at <http://arxiv.org/abs/1509.03328>.

- [11] C. Curto, E. Gross, J. Jeffries, K. Morrison, M. Omar, Z. Rosen, A. Shiu, and N. Youngs. What makes a neural code convex? Available online at <http://arxiv.org/abs/1508.00150>.
- [12] J. Cruz, C. Giusti, V. Itskov, and W. Kronholm. On open and closed convex codes. Available online at <https://arxiv.org/abs/1609.03502>.
- [13] M. 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.
- [14] A. Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge, 2002.
- [15] C. Curto and V. Itskov. Cell groups reveal structure of stimulus space. *PLoS Comput. Biol.*, 4(10):e1000205, 2008.
- [16] M. Tancer. Recognition of collapsible complexes is NP-complete. *Discrete Comput. Geom.*, 55(1):21–38, 2016.
- [17] A. Björner and F.H. Lutz. A 16-vertex triangulation of the Poincaré homology 3-sphere and non-PL spheres with few vertices. *Electronic Geometry Models*, No. 2003.04.001, 2003.
- [18] D. R. Grayson and M. E. Stillman. Macaulay2, a software system for research in algebraic geometry. Available at <http://www.math.uiuc.edu/Macaulay2/>.
- [19] M. Tancer. d -representability of simplicial complexes of fixed dimension. *J. Comput. Geom.*, 2(1):183–188, 2011.
- [20] E. Helly. Über Mengen konvexer Körper mit gemeinschaftlichen Punkten. *Jahresber. Dtsch. Math.-Ver.*, 32:175–176, 1923.
- [21] G. Kalai. Intersection patterns of convex sets. *Israel J. Math.*, 48(2-3):161–174, 1984.
- [22] G. Kalai. Characterization of f -vectors of families of convex sets in \mathbb{R}^d . I. Necessity of Eckhoff’s conditions. *Israel J. Math.*, 48(2-3):175–195, 1984.
- [23] G. Kalai. Characterization of f -vectors of families of convex sets in \mathbb{R}^d . II. Sufficiency of Eckhoff’s conditions. *J Combin. Theory Ser. A*, 41(2):167–188, 1986.

6 Supplementary Text

Contents

S1 Classification of convex codes for $n \leq 4$

S2 Hochster's formula

S3 Dimension bounds

S1 Classification of convex codes for $n \leq 4$

For $n = 1$ or $n = 2$, all codes are convex. The first non-convex codes appear for $n = 3$. Using our convention that all codes include the all-zeros codeword, there are a total of 40 permutation-inequivalent codes on 3 neurons [1]. Of these, only 6 are non-convex (see Table 1 and Lemma 6.1).

label	code \mathcal{C}	$\Delta(\mathcal{C})$
B3	000, 010, 001, 110, 101	L6
B5	000, 010, 110, 101	L6
B6	000, 110, 101	L6
E2	000, 100, 010, 110, 101, 011	L7
E3	000, 100, 110, 101, 011	L7
E4	000, 110, 101, 011	L7

Table 1: All non-convex codes on $n = 3$ neurons, up to permutation equivalence. Code labels are the same as in [1], and simplicial complex labels are as in Figures 5 and 6.

Lemma 6.1. There are 6 non-convex codes on $n \leq 3$ neurons, up to permutation equivalence. They are the codes shown in Table 1.

Proof. First, we show that all six codes in Table 1 are not convex. Let Δ_{L6} and Δ_{L7} be the simplicial complexes labeled L6 and L7 in Figures 5 and 6. It is easy to see that $\{1\} \in \mathcal{C}_{\min}(\Delta_{L6})$ and $\{1\}, \{2\}, \{3\} \in \mathcal{C}_{\min}(\Delta_{L7})$, because the corresponding links are non-contractible. Codes B3, B5, and B6 all have simplicial complex Δ_{L6} , but are missing the codeword 100, corresponding to $\{1\}$. Since $\{1\} \in \mathcal{C}_{\min}(\Delta_{L6})$, by Theorem 1.3 these codes each have a local obstruction and thus cannot be convex. Codes E2, E3, and E4 all have simplicial complex Δ_{L7} , but are missing the codeword 001, corresponding to $\{3\}$. Since $\{3\} \in \mathcal{C}_{\min}(\Delta_{L7})$, these codes cannot be convex. All remaining codes for $n = 3$ neurons were shown to be convex in [1], via explicit convex realizations in two dimensions. \square

We now consider $n \leq 4$. Figure 6 displays all simplicial complexes on $n \leq 4$ vertices, up to permutation equivalence, and highlights all intersections of two or more facets. By inspection, we see that every link corresponding to a non-empty max intersection is not contractible. We thus have the following lemma:

Lemma 6.2. Let \mathcal{C} be a neural code on $n \leq 4$ neurons. Then \mathcal{C} has no local obstructions if and only if \mathcal{C} is max \cap -complete.

