NEURAL CODES, UNDECIDABILITY, AND A NEW CLASS OF LOCAL OBSTRUCTIONS

AARON CHEN

ABSTRACT. Given an intersection pattern of open sets in Euclidean space, is it possible to tell if there is an arrangement so that the open sets are convex? This problem that can be classified as combinatorial/topological in nature, but surprisingly appears as a mathematical model for spatial cognition motivated by research in neuroscience on place cells in the brain (where the intersection patterns are denoted as "neural codes"). We prove that the notions of a neural code being locally good and a good cover code are in fact equivalent, and that the corresponding decision problem is undecidable. We also present a new type of local obstruction to convexity by considering collapsibility of links of missing codewords.

1. INTRODUCTION

The discovery of place cells by O'Keefe et al in 1971 was a major breakthrough in the field of neuroscience that led to a Nobel Prize in Medicine or Physiology in 2014 [6]. A place cell is a neuron that encodes the spatial information of an organism's surroundings by firing only when the organism is in a place field, which can be modeled by a convex open set. These place fields thus form an open cover of the organism's surroundings. One can describe the intersection patterns of n sets by a subset of $\{0,1\}^n$, where the ith coordinate represents the binary state of neuron i, which fires if and only if the organism is in the corresponding place field. Such intersection patterns have been coined *combinatorial neural codes*, which we will abbreviate as "neural codes," or when it is clear, just "codes." Each binary vector in such a neural code is called a codeword. Because codewords are binary vectors, they can be uniquely represented by their support as a subset of $[n] = \{1, ..., n\}$. Now given a (finite) collection of open sets in Euclidean space, it is straightforward to find the corresponding neural code that describes the intersection patterns of these sets. If we were to restrict the properties of the open sets (say, to require all sets to be convex), then the inverse problem of determining whether such an arrangement of open sets even exists is of particular interest. Indeed, intersection patterns of convex sets have been discussed in detail when the intersection patterns are precisely a simplicial complex (see [7] for an overview), but this more specific case has only caught attention recently.

2. BACKGROUND

Here we introduce some notation as well as basic definitions associated to the theory of neural codes. We will denote $[n] = \{1, ..., n\}$. We will reserve lowercase Greek letters to always be a subset of $[n]$ for some n, which usually refers to either a codeword in a neural code or a face in a simplicial complex. For shorthand, however, we will omit the braces and commas (e.g. if $\tau = \{1, 2, 3\}$ and $\sigma = \{2, 3, 4\}$ we write $\tau = 123, \sigma = 234, \tau \cap \sigma = 23$). Additionally, given a collection of sets $U_1, ..., U_n$, we define $U_\tau = \bigcap_{i \in \tau} U_i$.

2.1. Basic Definitions.

Consider a collection $\mathcal{U} = \{U_1, ..., U_n\}$ of open sets in $\mathbb{R}^d$, corresponding to locations where a neuron will fire.

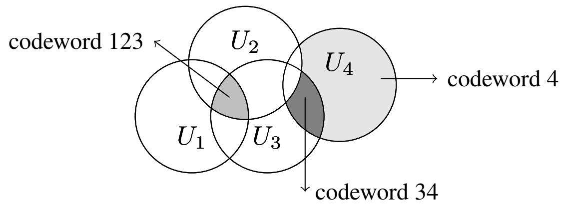

FIGURE 1. A realization of the code $\mathcal{C} = \{123, 234, 12, 23, 13, 24, 34, 1, 2, 3, 4, \emptyset\}$.

Definition 2.1. *A* ***neural code*** *$\mathcal{C}$ on n neurons is a subset of $2^{[n]}$* [1]*, and each $\sigma \in \mathcal{C}$ is called a* ***codeword****. Any codeword that is maximal in $\mathcal{C}$ with respect to set inclusion is a* ***maximal codeword****.*

Definition 2.2. *A code $\mathcal{C}$ is* ***realized*** *by open sets $\mathcal{U} = \{U_1, ..., U_n\}$ if*

$$\sigma \in \mathcal{C} \iff U_\sigma \setminus \bigcup_{j \notin \sigma} U_j \neq \emptyset. \tag{2.1}$$

Conversely, given a collection of open sets $\mathcal{U}$, the corresponding unique code that is realized by $\mathcal{U}$ will be denoted $\mathcal{C}(\mathcal{U})$.

