

NEURAL CODES, DECIDABILITY, AND A NEW LOCAL OBSTRUCTION TO CONVEXITY

AARON CHEN, FLORIAN FRICK, AND ANNE SHIU

ABSTRACT. Given an intersection pattern of arbitrary sets in Euclidean space, is there an arrangement of convex open sets in Euclidean space that exhibits the same intersections? This question is combinatorial and topological in nature, but is motivated by neuroscience. Specifically, we are interested in a type of neuron called a place cell, which fires precisely when an organism is in a certain region, usually convex, called a place field. The earlier question, therefore, can be rephrased as follows: Which neural codes, that is, patterns of neural activity, can arise from a collection of convex open sets? To address this question, Giusti and Itskov proved that convex neural codes have no “local obstructions,” which are defined via the topology of a code’s simplicial complex. Codes without local obstructions are called locally good, because the obstruction precludes the code from encoding the intersections of open sets that form a good cover. In other words, every good-cover code is locally good. Here we prove the converse: Every locally good code is a good-cover code. We also prove that the good-cover decision problem is undecidable. Finally, we reveal a stronger type of local obstruction that prevents a code from being convex, and prove that the corresponding decision problem is NP-hard. Our proofs use combinatorial and topological methods.

Keywords: convex, good cover, simplicial complex, decidability, nerve lemma, neural code, place cell


MSC codes: 05E45, 55U10, 92C20

1. INTRODUCTION

This work addresses the following question: Which binary codes arise from the regions cut out by a collection of convex open sets in some Euclidean space? One such code is

$$\mathcal{C} = \{1110, 0111, 1100, 0110, 1010, 0011, 1000, 0100, 0010, 0001\} = \{123, 234, 12, 23, 13, 24, 34, 1, 2, 3, 4\},$$

which arises from the following convex open sets $U_1, U_2, U_3,$ and U_4 :



Here, some of the regions are labeled by the corresponding *codewords* in \mathcal{C} . We can view each codeword as a vector in $\{0, 1\}^4$ or as its support set, here a subset of $\{1, 2, 3, 4\}$.

A closely related question is: Which *intersection patterns* arise from a collection of convex sets? This question asks only which sets U_i intersect, and not whether, for instance, there is a region where U_1 and

Date: September 25, 2018.

U_2 intersect outside of U_3 . This topic – intersection patterns of convex sets – has been studied extensively (see [24] for an overview), but the first question we posed has caught attention only recently [3–9, 12, 15, 17, 21, 31].

The recent interest in this area is motivated by neuroscience, specifically from the study of neurons called *place cells*. The discovery of place cells by O’Keefe *et al.* in 1971 was a major breakthrough that led to a shared 2014 Nobel Prize in Medicine or Physiology [20]. A place cell encodes spatial information about an organism’s surroundings by firing precisely when the organism is in the corresponding *place field*. In this context, a codeword represents the neural firing pattern that occurs when the organism is in the corresponding region of its environment: the i th coordinate is 1 if and only if the organism is in the place field of neuron i . The resulting set of codewords is called a *neural code*.

Place fields can be modeled by convex open sets [5], so we are interested in the following restatement of the question that opened this work: Which neural codes can arise from a collection of convex open sets? To address this problem, Giusti and Itskov identified a *local obstruction*, defined via the topology of a code’s simplicial complex, and proved that convex neural codes have no local obstructions [8]. Codes without local obstructions are called *locally good*, as the obstruction prevents the code from encoding the intersections of open sets that form a good cover (for instance, if the sets are convex). If such a good cover exists (for instance, from a collection of convex open sets), then the code is a *good-cover code*. Thus, we have:

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

The converse of the first implication is false [17]. The second implication is the starting point of our work. We prove that the implication is in fact an equivalence: every locally good code is a good-cover code (Theorem 3.13). We also prove that the good-cover decision problem is undecidable (Theorem 4.3).

Next, we discover a new, stronger type of local obstruction that precludes a code from being convex (Theorem 5.10). Like the prior obstruction, the new obstruction is defined in terms of a code’s simplicial complex, but in this case the link of “missing” codewords must be “collapsible” (which is implied by “contractible”, the condition in the original type of obstruction). We call codes without the new obstruction *locally great*, and examine the corresponding decision problem. We prove that the locally-great decidability problem is decidable, and in fact NP-hard (Theorem 5.18).

Thus, our results refine the implications we saw earlier, as follows:

$$\begin{array}{ccccccc} \mathcal{C} \text{ is convex} & \Rightarrow & \mathcal{C} \text{ is locally great} & \Rightarrow & \mathcal{C} \text{ is a good-cover code} & \Leftrightarrow & \mathcal{C} \text{ is locally good.} \\ & & \text{(NP-hard problem)} & & \text{(undecidable)} & & \end{array}$$

Finally, we add another implication to the end of those listed above, by noting that every locally good code can be realized by connected open sets, but not vice-versa (Proposition 2.19). Taken together, our results resolve fundamental questions in the theory of convex neural codes.

The outline of our work is as follows. Section 2 provides background on neural codes, local obstructions, and criteria for convexity. In Sections 3–5, we prove the results listed above, using classical tools from topology and combinatorics. Finally, our discussion in Section 6 lists open questions arising from our work.

2. BACKGROUND

Here we introduce notation as well as basic definitions in the theory of neural codes.

We define $[n] := \{1, 2, \dots, n\}$. We will reserve lowercase Greek letters (e.g., σ and τ) to denote subsets of $[n]$ (for some n). Such a subset usually refers to a codeword in a neural code (Definition 2.1) or a

face in a simplicial complex (Definition 2.6). For shorthand, we will omit the braces and commas; e.g., if $\tau = \{1, 2, 3\}$ and $\sigma = \{2, 3, 4\}$, we write $\tau = 123$, $\sigma = 234$, and $\tau \cap \sigma = 23$.

2.1. Codes, simplicial complexes, and the nerve theorem. Given a collection $\mathcal{U} = \{U_1, U_2, \dots, U_n\}$ of sets (place fields) in some *stimulus space* $X \subseteq \mathbb{R}^d$ and some $\tau \subseteq [n]$, let $U_\tau := \bigcap_{i \in \tau} U_i$, where $U_\emptyset := X$.

Definition 2.1.

- (1) A **neural code** \mathcal{C} on n neurons is a subset of $2^{[n]}$, and each $\sigma \in \mathcal{C}$ is a **codeword**. Any codeword that is maximal in \mathcal{C} with respect to set inclusion is a **maximal codeword**.
- (2) A code \mathcal{C} is **realized** by a collection of sets $\mathcal{U} = \{U_1, U_2, \dots, U_n\}$ in a stimulus space $X \subseteq \mathbb{R}^d$ if

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

Conversely, given a collection of sets \mathcal{U} , let $\mathcal{C}(\mathcal{U})$ denote the unique code realized by \mathcal{U} , via (2.1).