Proof. Let $\Delta = \Delta(\mathcal{C})$ be the simplicial complex on $n \leq 4$ vertices corresponding to the code \mathcal{C} . Recall from Theorem 1.3 that $\mathcal{C}_{\min}(\Delta)$ consists of all $\sigma \in \Delta$ such that $\text{Lk}_{\sigma}(\Delta)$ is not contractible. Since, for $n \leq 4$, all non-empty max intersections $\sigma \in \mathcal{F}_{\cap}(\Delta)$ have non-contractible links, it follows that $\mathcal{C}_{\min}(\Delta) = \mathcal{F}_{\cap}(\Delta)$. Therefore, by Theorem 1.3, \mathcal{C} has no local obstructions if and only if $\mathcal{C} \supseteq \mathcal{F}_{\cap}(\Delta)$. \square

Note that the proof of Lemma 6.2 relied on showing that $\mathcal{C}_{\min}(\Delta) = \mathcal{F}_{\cap}(\Delta)$. More generally, for any code \mathcal{C} whose simplicial complex Δ satisfies $\mathcal{C}_{\min}(\Delta) = \mathcal{F}_{\cap}(\Delta)$, we know that \mathcal{C} has no local obstructions if and only if it is max \cap -complete. Combining this with Proposition 1.8, we have the following:

Proposition 6.3. Let Δ be a simplicial complex satisfying $\mathcal{C}_{\min}(\Delta) = \mathcal{F}_{\cap}(\Delta)$. Then for any \mathcal{C} with $\Delta(\mathcal{C}) = \Delta$, \mathcal{C} is convex if and only if it is max \cap -complete.

As a corollary of Lemma 6.2 and Proposition 6.3, we obtain Proposition 1.7.

In fact, for $n \leq 4$ we have generated explicit convex realizations for all max \cap -complete codes. Figure 7 illustrates convex realizations for the max \cap -complete codes corresponding to most of the simplicial complexes L1-L28. The max \cap -complete codes for the omitted complexes L1-L5, L9-L10, and L12 all have obvious convex realizations, while those for L15 and L16 are obvious given the realizations for L7 and L8, respectively.

Table 2 summarizes our classification of $n \leq 4$ codes and provides minimal embedding dimensions $d(\mathcal{C})$ for all the convex codes. Note that $d(\mathcal{C})$ sometimes varies for codes with the same simplicial complex, depending on which optional codewords are included.





Figure 6: Classification of convex codes for $n \leq 4$. L1-L28 are the 28 possible simplicial complexes $\Delta = \Delta(\mathcal{C})$ that can arise, up to permutation equivalence. For each simplicial complex, faces highlighted in color correspond to pairwise (red) and triple (orange) intersections of facets (maximal faces). One can check that for each colored face σ , $\text{Lk}_\sigma(\Delta)$ is not contractible. By Proposition 1.7, any code \mathcal{C} with simplicial complex Δ is convex if and only if \mathcal{C} includes all codewords corresponding to the colored faces. All other non-maximal faces are optional codewords, whose inclusion or exclusion does not affect convexity. Simplicial complexes with no optional codewords are labeled +, while those with no (non-maximal) mandatory codewords are labeled \circ . When both + and \circ labels apply, as in L1, L2, and L4, we simply use \circ . The possible minimal embedding dimensions $d = d(\mathcal{C})$ that can arise for convex codes are shown in the lower right corners.