We note that in general any neural code can be realized by open sets. However, if we put more constraints on the open sets $\mathcal{U}$, then we can ask more interesting questions. The following are several properties of neural code that are of particular interest:

Definition 2.3. *A neural code $\mathcal{C}$ is said to be*

- ***convex*** *if $\mathcal{C}$ can be realized by convex open sets $\mathcal{U}$.*
- *a* ***good cover code*** *if $\mathcal{C}$ can be realized by contractible open sets $\mathcal{U}$ such that any intersection of sets in $\mathcal{U}$ is also contractible. Such a $\mathcal{U}$ is called a* ***good cover****.*
- ***max-intersection complete*** *if $\mathcal{C}$ is closed under taking intersections of maximal codewords.*
- ***k-sparse*** *if for all $\sigma \in \mathcal{C}$, $|\sigma| \leq k$.*

Example 2.4. *The codeword $\mathcal{C} = \{123, 234, 12, 23, 13, 24, 34, 1, 2, 3, 4, \emptyset\}$ is 3-sparse, has maximal codewords 123 and 234, and is realized in $\mathbb{R}^2$ above in Figure 1. $\mathcal{C}$ is convex (and therefore a good cover code) and max-intersection complete since $123 \cap 234 = 23$ is a codeword in $\mathcal{C}$.*

Definition 2.5. *An* ***abstract simplicial complex*** *Δ on n vertices is a subset of $2^{[n]}$ that is closed under taking subsets. Each $\sigma \in \Delta$ is called a* ***face*** *of Δ. We can topologically realize any simplicial complex in Euclidean space of sufficiently high dimension.*

Given a simplicial complex Δ, we will always use $|\Delta|$ to refer to the topological realization (which is unique up to homeomorphism). For σ a face in Δ, we use $[\sigma]$ to refer to the topological realization of σ as a simplex contained in $|\Delta|$. We define $\Delta(\mathcal{C})$ of a code $\mathcal{C}$ to be the smallest simplicial complex containing $\mathcal{C}$, which is uniquely defined by the maximal codewords of $\mathcal{C}$ (which are the maximal faces, or **facets**, of $\Delta(\mathcal{C})$).

Definition 2.6. *The* ***link*** *of a face σ in a simplicial complex Δ is defined as*

$$\mathrm{Lk}_\sigma(\Delta) = \{\tau \in \Delta \mid \sigma \cap \tau = \emptyset \text{ and } \sigma \cup \tau \in \Delta\}.$$

Note that $\sigma \cup \tau \in \Delta$ is the same as saying that $\sigma * \tau \subseteq |\Delta|$, where $*$ is the the topological join. We will only refer to the link in a topological sense, so $Lk_\sigma(\Delta)$ is a subset of Δ.

A closely related notion to $\Delta(\mathcal{C})$ is the nerve of a covering.

[1]We will follow the simplifying convention that the empty codeword is always in a neural code.

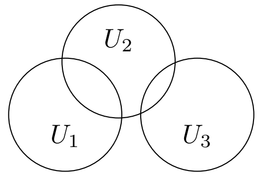

FIGURE 2. The code $\mathcal{C} = \{12, 23, 1, 2, 3, \emptyset\}$ is convex and a good-cover code. Its topological simplicial complex, $|\Delta(\mathcal{C})|$, is realized below.

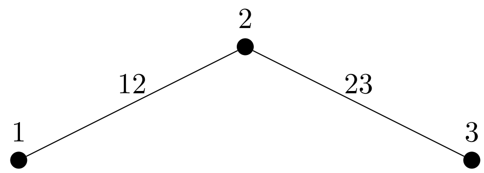

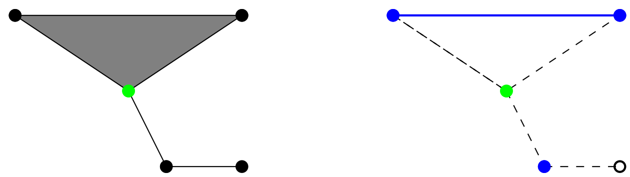

FIGURE 3. The link of the green vertex in the simplicial complex on the left is shown on the right in blue.

Definition 2.7. *Given a collection of sets $\mathcal{U} = \{U_1, ..., U_n\}$, the **nerve** of $\mathcal{U}$, denoted as $\mathcal{N}(\mathcal{U})$, is the simplicial complex Δ where*

$$\sigma \in \Delta \Longleftrightarrow U_\sigma \neq \emptyset$$

with the exception of $\emptyset$, which by definition is always in Δ.

Remark 2.8. $\mathcal{N}(\mathcal{U}) = \Delta(\mathcal{C}(\mathcal{U}))$.

The classical result referred to as the nerve theorem or nerve lemma, is a nice property of the nerve of an open covering that we will state in a more specific setting.