Remark 2.2. A neural code is often referred to as a *hypergraph* in the literature. Also, given a collection of subsets $U_1, U_2, \dots, U_n \subseteq X$, the neural code that is realized by these subsets can be defined as the collection of sets $U(x) = \{i \in [n] \mid x \in U_i\}$, where x varies over X .

Every neural code can be realized by open sets [5] and by convex sets [6]. We are interested, however, in realizing neural codes by sets that are both open and convex. This is because (biological) place fields are approximately convex and have positive measure, which are properties captured by convex open sets.

Definition 2.3. A neural code \mathcal{C} is:

- (1) **convex** if \mathcal{C} can be realized by a collection of convex open sets.
- (2) a **good-cover code** if \mathcal{C} can be realized by a collection of contractible open sets \mathcal{U} such that every nonempty intersection of sets in \mathcal{U} is also contractible. Such a \mathcal{U} is called a **good cover**.
- (3) **connected** if \mathcal{C} can be realized by a collection of connected open sets.

Example 2.4. We revisit the code $\mathcal{C} = \{123, 234, 12, 23, 13, 24, 34, 1, 2, 3, 4, \emptyset\}$ from the Introduction, where we saw that \mathcal{C} is convex. Hence, \mathcal{C} is a good-cover code and also a connected code.

Connected codes were classified recently by Mulas and Tran (not every code is connected) [19]. Good-cover codes are connected, but not vice-versa (see Proposition 2.19 later in this section).

Remark 2.5 (Codes and the empty set). A code \mathcal{C} is convex (respectively, a good-cover code) if and only if $\mathcal{C} \cup \{\emptyset\}$ is convex (respectively, a good-cover code) [4]. Indeed, if $\emptyset \notin \mathcal{C}$ and \mathcal{C} is realized by convex open sets (respectively, a good cover) \mathcal{U} , then $\mathcal{C} \cup \emptyset$ is realized by $\{U \cap B \mid U \in \mathcal{U}\}$, where B is an open ball that contains a point from each region cut out by \mathcal{U} . Conversely, if $\emptyset \notin \mathcal{C}$ and $\mathcal{C} \cup \emptyset$ is realized by \mathcal{U} (with respect to some stimulus space), then the code \mathcal{C} is realized by \mathcal{U} with respect to the stimulus space $X = \bigcup_{U \in \mathcal{U}} U$.

Convex codes on up to four neurons have been classified [4]. This classification was enabled by analyzing codes according to the simplicial complex they generate (see Definition 2.9).

Definition 2.6. An **abstract simplicial complex** Δ on $[n]$ is a subset of $2^{[n]}$ that is closed under taking subsets. Each $\sigma \in \Delta$ is a **face** of Δ . The **facets** of Δ are the faces that are maximal with respect to inclusion. The **dimension** of a face σ is $|\sigma| - 1$, and the **dimension** of a simplicial complex Δ , denoted by $\dim \Delta$, is the maximum dimension of the faces of Δ .

Every simplicial complex Δ can be realized geometrically in a Euclidean space of sufficiently high dimension, and we let $|\Delta|$ denote such a geometric realization (which is unique up to homeomorphism). Note that the dimension of a simplicial complex matches the dimension of its realization: $\dim \Delta = \dim |\Delta|$.

Definition 2.7. For a face σ of a simplicial complex Δ , the **restriction** of Δ to σ is the simplicial complex

$$\Delta|_{\sigma} := \{\tau \in \Delta \mid \tau \subseteq \sigma\}.$$

The **link** of σ in Δ is the simplicial complex:

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

Links are usually written as $\text{Lk}_{\Delta}(\sigma)$, instead of $\text{Lk}_{\sigma}(\Delta)$, but, following [4], we prefer to have σ in the subscript, because we often consider the link in several simplicial complexes. Note that $\sigma \cup \tau \in \Delta$ is the same as saying that $\sigma * \tau \subseteq |\Delta|$, where $*$ is the topological join.

Definition 2.8. Let Δ be a simplicial complex on $[n]$. The **cone over** Δ on a new vertex v is the following simplicial complex on $[n] \dot{\cup} \{v\}$:


$$\text{cone}_v(\Delta) := \{\sigma \cup \{v\} \mid \sigma \in \Delta\} \cup \Delta.$$

By construction, $\text{Lk}_v(\text{cone}_v(\Delta)) = \Delta$. We will use this fact throughout our work.


Definition 2.9. For a code \mathcal{C} on n neurons, the **simplicial complex of** \mathcal{C} , denoted by $\Delta(\mathcal{C})$, is the smallest simplicial complex on $[n]$ that contains \mathcal{C} .

Note that for $\mathcal{C} = \mathcal{C}(\mathcal{U})$, we have $\sigma \in \Delta(\mathcal{C})$ if and only if $U_{\sigma} \neq \emptyset$ (cf. (2.1)). Also, the facets of $\Delta(\mathcal{C})$ are the maximal codewords of \mathcal{C} .

Example 2.10. The code $\mathcal{C} = \{12, 23, 1, 2, 3, \emptyset\}$ is convex and realized here:



The simplicial complex of the code, $\Delta(\mathcal{C})$, is realized here:



A related notion to $\Delta(\mathcal{C})$ is the nerve of a cover (see Remark 2.12), which we define now.

Definition 2.11. Given a collection of sets $\mathcal{U} = \{U_1, U_2, \dots, U_n\}$, the **nerve** of \mathcal{U} , denoted by $\mathcal{N}(\mathcal{U})$, is the simplicial complex on $[n]$ defined by:

$$\sigma \in \mathcal{N}(\mathcal{U}) \iff U_{\sigma} \neq \emptyset.$$

Remark 2.12. $\mathcal{N}(\mathcal{U}) = \Delta(\mathcal{C}(\mathcal{U}))$.

Next, we recall the classical result called the nerve theorem or nerve lemma [29]. The version we state is [10, Corollary 4G.3]:

Proposition 2.13 (Nerve theorem). *If \mathcal{U} is a finite collection of nonempty and contractible open sets that cover a paracompact space S such that every intersection of sets is either empty or contractible, then S is homotopy equivalent to $\mathcal{N}(\mathcal{U})$.*

Metric spaces are paracompact [23], so good-cover realizations of codes satisfy the hypotheses of Proposition 2.13. Thus, if we determine that a code does not satisfy the conclusion of this proposition, then we conclude it is not convex. We turn to this topic now.

2.2. Local obstructions and criteria for convexity. One way to detect non-convexity of a neural code is to find what is known as a local obstruction.

Definition 2.14. Given a code \mathcal{C} that is realized by open sets \mathcal{U} , a **local obstruction** is a pair (σ, τ) for which

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

where $\tau \neq \emptyset$ and $\text{Lk}_\sigma(\Delta(\mathcal{C})|_{\sigma \cup \tau})$ is *not* contractible. A code \mathcal{C} with *no* local obstructions is **locally good**.