Δ	# non-convex codes	maximal codewords (facets)	mandatory codewords (non-maximal)	optional codewords in Δ	$d(\mathcal{C})$	picture	notes
L1 $^\circ$	0	1	none	none	1		
L2 $^\circ$	0	10, 01	none	none	1		
L3 $^\circ$	0	11	none	all (2)	1		
L4 $^\circ$	0	100, 010, 001	none	none	1		
L5 $^\circ$	0	110, 001	none	all (2)	1		
L6	3	110, 101	100	010, 001	1	L6-series	
L7 $^+$	3	110, 101, 011	all (3)	none	2	L7-series	convex $\Leftrightarrow \mathcal{C} = \Delta(\mathcal{C})$
L8 $^\circ$	0	111	none	all (6)	1, 2	L8-series	$d = 2 \Leftrightarrow \star$
L9 $^\circ$	0	1000, 0100 0010, 0001	none	none	1		
L10 $^\circ$	0	1100, 0010, 0001	none	all (2)	1		
L11	3	1100, 1010, 0001	1000	0100, 0010	1	L6-series	
L12 $^\circ$	0	1100, 0011	none	all (4)	1		
L13	7	1100, 1010, 0011	1000, 0010	1000, 0001	1	L6-series	
L14	4	1100, 1010, 1001	1000	0100, 0010, 0001	1, 2	L6-series	$d = 2 \Leftrightarrow \star$
L15 $^+$	3	1100, 1010 0110, 0001	all (3)	none	2	like L7	convex $\Leftrightarrow \mathcal{C} = \Delta(\mathcal{C})$
L16 $^\circ$	0	1110, 0001	none	all (6)	1, 2	like L8	$d = 2 \Leftrightarrow \star$
L17	10	1100, 1010 0110, 0101	1000, 0100 0010	0001	2	L7-series	
L18	40	1110, 0101	0100	1000, 0010, 0001 1100, 1010, 0110	1, 2	L8-series	$d = 2 \Leftrightarrow \mathcal{C}$ contains {1000, 0010}, or {1100, 0110, 1010}
L19 $^+$	5	1100, 1010 0101, 0011	all (4)	none	2	L6-series	convex $\Leftrightarrow \mathcal{C} = \Delta(\mathcal{C})$
L20 $^+$	8	1100, 1010, 1001 0110, 0011	all (4)	none	2	L7-series	convex $\Leftrightarrow \mathcal{C} = \Delta(\mathcal{C})$
L21	68	1110, 0101, 0011	0100, 0010 0001	1000, 1100 1010, 0110	2	L8-series	
L22	62	1110, 0111	0110	1000, 0100, 0010 0001, 1100, 1010 0101, 0011	1, 2	L22-series	$d = 2 \Leftrightarrow \star$, or {0101, 0011} $\subset \mathcal{C}$, or {1100, 1010} $\subset \mathcal{C}$
L23 $^+$	4	1100, 1010, 1001 0110, 0101, 0011	all (4)	none	2	L7-series	convex $\Leftrightarrow \mathcal{C} = \Delta(\mathcal{C})$
L24	36	1110, 1001 0101, 0011	1000, 0100 0010, 0001	1100, 1010, 0110	2, 3	L8-series d=3 codes	$d = 3 \Leftrightarrow$ no optional codewords
L25	168	1110, 1011, 0101	0100, 0001 1010	1000, 0010, 1100 1001, 0110, 0011	2	L22-series	
L26	407	1110, 1011, 0111	0010, 1010 0110, 0011	1000, 0100, 0001 1100, 1001, 0101	2	L22-series	
L27 $^+$	85	1110, 1101 1011, 0111	all (10)	none	3	d=3 codes	convex $\Leftrightarrow \mathcal{C} = \Delta(\mathcal{C})$
L28 $^\circ$	0	1111	none	all (14)	1, 2	L22-series	$d = 2 \Leftrightarrow \star$