Theorem 2.9 (Nerve Theorem, [10]). *Given a finite collection of open sets $\mathcal{U}$ in Euclidean space where each set and intersection of sets is either empty or contractible, the union of all sets in $\mathcal{U}$ is homotopy equivalent to $\mathcal{N}(\mathcal{U})$.*

2.2. **Local obstructions and criteria for convexity.** One way to detect non-convexity in a neural code is through the existence of what are known as local obstructions.

Definition 2.10. *Given $\mathcal{C}$ a code on n neurons, we define a **local obstruction** to be a pair (σ, τ) satisfying*

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

where $\tau \neq \emptyset$ and $\text{Lk}_\sigma(\Delta(\mathcal{C})|_{\sigma \cup \tau})$ is not contractible.

The reason local obstructions are named as such is due to the following result, proved by Curto et al. in [1].

Theorem 2.11. *If $\mathcal{C}$ has a local obstruction, then $\mathcal{C}$ is not a good cover code, and therefore not convex.*

Now if a code $\mathcal{C}$ has no local obstructions, we say that it is **locally good**. The next theorem gives a much more efficient way of characterizing local obstructions. As it turns out we only need to check the link of faces that are intersections of maximal faces with respect to $\Delta(\mathcal{C})$ as a whole.

Definition 2.12. *Given a simplicial complex Δ, we define $\mathcal{M}(\Delta) = \{\sigma \in \Delta \mid \mathrm{Lk}_\sigma(\Delta)$ is not contractible $\}$. $\mathcal{M}(\Delta)$ is called the set of* ***mandatory codewords*** *for any $\mathcal{C}$ such that $\Delta(\mathcal{C}) = \Delta$.*

Theorem 2.13. *Let $\mathcal{C}$ be a code. $\mathcal{C}$ is locally good if and only if it contains all its mandatory codewords. Furthermore, every mandatory codeword is an intersection of maximal codewords.*

See [1] for a detailed proof.

Corollary 2.14. *Max-intersection codes are locally good.*

In fact, max-intersection codes are convex, a result that was proved in [3].

Now Theorem 2.8 tells us that

$$\mathcal{C} \text{ is convex } \Rightarrow \mathcal{C} \text{ is a good cover code } \Rightarrow \mathcal{C} \text{ is locally good,} \tag{2.2}$$

and it is natural to ask whether the converse of either implication could be true. Lienkaemper et al. gave the first known counterexample to the conjecture that convex codes are locally good in [5].

Theorem 2.15. *The neural code $\mathcal{C} = \{2345, 123, 134, 145, 13, 14, 23, 34, 45, 3, 4, \emptyset\}$ is locally good and nonconvex.*

We will discuss this counterexample in more detail later in Section 5, as it turns out this *is* realizable if we use closed sets instead of open sets. We prove the converse to the latter implication of (2.2) in the next section.

We state one final result that for codes with the same simplicial complex, convexity is a monotone property with respect to inclusion.

Theorem 2.16. *If $\mathcal{C}$ is convex, and $\mathcal{C} \subseteq \mathcal{C}'$ where $\Delta(\mathcal{C}) = \Delta(\mathcal{C}')$, then $\mathcal{C}'$ is convex.*

This is proved in [3].

3. Locally good codes and good covers

In this section we will prove that the notions of being locally good and a good cover code are equivalent. To sketch the idea, we will be considering a generalization of $|\Delta(\mathcal{C})|$ where all codewords in $\Delta(\mathcal{C})$ but not in $\mathcal{C}$ correspond to a deleted "open face" in $\Delta(\mathcal{C})$.

We define the **code complex** of a code $\mathcal{C}$, $|\mathcal{C}|$, to be a modification of $|\Delta(\mathcal{C})|$ by deleting the interiors of faces of codewords not in $\mathcal{C}$. Recall that $[\tau]$ denotes the topological realization of τ as a closed simplex of dimension $|\tau| - 1$ that is contained in $|\Delta|$. Similarly, we define (τ) to be the interior of the topological space $[\tau]$. In particular, (τ) is a point if $|\tau| = 1$ and homeomorphic to the the open ball of dimension $|\tau| - 1$ otherwise. Now we formally define the code complex to be

$$|\mathcal{C}| = |\Delta(\mathcal{C})| \setminus \bigcup_{\tau \in \Delta(\mathcal{C}) \setminus \mathcal{C}} (\tau).$$

For example, the code complex of $\mathcal{C} = \{123, 12, 23, 1, 2, \emptyset\}$ is a 2-simplex with a point and an open edge missing, where the missing point corresponds to the lack of the codeword 3 and the missing edge corresponds to the lack of the codeword 13. It should be clear that $|\Delta(\mathcal{C})| = \overline{|\mathcal{C}|}$. We also note the following

Remark 3.1. $|\mathcal{C}| = \bigsqcup_{\sigma\in\mathcal{C}}(\sigma)$. *That is, even though we defined $|\mathcal{C}|$ by deleting faces, it can also be built from the disjoint union of all the faces in $\mathcal{C}$ in the manner that a CW-complex is constructed.*

The proof of the main theorem has two main steps. The first is to generate a realization of $\mathcal{C}$ with sets $\mathcal{V} = \{V_1, ..., V_n\}$ where $\mathcal{V}$ is an "almost good cover" in the sense that the sets are not necessarily open (but still have empty or contractible intersections). We then show how this construction extends naturally to a good cover realization with open sets.

