

What Makes a Neural Code Convex?*

Carina Curto[†], Elizabeth Gross[‡], Jack Jeffries[§], Katherine Morrison[¶], Mohamed Omar^{||},
Zvi Rosen[#], Anne Shiu^{††}, and Nora Youngs^{‡‡}

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 our understanding of the intrinsic combinatorial properties of convex codes.

Key words. neural coding, convex codes, simplicial complex, link, Nerve lemma, Hochster’s formula

AMS subject classifications. 92, 52, 05, 13

DOI. 10.1137/16M1073170

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

*Received by the editors May 3, 2016; accepted for publication (in revised form) December 21, 2016; published electronically March 28, 2017.

<http://www.siam.org/journals/siaga/1/M107317.html>

Funding: 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.

[†]Department of Mathematics, The Pennsylvania State University, University Park, PA 16802 (ccurto@psu.edu).

[‡]Department of Mathematics, San José State University, San José, CA 95192 (elizabeth.gross@sjsu.edu).

[§]Department of Mathematics, University of Utah, Salt Lake City, UT 84112. Current address: Department of Mathematics, University of Michigan, Ann Arbor, MI 48109 (jeffries45@gmail.com).

[¶]Department of Mathematics, The Pennsylvania State University, University Park, PA 16802. Current address: School of Mathematical Sciences, University of Northern Colorado, Greeley, CO 80639 (kmorris2@gmail.com).

^{||}Department of Mathematics, Harvey Mudd College, Claremont, CA 91711 (omar@g.hmc.edu).

[#]Department of Mathematics, The Pennsylvania State University, University Park, PA 16802, and Department of Mathematics, University of California, Berkeley, Berkeley, CA 94720. Current address: Department of Mathematics, University of Pennsylvania, Philadelphia, PA 19104 (zvhr@sas.upenn.edu).

^{††}Department of Mathematics, Texas A&M University, College Station, TX 77843 (annejls@math.tamu.edu).

^{‡‡}Department of Mathematics, Harvey Mudd College, Claremont, CA 91711. Current address: Department of Mathematics, Colby College, Waterville, ME 04901 (nyoungs@g.hmc.edu).

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 comprises a convex code for the animal’s position in space. Note that in the neuroscience literature, convex receptive fields are typically referred to as *unimodal*, emphasizing the presence of just one “hot spot” in a stimulus-response heat map such as those depicted in Figure 1.




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, the intrinsic combinatorial signatures of convex and nonconvex codes are not clear. 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. This may be particularly important for studying systems such as olfaction, where the underlying “olfactory space” is potentially high-dimensional and poorly understood. Having intrinsic signatures of convex codes is a critical step toward understanding whether something like convex coding may be going on in such areas. 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,

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

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 nonconvex 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 :

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

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

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

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 (M107317_01.pdf [local/web 669KB]), 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 following 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 (see [12, Theorem 4.4]). *If \mathcal{C} is $\max \cap$ -complete, then \mathcal{C} is convex.*

Finally, in Supplementary Text S3 (M107317_01.pdf [local/web 669KB]) 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 version of Helly’s 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 Supplementary Text S1, M107317_01.pdf [local/web 669KB]). 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 $d(\mathcal{C}) = 2$.

Example 2.1. Consider the open cover \mathcal{U} illustrated in Figure 2(a). 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 2(b)) 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.”




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).

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 3(a) 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 3(a)). 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 3(c).

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 (see [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, M107317.01.pdf [local/web 669KB]) and inspired aspects of the proof of Proposition 1.8, given in [12].

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 $d(\mathcal{C}) \leq 2$.*




Figure 4. A convex realization in \mathbb{R}^2 of the code in Example 2.4. (Left) Each nonmaximal 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.

Proof. Let $\rho \in \mathcal{C}$ be the unique maximal codeword, and let $m = |\mathcal{C}| - 2$ be the number of nonmaximal 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 nonmaximal codeword (excluding $00 \cdots 0$) 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 $\mathcal{C} = \mathcal{C}(\{U_i\})$. ■

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 nonoverlapping 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 the following 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. ■

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 nonempty intersection, $U_\sigma = \bigcap_{i \in \sigma} U_i$, is contractible.¹

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].

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

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 an 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 cannot 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 noncontractible 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 $|\tau| \leq 4$. Noncontractible 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 an 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. ■

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 noncontractible 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 subcomplex, 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. ■

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. ■

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.