Table 2: Convexity and dimension for codes on $n \leq 4$ neurons. For each simplicial complex Δ , labeled as in Figure 6, the second column is the number of non-convex codes \mathcal{C} such that $\Delta(\mathcal{C}) = \Delta$, up to permutation equivalence and including the all-zeros codeword, while the sixth column $d(\mathcal{C})$ displays the possible minimal embedding dimensions for convex codes only. The third column lists the codewords corresponding to facets of Δ ; these are automatically included in any code with simplicial complex Δ . The fourth column gives all other non-empty mandatory codewords – that is, elements of $\mathcal{C}_{\min}(\Delta)$ that are not facets of Δ . Optional codewords are elements of Δ whose presence or absence does not affect whether or not a code is convex, though they may alter the minimal embedding dimension $d(\mathcal{C})$. When all non-maximal codewords are mandatory or all are optional, their total number is given in parentheses. The picture column indicates the groupings used for the convex realizations in Figure 7. In the notes column, \star indicates that the set of optional codewords in \mathcal{C} can *not* form a 2-chain. A collection of codewords forms a *chain* if we can completely order the respective sets by containment – so {1111, 1100, 1000} is a chain, but {1110, 1000, 1101} is not. A collection of codewords can form a *2-chain* if it can be partitioned into two sets (possibly empty) which are both chains. + and \circ are the same as in Figure 6.


L6 - series




L7 - series



L8 - series



L22 - series



d=3 codes




Figure 7: Convex realizations for codes on $n \leq 4$ neurons. Each convex set U_i , for neuron i , is the union of all regions corresponding to codewords containing i (see also Figure 4). Note that each picture displays regions corresponding to mandatory codewords in various shades of gray, while optional codewords are in blue. A single picture thus shows convex realizations for all convex codes corresponding to the same simplicial complex, as blue regions can be included or excluded without affecting convexity.

S2 Alexander duality, the Stanley-Reisner ideal, and Hochster's formula

For any simplicial complex Δ on vertex set $[n]$, the *Alexander dual* is the related simplicial complex:

$$\Delta^* \stackrel{\text{def}}{=} \{\bar{\tau} \mid \tau \notin \Delta\},$$

where $\bar{\tau} = [n] \setminus \tau$ denotes the complement of τ in $[n]$. Note that $(\Delta^*)^* = \Delta$. Any Δ also has an associated ideal known as the *Stanley-Reisner ideal*:

$$I_\Delta \stackrel{\text{def}}{=} \left\{ \prod_{i \in \sigma} x_i \mid \sigma \notin \Delta \right\}.$$

Alexander duality relates the reduced homology of a simplicial complex to the cohomology of its Alexander dual,

$$\tilde{H}_i(\Delta^*, \mathbf{k}) \cong \tilde{H}^{n-i-3}(\Delta, \mathbf{k}),$$

where \mathbf{k} is a field. Meanwhile, Hochster's formula relates the nonzero Betti numbers from a minimal free resolution of the Stanley-Reisner ideal to the reduced cohomology of restricted simplicial complexes [9]. Specifically,

$$\beta_{i-1, \sigma}(I_\Delta) = \beta_{i, \sigma}(S/I_\Delta) = \dim \tilde{H}^{|\sigma|-i-1}(\Delta|_\sigma, \mathbf{k}), \quad (6)$$

where $S = \mathbf{k}[x_1, \dots, x_n]$ and the $\beta_{i, \sigma}$ s refer always to Betti numbers for a *minimal* free resolution.⁴

The following link lemma can be used to derive a dual version of Hochster's formula, similar to the version in [9, Corollary 1.40]. The dual formulation is more useful to us, as it allows us to compute the dimensions of all non-trivial homology groups for all links, $\text{Lk}_\sigma(\Delta)$, from a single free resolution.

Lemma 6.4. $\text{Lk}_\sigma(\Delta) = (\Delta^*|_{\bar{\sigma}})^*$.

Proof. First, observe that $\Delta^*|_{\bar{\sigma}} = \{\tau \mid \tau \subset \bar{\sigma} \text{ and } \tau \in \Delta^*\}$. The dual is thus $(\Delta^*|_{\bar{\sigma}})^* = \{\bar{\sigma} \setminus \tau \mid \tau \subset \bar{\sigma} \text{ and } \tau \notin \Delta^*\} = \{\omega \mid \omega \subset \bar{\sigma} \text{ and } \bar{\sigma} \setminus \omega \notin \Delta^*\} = \{\omega \mid \omega \cap \sigma = \emptyset \text{ and } \sigma \cup \omega \in \Delta\} = \text{Lk}_\sigma(\Delta)$. \square

Lemma 6.5 (Hochster's formula, dual version). $\dim \tilde{H}_i(\text{Lk}_\sigma(\Delta), \mathbf{k}) = \beta_{i+2, \bar{\sigma}}(S/I_{\Delta^*})$.