Theorem 3.2. *A code is locally good if and only if it is a good-cover code.*

Proof. The backwards direction is already established (Theorem 2.8), so we will prove the forward direction.

Let $\mathcal{C}$ be a locally good code on n neurons. First, we will give a realization of $\mathcal{C}$ with sets $V_1, ..., V_n$ that are not necessarily open, but satisfy V_τ being either empty or contractible for all $\tau \subseteq [n]$.

As $|\mathcal{C}|$ is the union of all open faces (τ) where $\tau \in \mathcal{C}$ as noted in in Remark 3.1, we set

$$V_i = \bigcup_{\tau\in\mathcal{C}, i\in\tau} (\tau).$$

For each $\tau \subseteq [n]$, we now need to check that V_τ is either empty or contractible. If V_τ is nonempty and $\tau \in \mathcal{C}$, V_i can deformation retract to (τ), which is contractible. It remains to check the cases when V_τ is nonempty and $\tau \notin \mathcal{C}$. First note that because $\mathcal{C}$ is locally good, then for any such τ, $Lk_\tau(\Delta(\mathcal{C}))$ is contractible. Now if τ was the only codeword in $\Delta(\mathcal{C}) \setminus \mathcal{C}$, our code complex $|\Delta(\mathcal{C}) \setminus \tau|$ is exactly $|\Delta| \setminus (\tau)$, and it is clear that V_τ deformation retracts to $Lk_\tau(\mathcal{C})$ which is contractible. In the general case, however, we note that other open faces could be missing, including faces in the link. Luckily, it is the case that given a contractible simplical complex and faces $\tau_1, ..., \tau_n$ where deleting any single one of the open faces (τ_i) preserves the homotopy type, we can remove all such open faces together from the simplicial complex while preserving the homotopy type, a fact that we relegate to Lemma 3.3. From this, we have completed the first step of the proof of achieving a non-open good cover realization.

We now go from $|\mathcal{C}|$ to an actual good cover realization through a "reverse deformation retract" of $|\mathcal{C}|$ in the following way. First, we realize our code complex $|\mathcal{C}|$ (which is a union of V_i's as described above) as a subset of the n-simplex in $\mathbb{R}^n$. Now each maximal face (of which there are $n + 1$) of the simplex uniquely defines a hyperplane of dimension $n - 1$. Moreover, the $n + 1$ hyperplanes partition $\mathbb{R}^n$ into 2^{n-1} regions, namely the interior of the n-simplex and regions that deformation retract onto each open face of the n-simplex. Letting R_τ be the region that deformation retracts onto (τ), we can set

$$W_i = \bigcup_{\tau\in\mathcal{C}, i\in\tau} R_\tau.$$

Now we take U_i to be the interior of W_i, as W_i may have limit points (see Figure 4). $U_1, ..., U_n$ is thus a good cover realization of $\mathcal{C}$. Additionally, if we wanted the sets to be finite, we could intersect each U_i with an open ball (of fixed size) containing the entire n-simplex. It is easy to see that each U_τ is homotopy equivalent to V_τ in our realization $|\mathcal{C}|$, completing the proof of the theorem. □

Lemma 3.3. *Given a contractible simplicial complex $|\Delta|$, suppose we have a collection of faces $\tau_1, ..., \tau_n$ where $|\Delta \setminus \tau_i|^2$ is still contractible for each $i \in [n]$. Then $|\Delta \setminus (\tau_1 \cup ... \cup \tau_n)|$ must still be contractible.*

[2] Note that we are taking $\Delta \setminus \tau$ as a neural code; by definition $|\Delta \setminus \tau| = |\Delta| \setminus (\tau)$.