Corollary 4.3. *Assume $v \notin \sigma$ and $\sigma \cap \tau = \emptyset$. If $\text{Lk}_{\sigma}(\Delta|_{\sigma \cup \tau})$ is not contractible, then (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 noncontractible link $\text{Lk}_{\sigma}(\Delta|_{\sigma \cup \tau})$, there exists a noncontractible “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 noncontractibility 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 noncontractible 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 nonempty. (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 nonempty 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 $\sigma \neq f_{\sigma}$. ■

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 noncontractible 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. ■

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 3(a) can be expressed as a pairwise intersection of facets 123 and 145 but is also contained in 134. As shown in Figure 3(b), 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 noncontractible.*

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 noncontractible. 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 $\sigma' \notin \mathcal{C}$. ■

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

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

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 noncontractible, 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 noncontractible links that can be detected via homology:

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

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 noncontractible.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,

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

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 [9]:

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

See Supplementary Text S2 (M107317_01.pdf [local/web 669KB]) for more details on Alexander duality, Hochster’s formula, and the Stanley–Reisner ideal.

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.

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

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

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:

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

Recalling from (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

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

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 noncontractible [17].

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 nontrivial $\tilde{H}_0(\text{Lk}_\sigma(\Delta), \mathbf{k})$. 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. 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 (2013), pp. 1571–1611.
- [2] D. H. HUBEL AND T. N. WIESEL, *Receptive fields of single neurons in the cat's striate cortex*, J. Physiol., 148 (1959), pp. 574–591.

- [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 (1971), pp. 171–175.
- [4] NOBEL MEDIA, *Physiology or Medicine 1981-Press Release*, Nobelprize.org, Nobel Media AB 2014, http://www.nobelprize.org/nobel_prizes/medicine/laureates/1981/press.html, 1981.
- [5] N. BURGESS, *The 2014 Nobel Prize in Physiology or Medicine: A Spatial Model for Cognitive Neuroscience*, Neuron, 84 (2014), pp. 1120–1125.
- [6] C. GIUSTI, E. PASTALKOVA, C. CURTO, AND V. ITS KOV, *Clique topology reveals intrinsic geometric structure in neural correlations*, Proc. Natl. Acad. Sci. USA, 112 (2015), pp. 13455–13460.
- [7] C. CURTO, V. ITS KOV, K. MORRISON, Z. ROTH, AND J. L. WALKER, *Combinatorial neural codes from a mathematical coding theory perspective*, Neural Comput., 25 (2013), pp. 1891–1925.
- [8] C. GIUSTI AND V. ITS KOV, *A no-go theorem for one-layer feedforward networks*, Neural Comput., 26 (2014), pp. 2527–2540.
- [9] E. MILLER AND B. STURMFELS, *Combinatorial Commutative Algebra*, Grad. Texts in Math. 227, Springer-Verlag, New York, 2005.
- [10] C. LIENKAEMPER, A. SHIU, AND Z. WOODSTOCK, *Obstructions to convexity in neural codes*, Adv. Appl. Math., 85 (2017), pp. 31–59.
- [11] C. CURTO, E. GROSS, J. JEFFRIES, K. MORRISON, M. OMAR, Z. ROSEN, A. SHIU, AND N. YOUNGS, *What Makes a Neural Code Convex?*, preprint, <https://arxiv.org/abs/1508.00150>, 2016.
- [12] J. CRUZ, C. GIUSTI, V. ITS KOV, AND W. KRONHOLM, *On Open and Closed Convex Codes*, preprint, <https://arxiv.org/abs/1609.03502>, 2016.
- [13] M. TANCER, *Intersection patterns of convex sets via simplicial complexes: A survey*, in Thirty Essays on Geometric Graph Theory, Springer, New York, 2013, pp. 521–540.
- [14] A. HATCHER, *Algebraic Topology*, Cambridge University Press, Cambridge, UK, 2002.
- [15] C. CURTO AND V. ITS KOV, *Cell groups reveal structure of stimulus space*, PLoS Comput. Biol., 4 (2008), e1000205.
- [16] M. TANCER, *Recognition of collapsible complexes is NP-complete*, Discrete Comput. Geom., 55 (2016), pp. 21–38.
- [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, 2003.04.001, <http://www.eg-models.de/>, 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 (2011), pp. 183–188.
- [20] E. HELLY, *Über Mengen konvexer Körper mit gemeinschaftlichen Punkten*, Jahresber. Deutsch. Math.-Verein., 32 (1923), pp. 175–176.
- [21] G. KALAI, *Intersection patterns of convex sets*, Israel J. Math., 48 (1984), pp. 161–174.
- [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 (1984), pp. 175–195.
- [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 (1986), pp. 167–188.