Proof. Using (in order) Lemma 6.4, Alexander duality, and the original version of Hochster's formula (6), we obtain:

$$\dim \tilde{H}_i(\text{Lk}_\sigma(\Delta), \mathbf{k}) = \dim \tilde{H}_i((\Delta^*|_{\bar{\sigma}})^*, \mathbf{k}) = \dim \tilde{H}^{|\bar{\sigma}|-i-3}(\Delta^*|_{\bar{\sigma}}, \mathbf{k}) = \beta_{i+2, \bar{\sigma}}(S/I_{\Delta^*}).$$

\square

It is important to note that if σ is a facet of Δ , then $\text{Lk}_\sigma(\Delta) = \emptyset$, which is non-contractible due to nontrivial homology in degree -1 . Hochster's formula thus detects facets of Δ via the nonzero Betti numbers $\beta_{1, \bar{\sigma}}$, as these correspond to σ such that $\dim \tilde{H}_{-1}(\text{Lk}_\sigma(\Delta), \mathbf{k}) > 0$. Note also that $\text{Lk}_\emptyset(\Delta) = \Delta$, so if Δ itself has nontrivial reduced homology in degree i , this will be detected as a nonzero Betti number $\beta_{i+2, [n]}$, where $\bar{\sigma} = [n]$ is the complement of $\sigma = \emptyset$.

S3 Bounds on the minimal embedding dimension of convex codes

We now turn to the problem of determining the *minimal embedding dimension* $d(\mathcal{C})$ of a convex code \mathcal{C} . There is no general method for computing $d(\mathcal{C})$, though bounds can be obtained from the information present in the simplicial complex $\Delta(\mathcal{C})$. In this section, we review known results on d -representability and Helly's theorem, and apply them to obtain lower bounds on $d(\mathcal{C})$. We also obtain an additional bound on $d(\mathcal{C})$ from the Fractional Helly theorem, and examine how it compares to the Helly's theorem bound.

⁴See [9, Chapter 1] for more details about free resolutions of monomial ideals.

Our dimension bounds all rely solely on features of the code captured by $\Delta(\mathcal{C})$, and do not take into account the finer structure of the code. Nevertheless, the presence or absence of a single codeword can have a significant effect on $d(\mathcal{C})$, even if the simplicial complex $\Delta = \Delta(\mathcal{C})$ is fixed (e.g., see Table 2 for L8, L14, L16, etc.). It remains an open question how to use this additional information in order to improve the bounds on $d(\mathcal{C})$.

Embedding dimension and d -representability

The problem of determining $d(\mathcal{C})$ for a convex code \mathcal{C} has not been directly addressed in the literature. However, the related problem of determining when a simplicial complex Δ can be realized as the nerve $\mathcal{N}(\mathcal{U})$ of a cover \mathcal{U} has received considerable attention (see [13, 19] and references therein). A simplicial complex Δ is said to be d -representable if there exists a collection of convex (not necessarily open) sets $\mathcal{U} = \{U_1, \dots, U_n\}$, with $U_i \subset \mathbb{R}^d$, such that $\Delta = \mathcal{N}(\mathcal{U})$. Note that for such a \mathcal{U} , the corresponding code $\mathcal{C}(\mathcal{U})$ need not be equal to Δ , though it is always true that $\Delta(\mathcal{C}(\mathcal{U})) = \Delta$.

While d -representability of $\Delta(\mathcal{C})$ does not tell us the value of $d(\mathcal{C})$, it does provide a lower bound. This motivates us to define the *nerve dimension* $d_{\mathcal{N}}(\mathcal{C})$ of a code \mathcal{C} to be the minimal d such that $\Delta(\mathcal{C})$ is d -representable. It immediately follows that

$$d(\mathcal{C}) \geq d_{\mathcal{N}}(\mathcal{C}),$$

because any embedding of \mathcal{C} in $\mathbb{R}^{d(\mathcal{C})}$ via a collection of convex open sets \mathcal{U} is also a realization of $\Delta(\mathcal{C})$ as the nerve $\mathcal{N}(\mathcal{U})$. Unfortunately, $d_{\mathcal{N}}(\mathcal{C})$ may be difficult to compute in general. In contrast, we can obtain a lower bound from Helly's theorem that is simple to read off from the structure of $\Delta(\mathcal{C})$.