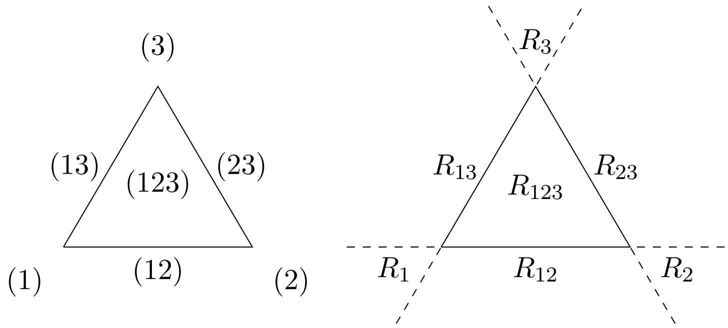

FIGURE 4. We show the correspondence between a 2-simplex and the regions that deformation retract to the open faces that make up the 2-simplex. In particular, note that R_{12} contains (12) and R_1 contains the dashed lines and the vertex (1) (making it in fact closed), which is why passing from W_i to its interior U_i is usually necessary.

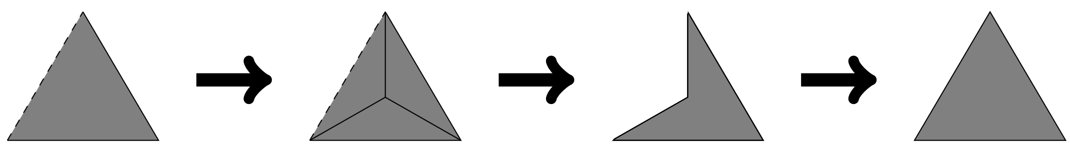

FIGURE 5. We start with a 2-simplex missing an open edge. We can deformation retract this to the third figure, which then is homeomorphic to the 2-simplex with the missing edge filled in. Note that if the 2-simplex was missing other open faces, the deformation retract/homeomorphism does not disturb them since it passes only though the interior of the maximal face.

Proof. For any σ where $|\Delta \setminus \sigma|$ is contractible (note that σ cannot be a maximal face in Δ), we describe the following "filling" procedure (see Figure 5 for an example). Let $M_1, ..., M_k$ be the maximal faces of Δ containing τ. We let $|M_i|$ denote the topological realization of each face M_i as a simplex contained in $|\Delta|$, and denote $b_{i,1}, ..., b_{i,\ell}$ be the barycenters of all faces η in M_i such that $\sigma \subseteq \eta$. Let $D_\sigma = \text{conv}(\sigma, b_{i,1}, ..., b_{i,\ell})$[3]. There is a natural deformation retract from $|M_i| \setminus (\sigma)$ to $\overline{|M_i| \setminus D_\sigma}$[4], the latter of which also has an natural homeomorphism back to $|M_i|$, "filling" up the deleted face (σ). We note that both the deformation retract and homeomorphism passes through faces strictly larger than σ. Moreover, this filling procedure can be done simultaneously with all M_i.

Now we take $|\Delta \setminus (\tau_1 \cup ... \cup \tau_n)|$, and without loss of generality suppose $|\tau_1| \geq |\tau_2| \geq ... \geq |\tau_n|$. We can fill in the deleted face (τ_1) by the described deformation retract and homeomorphism without affecting any of the other τ_i's, which gives us the homotopy equivalence

$$|\Delta \setminus (\tau_1 \cup ... \cup \tau_n)| \simeq |\Delta \setminus (\tau_2 \cup ... \cup \tau_n)|.$$

Repeating this $n-1$ more times for the faces $\tau_2, ..., \tau_n$ gives a homotopy equivalence between $|\Delta \setminus (\tau_1 \cup ... \cup \tau_n)|$ and Δ, which we know is contractible. □

[3] conv is the convex hull

[4] $\overline{|M_i| \setminus D_\sigma}$ is in fact the first barycentric subdivision of M_i with all sub-simplexes that have nonempty intersection with (σ) removed.

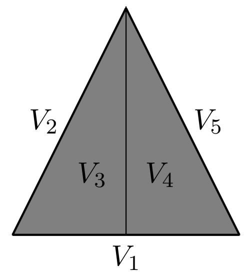

FIGURE 6. V_1, V_2, and V_5 are the closed line segments, while V_3 and V_4 are the closed triangles in the above realization. This is a closed convex realization of the counterexample code, $\mathcal{C} = \{2345, 123, 134, 145, 13, 14, 23, 34, 45, 3, 4, \emptyset\}$ [5].