Definition 2.14 does not depend on the choice of open sets \mathcal{U} [4].

The name “local obstruction” is due to the following result, which states that if a code has a local obstruction, then it is not a good-cover code, and therefore not convex.

Proposition 2.15 (Giusti and Itskov [8]; Curto *et al.* [4]).

$$\mathcal{C} \text{ is convex} \Rightarrow \mathcal{C} \text{ is a good-cover code} \Rightarrow \mathcal{C} \text{ is locally good.} \quad (2.2)$$

A natural question is: Do the converses of the implications in (2.2) hold? For the second implication, the converse is true; we will prove this in the next section (Theorem 3.13).

As for the first implication, the converse is in general false (but true for codes in which all codewords have size at most two [15] and codes on up to four neurons [4]). The first counterexample, which is a code on five neurons, was found by Lienkaemper, Shiu, and Woodstock [17]:

Proposition 2.16 (Counterexample code [17]). *The following neural code is a good-cover code, but non-convex: $\mathcal{C} = \{2345, 123, 134, 145, 13, 14, 23, 34, 45, 3, 4, \emptyset\}$.*

This code, it turns out, is realizable by *closed* convex sets instead of open convex sets [3].

Returning to the topic of detecting local obstructions, the next result gives a way to do so that is more efficient than simply applying the definition (Definition 2.14). As it turns out, we need to check only the links of faces that are intersections of facets of $\Delta(\mathcal{C})$.

Definition 2.17. The set of **mandatory codewords of a simplicial complex** Δ is $\mathcal{M}(\Delta) := \{\sigma \in \Delta \mid \text{Lk}_\sigma(\Delta) \text{ is not contractible}\}$; and the set of **non-mandatory codewords of** Δ is $\Delta \setminus \mathcal{M}(\Delta)$. The set of **mandatory codewords of a code** \mathcal{C} is $\mathcal{M}(\Delta(\mathcal{C}))$.

Proposition 2.18 (Curto *et al.* [4]). *A code \mathcal{C} is locally good if and only if it contains all its mandatory codewords (i.e., $\mathcal{M}(\Delta(\mathcal{C})) \subseteq \mathcal{C}$). Also, every mandatory codeword is an intersection of maximal codewords.*


As a corollary, *max-intersection complete* codes, those that are closed under taking intersections of maximal codewords, are locally good. In fact, these codes are convex [3]. Note that the (non-convex) counterexample code \mathcal{C} from Proposition 2.16 is *not* max-intersection complete, because the intersection of maximal codewords $123 \cap 134 \cap 145 = 1$ is missing from \mathcal{C} .

Another result pertaining to convexity, due to Cruz *et al.* [3], is as follows: for codes with the same simplicial complex, convexity is a monotone property with respect to inclusion. That is, *if \mathcal{C} is convex, and $\mathcal{C} \subseteq \mathcal{C}'$ where $\Delta(\mathcal{C}) = \Delta(\mathcal{C}')$, then \mathcal{C}' is convex.*


We end this section by showing that good-cover codes are realizable by connected open sets, but not vice-versa:

Proposition 2.19. *Every good-cover code is connected, but not every connected code is a good-cover code.*

Proof. Contractible sets are connected, so, by definition, good-cover codes are connected. As for the converse, consider the following code: $\mathcal{C} = \{124, 134, 234, 14, 24, 34, \emptyset\}$. The codeword 4, which is not in \mathcal{C} , is a mandatory codeword of \mathcal{C} , as $\text{Lk}_{\{4\}}(\Delta(\mathcal{C}))$ is the following non-contractible simplicial complex:



Thus, by definition, \mathcal{C} is not locally good, and so, by Proposition 2.18, is not a good-cover code. We complete the proof by displaying the following realization of \mathcal{C} by connected, open sets in \mathbb{R}^2 :



More precisely, for $i = 1, 2, 3, 4$, consider (the closures of) the regions above that are labeled by some σ for which $i \in \sigma$. Now let U_i be the interior of the union of all such regions. \square

3. LOCALLY GOOD CODES ARE GOOD-COVER CODES

In this section we will prove that being locally good is equivalent to being a good-cover code (Theorem 3.13). We accomplish this by constructing a good-cover realization of any locally good code. Our construction has two steps. We first realize *any* code via (not necessarily open) subsets of a geometric realization of its simplicial complex (Proposition 3.2), and then do a “reverse deformation retract” to obtain a realization by open sets (Proposition 3.10).

3.1. Code-complex realizations. The idea behind the following construction, which realizes any code \mathcal{C} by subsets of $|\Delta(\mathcal{C})|$, is to realize codewords in \mathcal{C} by the corresponding faces of $|\Delta(\mathcal{C})|$. Thus, we will simply delete faces corresponding to codewords that are *not* in \mathcal{C} . Accordingly, for any simplicial complex Δ and any $\sigma \in \Delta$, let $\text{int}(\sigma)$ denote the relative interior of the realization of the face σ within $|\Delta|$ (if v is a vertex, then $\text{int}(v) = \{v\}$). See the left-hand side of Figure 1. It follows that $|\Delta| = \dot{\cup}_{\emptyset \neq \sigma \in \Delta} \text{int}(\sigma)$. Thus, $|\Delta|$ is a CW-complex built from the $\text{int}(\sigma)$'s.

Definition 3.1. Let \mathcal{C} be a code on n neurons. For each $i \in [n]$, consider the following subset of $|\Delta(\mathcal{C})|$:

$$V_i := \bigcup_{\sigma \in \mathcal{C}} \text{int}(\sigma).$$

Then $\{V_1, V_2, \dots, V_n\}$ is the **code-complex realization** of \mathcal{C} .




FIGURE 1. The correspondence between the face-interiors $\text{int}(\sigma)$ that make up the 2-simplex and the regions R_σ that deformation retract to the $\text{int}(\sigma)$'s (see equation (3.4)). We define the region R_{12} to contain the face-interior labeled by (12), and R_1 to contain the dashed lines of its boundary and the vertex (1) (so it is closed), which is why in Proposition 3.10 we must pass from the unions of regions R_σ to their relative interiors.


Proposition 3.2. *Let \mathcal{C} be a code on n neurons. Then the code-complex realization $\{V_1, V_2, \dots, V_n\}$ of \mathcal{C} given in Definition 3.1 is a realization of \mathcal{C} . Moreover, if \mathcal{C} is locally good, then every nonempty intersection of sets from $\{V_1, V_2, \dots, V_n\}$ is contractible.*

We postpone the proof of Proposition 3.2 to the end of this subsection.


Remark 3.3. Code-complex realizations (Definition 3.1) are in general not simplicial complexes or even CW-complexes. See Examples 3.5 and 3.6, where the sets V_i and their unions are not even closed.

Remark 3.4. Code-complex realizations (Definition 3.1) are somewhat similar to prior constructions that also specify regions that correspond to codewords [3, 4] or intersection patterns [24].