Bounds from Helly's theorem

One common tool used to address d -representability of simplicial complexes is Helly's theorem.

Lemma 6.6 (Helly's theorem [20]). Let $\mathcal{U} = \{U_1, \dots, U_n\}$ be a collection of convex open sets in \mathbb{R}^d . If for every $d + 1$ sets in \mathcal{U} , the intersection is non-empty, then the full intersection $\bigcap_{i=1}^n U_i$ is also non-empty.

Helly's theorem implies that if Δ is d -representable and Δ contains all possible d -dimensional faces, then Δ must be the full simplex. On the other hand, if Δ contains all possible d -dimensional faces but is not the full simplex, then it is *not* d -representable. This immediately yields examples where the presence or absence of a single codeword can have a large effect on $d(\mathcal{C})$.

Proposition 6.7. Let \mathcal{C} be a code on n neurons, and suppose that for some k with $1 \leq k < n$, $\Delta(\mathcal{C})$ contains all k -dimensional faces. If $11 \cdots 1 \in \mathcal{C}$, then $d(\mathcal{C}) \leq 2$; otherwise, $d(\mathcal{C}) > k$.

Proof. In the first case, where $11 \cdots 1 \in \mathcal{C}$, the fact that \mathcal{C} is convex and $d(\mathcal{C}) \leq 2$ follows from Lemma 2.5. For the second case, where $11 \cdots 1 \notin \mathcal{C}$, suppose \mathcal{C} is realizable as a convex code in \mathbb{R}^d for some $d \leq k$, so that $\mathcal{C} = \mathcal{C}(\mathcal{U})$ for some collection of convex open sets $\mathcal{U} = \{U_1, U_2, \dots, U_n\}$, with each $U_i \subset \mathbb{R}^d$. Since, by hypothesis, $\Delta(\mathcal{C})$ contains all k -dimensional faces, it also contains all d -dimensional faces, and so the intersection of every collection of $d + 1$ subsets in \mathcal{U} is non-empty. Thus, by Helly's Theorem, the full intersection of all sets in \mathcal{U} is non-empty, and so $11 \cdots 1 \in \mathcal{C}$. This contradicts the fact that $11 \cdots 1 \notin \mathcal{C}$; hence, we must have $d(\mathcal{C}) > k$. \square

We can also apply Helly's theorem to every subcollection $\{U_{i_1}, \dots, U_{i_m}\} \subset \mathcal{U}$, or equivalently to the induced subcomplex on elements i_1, \dots, i_m , to see that if all the d -dimensional faces of this subcomplex are present, then the top-dimensional face must also be present in order for Δ to be d -representable. This leads us to the following definitions. A simplicial complex is said to contain an *induced k -dimensional simplicial hole* if it contains $k + 1$ vertices such that the induced subcomplex

on those vertices is isomorphic to a hollow simplex (see Section 1.2, and [13]). We define the *Helly dimension*⁵ of \mathcal{C} , denoted $d_H(\mathcal{C})$, to be the dimension of the largest induced simplicial hole of $\Delta(\mathcal{C})$:

$$d_H(\mathcal{C}) \stackrel{\text{def}}{=} \max\{k \mid \Delta(\mathcal{C}) \text{ has a } k\text{-dimensional induced simplicial hole}\}. \quad (7)$$

Clearly, $d_H(\mathcal{C}) \leq d(\mathcal{C})$.

Bounds from the Fractional Helly theorem

The Fractional Helly theorem is a well-known extension of Helly's theorem that provides new bounds on $d(\mathcal{C})$, though they are not always better.

Lemma 6.8 (Fractional Helly theorem⁶, Theorem 6.7 of [13]). Let $\alpha > 0$, and $\mathcal{U} = \{U_1, U_2, \dots, U_n\}$ be a collection of convex open sets in \mathbb{R}^d such that at least $\alpha \binom{n}{d+1}$ of the $(d+1)$ -tuples of sets in \mathcal{U} have non-empty intersections. Then there exists $\sigma \subseteq [n]$ such that $|\sigma| > \frac{\alpha}{d+1}n$, and $\bigcap_{i \in \sigma} U_i$ is non-empty.