4. NEURAL CODES AND COMPUTABILITY

Given this equivalence between being a good cover code and locally good (Theorem 3.2), we now show that both are undecidable problems. The reduction is quite simple and hinges on the undecidability of determining whether a homology ball (of sufficiently high dimension) is contractible:

Lemma 4.1 (Tancer [8])**.** *It is undecidable whether a given 4-dimensional simplicial complex is contractible.*

Theorem 4.2. *The problem of deciding whether a 6-sparse code has a good cover is undecidable.*

Proof. Given any 4-dimensional simplicial complex Δ, consider the cone over Δ on a new vertex **v**. This is itself a simplicial complex, which we denote by Δ'. Now let $\mathcal{C}$ be the 5-sparse neural code $\Delta' \setminus \mathbf{v}$. The only possible local obstruction is at **v**, so our code has a good cover if and only if $\Delta = Lk_v(\Delta')$ is contractible. Thus, any algorithm that could decide whether $\mathcal{C}$ is locally good would also decide whether Δ is contractible, which is impossible by Lemma 4.1. □

We note that it is decidable to tell whether a graph is contractible, so we can always tell whether a 3-sparse code is locally good. The problem of whether we can lower the dimension in Lemma 4.1 to 3 or 2 is still open (which would let us to drop the dimension in Theorem 4.2 to 5 or 4, respectively). One might ask in the case where $\mathcal{C}$ is a cone on a contractible simplicial complex without the vertex like the one described in the proof, whether $\mathcal{C}$ is convex (if this were the case, we would immediately get undecidibility of the convexity problem by the same argument). This is not always the case, which we will see in the next section.

5. CLOSED CONVEX REALIZATIONS AND A NEW OBSTRUCTION TO CONVEXITY

In [5] the first counterexample to the conjecture that “Locally good codes are convex” is given. The example in question, however, is realizable by convex *closed* sets as shown in Figure 6.

Note that in this case, there are sets that have measure zero, which is not possible in an open realization. It would be natural for one to ask whether a locally good code might always have a closed convex realization, but the answer again is no.

We present a new type of obstruction to convexity, that is, a code that is a good cover code but not convex, even if we allow for closed realizations.

Furthermore, the counterexample code in [5] could be resolved in two different ways, that is, there are two distinct neural codes with the same simplicial complex as that of the counterexample code with neither being contained in the other. We present a new family of locally good nonconvex codes where each code $\mathcal{C}$ satisfies $|\Delta(\mathcal{C}) \setminus \mathcal{C}| = 1$, that is, it is missing only one codeword from its simplicial complex. By the monotonicity property, this means any code with the same simplicial complex as $\mathcal{C}$ must contain the missing codeword.

First we make note of an overloaded definition in the literature. The notion of a *collapse* of a simplicial complex was first posed by Whitehead in [11]. The term of *d-collpasibility* was coined by Wegner much later [9], and is a closely related idea.

Definition 5.1. *Let Δ be a simplicial complex and let $\mathcal{M}$ be its maximal faces. Now for any face σ such that there is a unique $\tau \in \mathcal{M}$ where $\sigma \subsetneq \tau$, we define*

$$\Delta' = \Delta \setminus \{\nu \in \mathcal{M} \mid \sigma \subseteq \nu\}$$

*and say that Δ' is an **elementary collapse** of Δ induced by σ, and give it the notation $\Delta \to \Delta'$. We call a sequence where*

$$\Delta \to \Delta_1 \to \Delta_2 \to ... \to \Delta_n$$

*a **collapse** of Δ to Δ_n. Moreover, Δ is **collapsible** if it collapses to a point.*

Definition 5.2. *Let Δ be a simplicial complex and let $\mathcal{M}$ be its maximal faces. For any face σ such that there is a unique $\tau \in \mathcal{M}$ where $\sigma \subseteq \tau$ (note the difference from Definition 5.1 here), we define*

$$\Delta' = \Delta \setminus \{\nu \in \mathcal{M} \mid \sigma \subseteq \nu\}$$

*and say that Δ' is an **elementary d-collapse** of Δ induced by σ (d refers to the constraint that dim $\sigma \leq d$, but we will not be concerned with any particular value of d here and let d be arbitrarily high when we invoke the term "d-collapsibility"). We call a sequence of elementary d-collapses starting with Δ and ending with Δ'' a **d-collapse** of Δ to Δ''.*

While it may seem like a subtle difference between the two definitions above, it is important to note that the homotopy type is preserved throughout a collapse, but not necessarily throughout a d-collapse, as we see in the following lemma.