Example 3.5. Consider the (locally good) code $\mathcal{C} = \{123, 12, 23, 1, 2, \emptyset\}$. The code-complex realization of \mathcal{C} given in Definition 3.1 is depicted here, along with $|\Delta(\mathcal{C})|$:



Example 3.6. For the (not locally good) code $\mathcal{C} = \{13, 23, 1, \emptyset\}$, the code-complex realization is:



Comparing Examples 3.5 and 3.6, note that, consistent with Proposition 3.2, the sets V_i in the first example are contractible and have contractible intersections, but not in the second example (V_3 is not connected).

The proof of Proposition 3.2 relies on the following lemma, which may be of independent interest.

Lemma 3.7. *Let Σ be a simplicial complex, and let Ω be a subset of the non-mandatory codewords of Σ . Then, if Σ is contractible, then so is the code-complex realization of the code $\Sigma \setminus \Omega$.*

We prove Lemma 3.7 in Appendix A.

Proof of Proposition 3.2. By construction, we have

$$\bigcup_{i=1}^n V_i = \bigcup_{\emptyset \neq \sigma \in \mathcal{C}} \text{int}(\sigma),$$

and the subset $\text{int}(\sigma)$ realizes the codeword σ . If $\emptyset \notin \mathcal{C}$, define the stimulus space $X = \bigcup_{i=1}^n V_i$; otherwise, define X to be the ambient Euclidean space. Thus, $\{V_1, V_2, \dots, V_n\}$ realizes \mathcal{C} with respect to X .

Now assume that \mathcal{C} is locally good. We must show that every intersection of the V_i 's is empty or contractible. To this end, let $\emptyset \neq \tau \subseteq [n]$, and assume that V_τ is nonempty. Note that

$$V_\tau = \bigcup_{\tau \subseteq \sigma \in \mathcal{C}} \text{int}(\sigma).$$

We consider two cases. If $\tau \in \mathcal{C}$, then $V_\tau = \text{int}(\tau) \cup \bigcup_{\tau \subsetneq \sigma \in \mathcal{C}} \text{int}(\sigma)$. Now consider the deformation retract of $|\Delta^n|$ to the face $\bar{\tau}$, arising from orthogonal projection to that face. It is straightforward to check that this restricts to a deformation retraction of V_τ to $\text{int}(\tau)$, which is contractible. Thus, V_τ is contractible.

Consider the remaining case, when $\tau \notin \mathcal{C}$. Then $V_\tau = \bigcup_{\tau \subsetneq \sigma \in \mathcal{C}} \text{int}(\sigma)$. We will show that V_τ is contractible by proving first that V_τ is homotopy equivalent to the code-complex realization of the following *link of a code* [3, Definition 2.3] or, for short, “code-link”:

$$\text{Lk}_\tau(\mathcal{C}) := \{\eta \subseteq [n] \mid \tau \cap \eta = \emptyset \text{ and } \tau \cup \eta \in \mathcal{C}\}, \quad (3.1)$$

and then proving that this code-complex realization, which we denote by \mathcal{L} , is contractible.

To see that $V_\tau \simeq \mathcal{L}$, we compute:

$$V_\tau = \bigcup_{\tau \subsetneq \sigma \in \mathcal{C}} \text{int}(\sigma) = \bigcup_{\emptyset \neq \eta \in \text{Lk}_\tau(\mathcal{C})} \text{int}(\tau \dot{\cup} \eta), \quad (3.2)$$

and

$$\mathcal{L} = \bigcup_{\emptyset \neq \eta \in \text{Lk}_\tau(\mathcal{C})} \text{int}(\eta). \quad (3.3)$$

It is straightforward to see, from (3.2) and (3.3), that V_τ deformation retracts to a copy of \mathcal{L} .

To see that \mathcal{L} is contractible, we will appeal to Lemma 3.7 (where $\Sigma = \text{Lk}_\tau(\Delta(\mathcal{C}))$ and $\Sigma \setminus \Omega$ is the code-link in equation (3.1)). To apply Lemma 3.7, we must show that every σ in $\text{Lk}_\tau(\Delta(\mathcal{C}))$ but not in the code-link (3.1) is *not* a mandatory codeword of $\text{Lk}_\tau(\Delta(\mathcal{C}))$. To this end, note that such a σ satisfies $\sigma \cap \tau = \emptyset$ and $\sigma \cup \tau \in \Delta(\mathcal{C}) \setminus \mathcal{C}$. Thus, because \mathcal{C} is locally good, the link $\text{Lk}_{\sigma \cup \tau}(\Delta(\mathcal{C})) = \text{Lk}_\sigma(\text{Lk}_\tau(\Delta(\mathcal{C})))$ is contractible, and thus σ is *not* a mandatory codeword of $\text{Lk}_\tau(\Delta(\mathcal{C}))$.

We now apply Lemma 3.7: the link $\text{Lk}_\tau(\Delta(\mathcal{C}))$ is contractible (because \mathcal{C} is locally good and $\tau \notin \mathcal{C}$), so \mathcal{L} also is contractible. Hence, V_τ is contractible. \square

Remark 3.8. In the proof of Proposition 3.2, the only place where we use the fact that \mathcal{C} is locally good, is at the end, when we show that \mathcal{L} is contractible.

3.2. Main result. Our next step is to modify, via a “reverse deformation retract”, the realization from Proposition 3.2 so that the sets are open (Proposition 3.10).

To this end, for any n , let $\Delta^n := 2^{[n]}$ be the complete simplicial complex on $[n]$, and consider an embedding of $|\Delta^n|$ in \mathbb{R}^{n-1} . The facet-defining hyperplanes of $|\Delta^n|$ form a hyperplane arrangement whose open chambers (in the complement of the hyperplanes) correspond to nonempty subsets σ of $[n]$ (see Figure 1). Let C_σ denote the closure of the chamber that corresponds to $\text{int}(\sigma)$. Choose an ordering of the nonempty faces of Δ^n so that they are nondecreasing in dimension: $0 = \dim \sigma_1 \leq \dim \sigma_2 \leq \dots \leq \dim \sigma_{2^n-1}$. We define recursively the following sets (see the right-hand side of Figure 1):

$$R_{\sigma_k} := C_{\sigma_k} \setminus \left(\bigcup_{j=1}^{k-1} C_{\sigma_j} \right). \quad (3.4)$$

We claim that the order of the σ_i 's does not matter. Indeed, if $\dim \sigma_i = \dim \sigma_j$ (with $i \neq j$), then the intersection $C_{\sigma_i} \cap C_{\sigma_j}$ is the (possibly empty) face $\sigma_i \cap \sigma_j$ of $|\Delta^n|$. This face $\sigma_i \cap \sigma_j = \sigma_l$ is indexed by some $l < i, j$, and so this face, regardless of whether $i < j$ or $j < i$, is neither in R_{σ_i} nor R_{σ_j} .