The Fractional Helly theorem indicates that if a code \mathcal{C} can be embedded in \mathbb{R}^d , and the simplicial complex $\Delta(\mathcal{C})$ has many d -dimensional faces, then $\Delta(\mathcal{C})$ must have some sufficiently high-dimensional face. The following lemma quantifies these observations in our context.

Lemma 6.9. Let Δ be a k -dimensional simplicial complex on n elements, and let $f_d(\Delta)$ be the number of d -dimensional faces in Δ for $1 \leq d < n$. If Δ is d -representable, then $k + 1 > f_d(\Delta) / \binom{n-1}{d}$.

Proof. By definition of d -representable, we have $\Delta = \mathcal{N}(\mathcal{U})$ for some $\mathcal{U} = \{U_1, U_2, \dots, U_n\}$, where each $U_i \subset \mathbb{R}^d$. Since each d -dimensional face of Δ corresponds to an intersection of $(d+1)$ of the U_i s, we have that $f_d(\Delta(\mathcal{C}))$ of the $(d+1)$ -tuples have non-empty intersections. By the Fractional Helly theorem, there is some $\sigma \subseteq [n]$ with $|\sigma| > \frac{\alpha}{d+1}n$ such that $\bigcap_{i \in \sigma} U_i \neq \emptyset$, where $\alpha = \frac{f_d(\Delta(\mathcal{C}))}{\binom{n}{d+1}}$. Since $\Delta(\mathcal{C})$ is k -dimensional, it follows that $|\sigma| \leq k + 1$, and so $k + 1 > \frac{\alpha}{d+1}n = f_d(\Delta(\mathcal{C})) / \binom{n-1}{d}$. \square

This leads us to the following definition. Let \mathcal{C} be a code on n neurons with a k -dimensional simplicial complex $\Delta(\mathcal{C})$, and let $f_d(\Delta(\mathcal{C}))$ be the number of d -dimensional faces in $\Delta(\mathcal{C})$ for $1 \leq d < n$. The *Fractional Helly dimension* $d_{FH}(\mathcal{C})$ of \mathcal{C} is given by:

$$d_{FH}(\mathcal{C}) \stackrel{\text{def}}{=} 1 + \max \left\{ d \mid f_d(\Delta(\mathcal{C})) \geq (k+1) \cdot \binom{n-1}{d}, 1 \leq d < n \right\}. \quad (8)$$

Further dimension bounds based on the f -vector $\{f_i(\Delta)\}$ can be obtained from the results in [22, 23].

Comparison of dimension bounds

How do the Helly and Fractional Helly dimensions $d_H(\mathcal{C})$ and $d_{FH}(\mathcal{C})$ compare to each other and to the nerve dimension $d_{\mathcal{N}}(\mathcal{C})$? First, note that although a simplicial complex Δ cannot be represented in any dimension less than its Helly dimension $d_H(\Delta)$, Helly's theorem does not guarantee that Δ is d_H -representable. Thus we have $d_H(\mathcal{C}) \leq d_{\mathcal{N}}(\mathcal{C}) \leq d(\mathcal{C})$. By the same reasoning, $d_{FH}(\mathcal{C}) \leq d_{\mathcal{N}}(\mathcal{C}) \leq d(\mathcal{C})$. The next examples show that we can have $d_H(\mathcal{C}) < d_{\mathcal{N}}(\mathcal{C})$ and $d_{FH}(\mathcal{C}) < d_{\mathcal{N}}(\mathcal{C})$; i.e., the nerve dimension can provide a strictly stronger lower bound.

⁵The closely-related notion of *Helly number* for a simplicial complex was previously introduced in the literature. Specifically, the Helly number of $\Delta(\mathcal{C})$ is $d_H(\mathcal{C}) + 1$.

⁶In [21], Kalai gave a sharp version of Theorem 6.8: If $\alpha \binom{n}{d+1}$ of the $(d+1)$ -tuples have non-empty intersections, then there exists σ with $|\sigma| > (1 - (1 - \alpha)^{1/(d+1)})n$ such that the corresponding intersection is non-empty. The precise value of β for which $|\sigma| > \beta n$ is not essential to the results in the remainder of our paper. Therefore, for ease of notation and computation, we prefer $\beta = \alpha/(d+1)$ as in Theorem 6.8, while keeping in mind that a tighter bound exists.




Figure 8: **(a)** A simplicial complex Δ , with facets 125, 235, 345, 145. **(b)** A convex realization of Δ in \mathbb{R}^2 .