Lemma 5.3. *Let Δ be contractible and $\Delta \to \Delta'$ an elementary d-collapse induced by σ where Δ' is nonempty. σ is a maximal face of Δ if and only if Δ' is not contractible.*

Proof. We make use of the fact that a simplicial complex is contractible if and only if its fundamental group along with all reduced homology groups vanish. For the forward direction, suppose that σ is a maximal face. $[\sigma] \cap |\Delta'|$ is the boundary of $[\sigma]$, which is homeomorphic to a sphere, and $[\sigma]$ is naturally contractible because it is a simplex. Applying the Mayer-Vietoris sequence for homology to $|\Delta| = [\sigma] \cup |\Delta'|$, we see that $|\Delta'|$ must have a non-vanishing homology group, and therefore cannot be contractible. On the other hand, suppose σ is not a maximal face. Then it is not hard to see that $[\sigma] \cap |\Delta'|$ must be contractible, and Mayer-Vietoris tells us that $|\Delta'|$ must have vanishing reduced homology groups An analogous application of the Seifert-van Kampen theorem tells us that the fundamental group of $|\Delta'|$ must vanish as well, so $|\Delta'|$ is contractible. □

Corollary 5.4. *The following are equivalent*

(1) *Δ is collapsible.*
(2) *There exists a d-collapse of Δ to a simplex such that none of the elementary d-collapses are induced by a maximal face.*
(3) *There exists a d-collapse of Δ to a simplex that preserves the homotopy type of Δ.*

While all collapsible simplicial complexes are contractible by definition, a counterexample to the converse statement is any triangulation of the 2-dimensional topological space known as "Bing's house with two rooms."

We can now prove the following theorem.

Theorem 5.5. *Let Λ be a contractible simplicial complex. Set Δ to be a cone of Λ over a new vertex $\mathbf{v}$ and let $\mathcal{C} = \Delta \setminus \{\mathbf{v}\}$. If $\mathcal{C}$ is convex, then Λ is collapsible.*

Proof. Let $\mathcal{C}$ be a convex code on $n+1$ neurons realized by sets $U_1, ..., U_n, U_{\mathbf{v}}$, where $U_1, ..., U_n$ realize Λ (taken itself as a neural code). Now it is necessarily the case that $U_{\mathbf{v}} = \bigcup_{i \in [n]} U_i$. Consider sliding a hyperplane along a line across the ambient space of the realization (where $U_1, ..., U_n$ are realized) and deleting everything on one side. We do this until an intersection region corresponding to a maximal face in Λ has been removed. This deletion will remove a maximal face in Λ, and in fact induces an elementary d-collapse $\Lambda \to \Lambda'$. By the nerve theorem the homotopy type of Λ' is the same as that of $U_{\mathbf{v}}$ with part of it deleted by the hyperplane. Because the convexity of $U_{\mathbf{v}}$ is preserved when we delete part of it with the hyperplane sliding, it remains contractible, so Λ' is contractible. Repeating this hyperplane-sliding deletion procedure until only a single maximal face in Λ remains induces to a d-collapse of Λ to Λ'' where Λ'' is a simplex. Now the homotopy type of Λ is preserved throughout the d-collapse to Λ''. Corollary 5.4 tells us that that Λ is collapsible, completing the proof. The explicit construction of sweeping the hyperplane and the resulting d-collapse is given in [9] and elaborated upon in [7], so we will not repeat the details. □

Example 5.6. *Let Δ be a cone over a triangulation of Bing's house, which is a 2-dimensional simplicial complex. Now take $\mathcal{C}$ to be Δ missing the codeword corresponding to the vertex of the cone. $\mathcal{C}$ is a good cover code but not convex.*

The following strengthening of the previous theorem gives us a new class of local obstructions.

Theorem 5.7. *For any convex neural code $\mathcal{C}$ where $\Delta(\mathcal{C}) = \Delta$, if $\sigma \in \Delta \setminus \mathcal{C}$, then $Lk_\sigma(\Delta)$ is collapsible.*

Proof. Suppose $\mathcal{C}$ is convex and we have $\sigma \in \Delta \setminus \mathcal{C}$. In particular, it must be the case that U_σ is convex. We note that $U_\sigma = \bigcup_{\tau \supsetneq \sigma} U_\tau$. Now consider the neural code $\mathcal{C}' = \{\tau \mid \sigma \subseteq \tau\}$. Such a code must also be convex. Moreover, because each codeword in $\mathcal{C}$ includes σ, we can replace σ with a vertex $\mathbf{v}$. Now note that $\Delta(\mathcal{C}')$ is precisely a cone over $Lk_\sigma(\Delta(\mathcal{C}))$, and $\mathcal{C}' = \Delta(\mathcal{C}') \setminus \mathbf{v}$. By the previous theorem, we have a contradiction. □

To make the distinction between the For any $\sigma \in \Delta(\mathcal{C}) \setminus \mathcal{C}$ such that $Lk_\sigma(\Delta(\mathcal{C}))$ is not collapsible, we say that σ is a **local obstruction of the second kind**. Moreover, if $\mathcal{C}$ has no local obstructions of the second kind, we say that $\mathcal{C}$ is **locally great**.