We make the following several observations. First, the R_σ 's are disjoint, and each R_σ is convex and full-dimensional. Also, R_σ deformation retracts to $\text{int}(\sigma)$ via a deformation retract of \mathbb{R}^{n-1} to $|\Delta^n|$. Finally, the interior of R_σ is the open chamber C_σ of the hyperplane arrangement of $|\Delta^n|$'s facet-defining hyperplanes.

In analogy to how we built the sets V_i from the sets $\text{int}(\sigma)$ (in Definition 3.1), we now build sets W_i from the sets R_σ (which, we noted, deformation retract to the sets $\text{int}(\sigma)$).

Proposition 3.9. *Let \mathcal{C} be a code on n neurons. For each $i \in [n]$, consider the following subset of \mathbb{R}^{n-1} :*

$$W_i := \bigcup_{\sigma \in \mathcal{C}} R_\sigma, \quad (3.5)$$

where R_σ is as in (3.4). Then $\{W_1, W_2, \dots, W_n\}$ is a realization of $\mathcal{C} \setminus \{\emptyset\}$.

Proof. By construction, we have

$$\bigcup_{i=1}^n W_i = \bigcup_{\emptyset \neq \sigma \in \mathcal{C}} R_\sigma,$$

and the subset R_σ realizes the codeword σ . Define the stimulus space $X = \bigcup_{i=1}^n W_i$. Then the collection $\{W_1, W_2, \dots, W_n\}$ realizes $\mathcal{C} \setminus \{\emptyset\}$ with respect to X . \square

The sets W_i in Proposition 3.9 are in general not open sets (see Figure 1), but they are unions of full-dimensional convex sets, so we now consider the interiors.


Proposition 3.10. *Let \mathcal{C} be a code on n neurons. For each $i \in [n]$, consider the following subset of \mathbb{R}^{n-1} :*

$$U_i := \text{int}(W_i),$$

where W_i is as in (3.5). Then $\{U_1, U_2, \dots, U_n\}$ is a realization of \mathcal{C} . Moreover, if \mathcal{C} is locally good, then $\{U_1, U_2, \dots, U_n\}$ is a good-cover realization of $\mathcal{C} \setminus \{\emptyset\}$.

Before proving Proposition 3.10, we give an example of the realization given in the proposition.

Example 3.11. We return to the locally good code $\mathcal{C} = \{123, 12, 23, 1, 2, \emptyset\}$ from Example 3.5. The good-cover realization $\{U_1, U_2, U_3\}$ of \mathcal{C} from Proposition 3.10, along with $U_1 \cup U_2 \cup U_3$, is depicted here:



Proof of Proposition 3.10. By construction, the U_i 's are open. So, we must show that (1) the U_i 's realize $\mathcal{C} \setminus \{\emptyset\}$, that is, $\mathcal{C}(\{U_i\}) = \mathcal{C} \setminus \{\emptyset\}$, and (2) if \mathcal{C} is locally good, then $\{U_1, U_2, \dots, U_n\}$ is a good cover.

Proof of (1). We begin by proving the containment $\mathcal{C}(\{U_i\}) \supseteq \mathcal{C} \setminus \{\emptyset\}$, where we define the stimulus space to be $X = \cup_{i=1}^n U_i$ (so, $\emptyset \notin \mathcal{C}(\{U_i\})$). Take $\emptyset \neq \sigma \in \mathcal{C}$. We must show that $\sigma \in \mathcal{C}(\{U_i\})$. Let x be in the interior of R_σ . Then, by construction, $x \in U_\sigma \setminus (\cup_{j \neq \sigma} U_j)$, so $\sigma \in \mathcal{C}(\{U_i\})$.

To prove the remaining containment, $\mathcal{C}(\{U_i\}) \subseteq \mathcal{C} \setminus \{\emptyset\}$, let $\sigma \in \mathcal{C}(\{U_i\})$. As explained above, $\sigma \neq \emptyset$. Let $x \in U_\sigma \setminus (\cup_{j \neq \sigma} U_j)$. We consider two cases. If x is *not* on a facet-defining hyperplane of $|\Delta^n|$, then x is in the interior of some R_τ . It is straightforward to check that $\tau = \sigma$, so $\text{int}(R_\sigma) \subseteq U_\sigma$, and thus $\sigma \in \mathcal{C}$.

In the second case, x is on exactly m facet-defining hyperplanes of $|\Delta^n|$ (for some $m \geq 1$). Crossing exactly one such hyperplane means going from some region R_α to some region $R_{\alpha \cup \{i\}}$ (for some $i \notin \alpha$), or vice-versa. Thus, a small neighborhood B of x intersects exactly 2^m regions R_τ : precisely those with

$$\eta \subseteq \tau \subseteq \tilde{\eta}, \quad (3.6)$$

for some $\eta \subseteq \tilde{\eta} \subseteq [n]$ with $|\eta| + m = |\tilde{\eta}|$.

Recall that x is in the open set U_σ , so $B \subseteq U_\sigma$, and thus the interiors of the 2^m regions R_τ are contained in U_σ . So, all sets τ given in (3.6) are in the code \mathcal{C} . Thus, to show that $\sigma \in \mathcal{C}$, it suffices to show that $\eta = \sigma$. The containment $\eta \supseteq \sigma$ follows from the fact that $\text{int}(R_\eta) \subseteq U_\sigma$ (as explained above). We prove $\eta \subseteq \sigma$ by contradiction. Assume that there exists $k \in \eta \setminus \sigma$. Then we have:

$$x \in B \subseteq \cup_{\eta \subseteq \tau \subseteq \tilde{\eta}} R_\tau \subseteq U_\eta \subseteq U_k.$$

Thus, $x \in U_k$, where $k \notin \sigma$, which contradicts the choice of x . So, $\eta \subseteq \sigma$ holds.

Proof of (2). Assume that \mathcal{C} is locally good. Let $\emptyset \neq \tau \subseteq [n]$, and assume that U_τ is nonempty. We must show that U_τ is contractible. By Proposition 3.2, V_τ is contractible, so it is enough to show that $U_\tau \simeq V_\tau$. We will prove this by showing that the set $Z_\tau := U_\tau \cup V_\tau$ deformation retracts to V_τ and also

weak deformation retracts [10] to U_τ . It is straightforward to check that $U_\tau = \text{int}(W_\tau)$, and W_τ is a (full-dimensional) set that deformation retracts to V_τ . Thus, $Z_\tau = U_\tau \cup V_\tau$ deformation retracts to V_τ . Finally, we obtain a weak deformation retract of Z_τ to U_τ via the following homotopy: we simultaneously translate each facet-defining hyperplane of the simplex $|\Delta^n|$ a small distance away from the simplex, so that the subset of Z_τ that is “swept up” is pushed into U_τ . \square

It is natural to ask whether a closed-set version of Proposition 3.10 holds:

Question 3.12. Is every locally good code a closed-good-cover code?