Example 6.10. Consider any code \mathcal{C} such that $\Delta(\mathcal{C})$ is the simplicial complex in Figure 8a. We obtain $d_H(\mathcal{C}) = 1$, because the two maximal induced simplicial holes of $\Delta(\mathcal{C})$ arise from the subsets $\{1, 3\}$ and $\{2, 4\}$, which both have dimension 1. Although $\Delta(\mathcal{C})$ is contractible, the induced subcomplex on $\{1, 2, 3, 4\}$ is a 1-cycle, so $\Delta(\mathcal{C})$ is at best 2-representable and $d_{\mathcal{N}}(\mathcal{C}) \geq 2$. Thus, $d_{\mathcal{N}}(\mathcal{C}) > d_H(\mathcal{C})$. Figure 8b shows that the minimal embedding dimension is, in fact, $d(\mathcal{C}) = 2$.

Our final example shows not only that we can have $d_{FH}(\mathcal{C}) < d_{\mathcal{N}}(\mathcal{C})$, but also that it is possible for $d_{FH}(\mathcal{C}) < d_H(\mathcal{C})$ or for $d_H(\mathcal{C}) < d_{FH}(\mathcal{C})$, depending on the code. So although $d_{\mathcal{N}}(\mathcal{C})$ is always the strongest of the three bounds, neither of the easier-to-compute $d_H(\mathcal{C})$ and $d_{FH}(\mathcal{C})$ bounds is universally stronger than the other. In other words, all we can say in general is that the minimal embedding dimension satisfies:

$$d(\mathcal{C}) \geq d_{\mathcal{N}}(\mathcal{C}) \geq \max\{d_H(\mathcal{C}), d_{FH}(\mathcal{C})\}. \quad (9)$$




Figure 9: **(a)** The simplicial complex of code \mathcal{C}_1 from Example 6.11. **(b)** The simplicial complex of code \mathcal{C}_2 from Example 6.11, for $r = 4$.

Example 6.11. (a) Let $n = 3$, and consider the code $\mathcal{C}_1 = \{000, 110, 101, 011\}$, whose simplicial complex $\Delta(\mathcal{C}_1)$ is the empty triangle, as shown in Figure 9a. This is a hollow 2-simplex, so we have $d_H(\mathcal{C}_1) = 2$ and $d_{\mathcal{N}}(\mathcal{C}_1) = 2$. In contrast, a quick computation yields $d_{FH}(\mathcal{C}_1) = 1$, and so $d_H(\mathcal{C}_1) = d_{\mathcal{N}}(\mathcal{C}_1) > d_{FH}(\mathcal{C}_1)$.

(b) Let $n = 2r$, where $r \geq 4$, and suppose $\Delta(\mathcal{C}_2) = K_{r,r}$, the complete bipartite graph on $2r$ vertices, as shown in Figure 9b. This graph contains no triangles, so the largest induced simplicial holes result from missing edges in $\Delta(\mathcal{C}_2)$, which have dimension 1. Thus, $d_H(\mathcal{C}_2) = 1$. To compute $d_{FH}(\mathcal{C}_2)$, we first find the f -vector of $\Delta(\mathcal{C}_2)$. Observe that $f_0(\Delta(\mathcal{C}_2)) = n$, $f_1(\Delta(\mathcal{C}_2)) = r^2$, and $f_i(\Delta(\mathcal{C}_2)) = 0$ for $2 \leq i < n$. Note also that $k = 1$, since $\Delta(\mathcal{C}_2)$ is 1-dimensional. Since $f_d(\Delta(\mathcal{C}_2)) = 0$ for $d \geq 2$, the inequality $f_d(\Delta(\mathcal{C}_2)) \geq (k+1) \cdot \binom{n-1}{d}$ is *not* satisfied for these values of d . However, we have

$$f_1(\Delta(\mathcal{C}_2)) = r^2 \geq 2(2r-1) = 2(n-1) = (k+1) \cdot \binom{n-1}{1},$$

with the inequality being valid for all $r \geq 4$. Thus, directly from the definition, we find $d_{FH}(\mathcal{C}_2) = 2$. Hence, $d_{FH}(\mathcal{C}_2) > d_H(\mathcal{C}_2)$.