Remark 5.8. *Theorem 5.7 still fails to give a sufficient condition for convexity, precisely by the counterexample code from [5].*

Remark 5.9. *The question of whether a simplicial complex is collapsible is decidable, and is in fact NP-complete in general [8].*

As it turns out, our proof of the undecidability of the good-cover problem does not directly extend to give us any results about the decidability of whether a code is convex, as we have presented a strictly stronger criterion (than that of being locally good) of being locally great that is in fact decidable.

6. Discussion

In summary, we now know the following:

$$\mathcal{C} \text{ is convex } \Rightarrow \mathcal{C} \text{ is locally great } \Rightarrow \mathcal{C} \text{ is a good cover code } \Leftrightarrow \mathcal{C} \text{ is locally good.}$$

In particular, it is undecidable to tell whether an arbitrary neural code $\mathcal{C}$ is locally good/a good cover code, decidable to tell whether $\mathcal{C}$ is locally great, and the question remains open for determining whether $\mathcal{C}$ is convex.

We note that based on Theorem 5.7, any code where the link of a missing face is not collapsible (but still contractible) requires that the code be at least 4-sparse, like the counterexample code discovered in [5]. As we know that convexity is equivalent to being locally good for 2-sparse codes, the question of whether 3-sparse locally good codes are convex remains open. Indeed, it is not too hard to see that for $\mathcal{C}$ a 3-sparse code, the only mandatory codewords are ones of size one, and that the links of any vertex in $\Delta(\mathcal{C})$ is a graph, where contractibility and collapsibility are equivalent.

Neural codes have also been studied from an algebraic standpoint through neural ideals and rings, which are closely related to the Stanley-Reisner ring/ideal of a simplicial complex [2]. Using these algebraic tools, it is possible to find local obstructions that can be detected by homology, which has sufficed to determine contractibility for small simplicial complexes. It may be worthwhile to see if collapsibility can be characterized in a similar way, as unlike contractibility the former is decidable.

The decision problem of convexity is also still unresolved. Indeed, as being convex is a stronger notion than being locally good, it is entirely possible that the decision problem of determining whether a code is complex is decidable. It would also be interesting to investigate how the theory changes when considering closed verses open realizations, as we now have codes that are neither closed nor open realizable, codes that are closed but not open realizable, and codes that are open and closed realizable. We conclude with the following conjecture:

Conjecture 6.1. *If $\mathcal{C}$ is locally great, then $\mathcal{C}$ is realizable by convex sets in Euclidean space (note that we make no assumptions on closedness or openness here).*

Acknowledgements

This work was conducted at a 2016 REU at Texas A&M, "Algebraic Methods in Computational Biology," which was supported by NSF DMS-1460766. The author would like to thank Anne Shiu for her guidance and insight. Additionally, Martin Tancer provided some very insightful comments on convex intersection theory, including an explanation of the hyperplane sliding argument.

References

[1] Carina Curto, Elizabeth Gross, Jack Jeffries, Katie Morrison, Mohamed Omar, Zvi Rosen, Anne Shiu, and Nora Youngs. What makes a neural code convex? Available at `arXiv:1508.00150`.

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

[3] Chad Giusti, Vladimir Itskov, and William Kronholm. On convex codes and intersection violators. In preparation.

[4] R. Amzi Jeffs, Mohamed Omar, Natchanon Suaysom, Aleina Wachtel, and Nora Youngs. Sparse neural codes and convexity. Available at `arXiv:1511.00283`, 2015.

[5] Caitlin Lienkaemper, Anne Shiu, and Zev Woodstock. Obstructions to convexity in neural codes. Available at `arXiv:1509.03328`, 2015.

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

[7] Martin Tancer. Intersection patterns of convex sets via simplicial complexes: a survey. In *Thirty essays on geometric graph theory*, pages 521–540. Springer, New York, 2013.
[8] Martin Tancer. Recognition of collapsible complexes is NP-complete. *Discrete & Computational Geometry*, 55(1):21–38, 2016.
[9] Gerd Wegner. d-collapsing and nerves of families of convex sets. *Archiv der Mathematik*, 26(1):317–321, 1975.
[10] André Weil. Sur les théoròmes de de rham. *Commentarii mathematici Helvetici*, 26:119–145, 1952.
[11] J.H.C. Whitehead. Simplical spaces, nuclei and m-groups. *Proceedings of the London Mathematical Society*, 45:243–327, 1938.

E-mail address: ahc232@cornell.edu

DEPARTMENT OF MATHEMATICS, CORNELL UNIVERSITY, ITHACA, NY 14850