We could try to resolve this question by “closing” our open-good-cover realizations; we ask, For a locally good code \mathcal{C} , is $\{\overline{U}_i\}$ a closed-good-cover realization of \mathcal{C} , where $\{U_i\}$ is as in Proposition 3.10? However, this is not true. Indeed, it is easy to check that for the locally good code $\mathcal{C} = \{1, 12, 13\}$, the open realization $\{U_1, U_2, U_3\}$, when we take closures, yields a code $\mathcal{C}(\{\overline{U}_1, \overline{U}_2, \overline{U}_3\})$ with the “extra” codeword 123.

Theorem 3.13. *A code is locally good if and only if it is a good-cover code.*

Proof. The backward direction is in Proposition 2.15. The forward direction follows from Proposition 3.10 and the fact that if $\mathcal{C} \cup \{\emptyset\}$ is a good-cover code, then so is \mathcal{C} (Remark 2.5). \square

Finally, note that the good-cover realizations from Proposition 3.10 are embedded in an $(n-1)$ -dimensional Euclidean space, where n is the number of neurons. The following result shows that this embedding dimension can not, in general, be improved.

Proposition 3.14. *The code on n neurons comprising all proper subsets of $[n]$,*

$$\mathcal{C}_n := 2^{[n]} \setminus \{123 \dots n\},$$

is locally good, and can not be realized by a good cover in \mathbb{R}^{n-2} .

Proof. The code \mathcal{C}_n is locally good, as it is a simplicial complex and hence contains all its mandatory codewords. Next, suppose for contradiction that $\mathcal{U} = \{U_1, U_2, \dots, U_n\}$ forms a good-cover realization of \mathcal{C}_n in \mathbb{R}^{n-2} . By the nerve lemma (Proposition 2.13), the union $\cup_{i=1}^n U_i$ is homotopy equivalent to the nerve of \mathcal{U} , which is the boundary of the $(n-1)$ -simplex. However, no subset of \mathbb{R}^{n-2} is homotopy equivalent to an $(n-2)$ -sphere. We have reached a contradiction. \square

4. UNDECIDABILITY OF THE GOOD-COVER DECISION PROBLEM

Having shown that being locally good is equivalent to being a good-cover code (Theorem 3.13), we now prove that the corresponding decision problem is undecidable (Theorem 4.3). The proof hinges on the undecidability of determining whether a homology ball is contractible (Lemma 4.2) [25].

We say that a code \mathcal{C} is k -sparse if its simplicial complex $\Delta(\mathcal{C})$ has dimension at most $k-1$:

Definition 4.1. A code \mathcal{C} is k -sparse if $|\sigma| \leq k$ for all $\sigma \in \mathcal{C}$.

Lemma 4.2 (Tancer [25]). *The problem of deciding whether a given 4-dimensional simplicial complex is contractible is undecidable.*

Theorem 4.3 (Undecidability of the good-cover decision problem). *The problem of deciding whether a 5-sparse code is locally good (or, equivalently, 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 cone is itself a simplicial complex, which we denote by Δ' . Let \mathcal{C} denote the neural code $\Delta' \setminus \{v\}$, which is 5-sparse. Note that $\Delta(\mathcal{C}) = \Delta'$, so the only codeword in $\Delta(\mathcal{C})$ that is missing from \mathcal{C} is v . So, by Proposition 2.18, the code \mathcal{C} is locally good if and only if $\text{Lk}_v(\Delta') = \Delta$ 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 4.2. \square

4.1. Decidability for 3-sparse and 4-sparse codes. Can the condition of 5-sparsity in Theorem 4.3 can be extended to 4-sparsity or even 3-sparsity? For 3-sparsity, the answer is “no”:

Proposition 4.4. *The problem of determining whether a 3-sparse code is locally good (or, equivalently, has a good cover) is decidable.*

Proof. It is straightforward to write an algorithm that, for a given code \mathcal{C} , enumerates the nonempty intersections of maximal codewords. By Proposition 2.18, \mathcal{C} is locally good if and only if the links of these intersections are all contractible. Hence, we are interested in the decision problem for determining whether these links are contractible. When \mathcal{C} is 3-sparse, these links are 2-sparse, i.e., (undirected) graphs. Contractible graphs are precisely trees (connected graphs without cycles), and the problem of determining whether a graph is a tree is decidable. \square

As for extending Theorem 4.3 to 4-sparsity, this problem is open. This is because it is unknown whether the dimension in Lemma 4.2 can be lowered from 4 to 3 [25, Appendix A].

4.2. Relation to the convexity decision problem. For 2-sparse codes, being convex is equivalent to being locally good [15]. So, by Proposition 4.4, the convexity decision problem for these codes is decidable.

For codes without restriction on the sparsity, we revisit, from the proof of Theorem 4.3, codes of the form $\mathcal{C} = \Delta' \setminus \{v\}$, where Δ' is a cone over a simplicial complex Δ on a new vertex v . Consider the case when Δ is contractible. Does it follow that \mathcal{C} is convex? If it were, then by an argument analogous to the proof of Theorem 4.3, the convexity decision problem would be undecidable. However, we will see in the next section that there exist such codes \mathcal{C} that are non-convex (Example 5.14). Indeed, the convexity decision problem is unresolved, so we pose it here.

Question 4.5. Is the problem of determining whether a code is convex, decidable?

In the next section, we will introduce a superset of all convex codes, the “locally great” codes, and show that the corresponding decision problem is decidable (Theorem 5.18).

5. A NEW, STRONGER LOCAL OBSTRUCTION TO CONVEXITY

Recall that a code \mathcal{C} has no local obstructions if and only if it contains all its mandatory codewords (Proposition 2.18), and these mandatory codewords are precisely the faces of the simplicial complex $\Delta(\mathcal{C})$ whose link is non-contractible. We prove in this section that by replacing “non-contractible” by “non-collapsible” (Definition 5.1) we obtain a stronger type of local obstruction to convexity (Theorem 5.10).

This result yields a new family of codes that, like the earlier counterexample code (Proposition 2.16), are locally good, but not convex. This new family comprises codes of the form $\mathcal{C} = \tilde{\Delta} \setminus \{v\}$, where $\tilde{\Delta}$ is a cone, on a new vertex v , over a contractible but non-collapsible simplicial complex Δ (see Example 5.14). Such a code \mathcal{C} (as we saw in the proof of Theorem 4.3) is missing only one codeword (namely, v) from its simplicial complex (because $\Delta(\mathcal{C}) = \tilde{\Delta}$). We will use this several times in this section.

5.1. Background on collapses. First we make note of a somewhat overloaded definition in the literature. The notion of a *collapse* of a simplicial complex was introduced by Whitehead in 1938 [30]. The more general concept of *d-collapse* was introduced by Wegner in 1975 [28].

Definition 5.1. Let Δ be a simplicial complex, and let \mathcal{M} be the set of its facets.

- (1) For any face σ of Δ such that there is a unique $\tau \in \mathcal{M}$ for which $\sigma \subseteq \tau$, we define

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

and say that Δ' is an **elementary d-collapse** of Δ induced by σ (or (σ, τ)). This elementary *d-collapse* is denoted by $\Delta \rightarrow \Delta'$. (Here *d* refers to the constraint $\dim \sigma \leq d - 1$, but in this work we will let *d* be arbitrarily high when we use the term “*d-collapse*”). A sequence of elementary *d-collapses*


$$\Delta \rightarrow \Delta_1 \rightarrow \Delta_2 \rightarrow \cdots \rightarrow \Delta_n$$

is a **d-collapse** of Δ to Δ_n .


- (2) An **elementary collapse** is an elementary *d-collapse* induced by a face σ that is *not* a facet (i.e., $\sigma \subsetneq \tau$). A sequence of elementary collapses starting with Δ and ending with Δ' is a **collapse** of Δ to Δ' . Finally, a simplicial complex is **collapsible** if it collapses to a point (via some sequence).

Remark 5.2. Following [25], our Definition 5.1(2) characterizes “collapsible” in terms of elementary collapses induced by pairs (σ, τ) with $\dim \sigma < \dim \tau$. An equivalent definition, also used in the literature [1,2], is via elementary collapses under a stronger condition: $\dim \sigma = \dim \tau - 1$. For completeness, we prove in Appendix B that these two definitions of collapsible are equivalent (Proposition B.2).

Example 5.3 (A *d-collapse*). The following is an elementary *d-collapse* defined by $\sigma = \tau = 123$:



Example 5.4 (A collapse). The following is a collapse to a point:



This collapse arises from the elementary collapses induced by, respectively, $(\sigma, \tau) = (1, 123)$, $(\sigma, \tau) = (2, 23)$, and $(\sigma, \tau) = (3, 34)$.

Comparing Examples 5.3 and 5.4, note that the homotopy type was *not* preserved throughout the *d-collapse*, but was preserved through the collapse. This is explained in the next two results, the first of which is well known (see, e.g., [2, Ch. 1.2] or [1]).

Lemma 5.5. *Elementary collapses preserve homotopy type. Thus, collapsible implies contractible.*

However, not every contractible simplicial complex is collapsible. One example is any triangulation of the 2-dimensional topological space known as Bing’s house with two rooms [22, Ex. 1.6.1]. Another example is the dunce hat [32].

Lemma 5.5 extends as follows:

Lemma 5.6. *Let Δ be a contractible simplicial complex, and let $\Delta \rightarrow \Delta'$ be an elementary d -collapse induced by a face σ . Assume that Δ' contains a nonempty face. Then Δ' is contractible if and only if σ is not a facet (i.e., $\Delta \rightarrow \Delta'$ is an elementary collapse).*

Proof. The backward direction is immediate from Lemma 5.5. We prove the contrapositive of the forward direction. Suppose that σ is a facet of Δ . Then $|\Delta| = [\sigma] \cup |\Delta'|$, where $[\sigma]$ denotes the topological realization of σ as a facet of $|\Delta|$. With an eye toward using the Mayer-Vietoris sequence, we note that the intersection $[\sigma] \cap |\Delta'|$ is the boundary of the simplex $[\sigma]$ and thus is a sphere and so has a non-vanishing homology group. Applying the Mayer-Vietoris sequence for homology to $|\Delta| = [\sigma] \cup |\Delta'|$, and using the hypothesis that Δ is contractible, we conclude that $|\Delta'|$ has a non-vanishing homology group, and therefore is not contractible. \square

Finally, we state the following result of Wegner [28]; see also Tancer’s description in [24, §2].

Lemma 5.7 (Wegner [28]). *Let $\mathcal{W} = \{W_1, W_2, \dots, W_n\}$ be a collection of nonempty convex (not necessarily open) sets in \mathbb{R}^m . Let Λ be the nerve of \mathcal{W} . Then there exists an open halfspace H in \mathbb{R}^m such that, letting Λ' denote the nerve of the collection $\{W_1 \cap H, W_2 \cap H, \dots, W_n \cap H\}$, the following is an elementary d -collapse: $\Lambda \rightarrow \Lambda'$.*

Informally speaking, Wegner proved Lemma 5.7 by sweeping a hyperplane from infinity across \mathbb{R}^m , deleting everything in its path, until an intersection region corresponding to a facet τ of Λ has been removed. This yields the elementary d -collapse $\Lambda \rightarrow \Lambda'$ induced by some face contained in τ .

Remark 5.8. Recent work of Itskov, Kunin, and Rosen is similar in spirit to ours: they show that the collapsibility of a certain simplicial complex associated to a code (the “polar complex”) ensures that the code avoids certain obstructions to being a certain type of convex code (namely, a code arising from a “non-degenerate” hyperplane arrangement) [12, §6.5].

5.2. A key lemma. This subsection contains the key result (Lemma 5.9) that allows us to establish, via Theorem 5.10, our new local obstruction. Lemma 5.9 states that for an open cover by convex sets of a set that is itself convex, the corresponding nerve is collapsible. The original local obstruction (Proposition 2.15) relied on a weaker version of this result, which states only that such a nerve is contractible. Accordingly, the way we use Lemma 5.9 to prove Theorem 5.10 is analogous to how the authors of [4, 8] used the “contractible” version of the lemma to establish the original notion of local obstruction.

Lemma 5.9. *Let $\mathcal{W} = \{W_1, W_2, \dots, W_n\}$ be a collection of convex open sets in \mathbb{R}^m such that their union $W_1 \cup W_2 \cup \dots \cup W_n$ is nonempty and convex. Then the nerve of \mathcal{W} is collapsible.*

Proof. Let Λ denote the nerve of \mathcal{W} . Let p denote the number of nonempty faces of Λ . Note that $p \geq 1$, as the union of the W_i ’s is nonempty. We proceed by induction on p .

Base case: $p = 1$. Then Λ is a point, and thus is collapsible.

Inductive step: $p \geq 2$. Assume that the lemma is true for all nerves Λ with at most $p - 1$ nonempty faces. First consider the case when Λ has only one facet. Then Λ is a simplex, and every simplex is collapsible.

Now consider the remaining case, when Λ has at least two facets. Without loss of generality, each W_i is nonempty (deleting W_i ’s that are empty does not affect the union or the nerve). So, by Lemma 5.7, there exists an open halfspace H such that $\Lambda \rightarrow \Lambda'$ is an elementary d -collapse, where Λ' is the nerve of $\mathcal{V} := \{W_1 \cap H, W_2 \cap H, \dots, W_n \cap H\}$. We see that \mathcal{V} is a collection of convex open sets whose union,

$(\cup_{i \in [n]} W_i) \cap H$, is convex (because both $\cup_{i \in [n]} W_i$ and H are convex) and nonempty (indeed, the nerve Λ' is nonempty, because Λ has at least two facets and so at least one was unaffected by the elementary d -collapse $\Lambda \rightarrow \Lambda'$). Thus, by the induction hypothesis, Λ' is collapsible.

Thus, by definition of collapsible, we need only show that the elementary d -collapse $\Lambda \rightarrow \Lambda'$ was in fact an elementary collapse. To see this, note that the nerve theorem (Proposition 2.13) implies that the nerve Λ' is homotopy equivalent to $(\cup_{i \in [n]} W_i) \cap H$, which we saw above is convex (and nonempty) and thus contractible. Hence, by Lemma 5.6, $\Lambda \rightarrow \Lambda'$ is an elementary collapse. This completes the proof. \square

5.3. Locally great codes. The following result gives a new class of local obstructions.

Theorem 5.10. *Let \mathcal{C} be a convex code. Then for any $\sigma \in \Delta(\mathcal{C}) \setminus \mathcal{C}$, the link $\text{Lk}_\sigma(\Delta(\mathcal{C}))$ is collapsible.*

Proof. Let U_1, U_2, \dots, U_n be convex open sets (in some stimulus space $X \subseteq \mathbb{R}^m$) that realize \mathcal{C} . Let $\sigma \in \Delta(\mathcal{C}) \setminus \mathcal{C}$. Then, by definition, $\emptyset \neq U_\sigma \subseteq \cup_{i \notin \sigma} U_i$. Thus, $\mathcal{W} := \{U_i \cap U_\sigma\}_{i \notin \sigma}$ is a collection of convex open sets such that their union equals U_σ (and thus this union is nonempty and convex). So, by Lemma 5.9, the nerve of \mathcal{W} is collapsible. It is straightforward to check that this nerve equals $\text{Lk}_\sigma(\Delta(\mathcal{C}))$. So, as desired, the link is collapsible. \square

Definition 5.11. A code \mathcal{C} has a **local obstruction of the second kind** if there exists $\sigma \in \Delta(\mathcal{C}) \setminus \mathcal{C}$ such that the link $\text{Lk}_\sigma(\Delta(\mathcal{C}))$ is *not* collapsible. If \mathcal{C} has *no* local obstructions of the second kind, \mathcal{C} is **locally great**.

The next result follows from Theorem 5.10 and the fact that collapsible implies contractible (Lemma 5.5).

Corollary 5.12. \mathcal{C} is convex $\Rightarrow \mathcal{C}$ is locally great $\Rightarrow \mathcal{C}$ is locally good.

Neither implication in Corollary 5.12 is an equivalence, as we see in the following examples.

Example 5.13 (Locally great does *not* imply convex). The counterexample code from Proposition 2.16 is non-convex [17], but, we claim, locally great. To verify this, we must check that for every missing codeword $\sigma \in \Delta(\mathcal{C}) \setminus \mathcal{C} = \{234, 235, 245, 345, 12, 15, 24, 25, 35, 1, 2, 5\}$, the corresponding link $\text{Lk}_\sigma(\Delta(\mathcal{C}))$ is collapsible. We accomplish this as follows:

- When $\sigma \in \{234, 235, 245, 345, 12, 15\}$, the link is a point and thus collapsible.
- When $\sigma \in \{24, 25, 35, 1\}$, the link is a single edge and thus collapsible (cf. Example 5.4).
- When $\sigma \in \{2, 5\}$, the link is, up to relabeling, the simplicial complex in Example 5.4, which we showed is collapsible.

Example 5.14 (Locally good does *not* imply locally great). Let Δ be a triangulation of Bing's house, so Δ is a 2-dimensional simplicial complex that is contractible but not collapsible [22, Ex. 1.6.1]. Let $\tilde{\Delta}$ be a cone over Δ on a new vertex v , and consider the code $\mathcal{C} := \tilde{\Delta} \setminus \{v\}$.

We claim that \mathcal{C} is locally good, but not locally great. Indeed, the only “missing” codeword of \mathcal{C} is v , and its link is $\text{Lk}_v(\Delta(\mathcal{C})) = \Delta$, which is contractible but not collapsible. Thus, \mathcal{C} is locally good (by Proposition 2.18), but not locally great (by definition).

We now consider the question of whether there might be an even stronger local obstruction beyond “locally great” (perhaps “locally excellent”?). That is, can the conclusion of Lemma 5.9 be strengthened from “collapsible” to some more general property of simplicial complexes? Or, on the contrary, does the converse of Lemma 5.9 hold? Accordingly, we pose the following question:

Question 5.15. Is every collapsible simplicial complex the nerve of a convex open cover of some convex set?

After our manuscript appeared on arXiv, Question 5.15 was resolved in the negative by Jeffs and Novik [14], who also proved further results about neural codes that can be realized by convex open sets whose union is convex.

Remark 5.16. Some decision problems related to Question 5.15 have been resolved. The problem of deciding whether a simplicial complex is the nerve of a convex open cover of some (not necessarily convex) subset of \mathbb{R}^d is decidable [24, 27]. On the other hand, the problem of deciding whether a simplicial complex is the nerve of a good cover of some subset of \mathbb{R}^d , when $d \geq 5$, is undecidable [26].

5.4. The locally-great decision problem. In this subsection, we prove that the locally-great decision problem is decidable. This result relies on the following lemma.

Lemma 5.17 (Tancer [25]). *The problem of deciding whether a given simplicial complex is collapsible is NP-complete.*

Theorem 5.18 (Decidability of the locally-great decision problem). *The problem of deciding whether a neural code is locally great is NP-hard.*

Proof. By Definition 5.11 and Lemma 5.17, the following steps form an algorithm that determines whether a code is locally great: enumerate $\Delta(\mathcal{C}) \setminus \mathcal{C}$ and the corresponding links, and then check whether any of these links is collapsible.

To show this is NP-hard, it suffices to reduce (in polynomial time) the problem of deciding whether a simplicial complex is collapsible (which is NP-complete by Lemma 5.17) to the locally-great decision problem. To this end, we proceed as in the proof of Theorem 4.3: given any simplicial complex Δ , consider the cone over Δ on a new vertex v . This cone is itself a simplicial complex, which we denote by Δ' . Let $\mathcal{C} = \Delta' \setminus \{v\}$. Then, v is the only codeword in $\Delta(\mathcal{C})$ that is missing from \mathcal{C} . So, by Definition 5.11, the original simplicial complex $\Delta = \text{Lk}_v(\Delta')$ is collapsible if and only if the code \mathcal{C} is locally great. \square

6. DISCUSSION

We return to the question that opened this work, Which neural codes are convex? There is a growing literature tackling this question, and here we resolved some foundational problems in this developing theory. In summary, we now know the following:

\mathcal{C} is convex $\xrightarrow{(a)}$ \mathcal{C} is locally great $\xrightarrow{(b)}$ \mathcal{C} is a good-cover code \Leftrightarrow \mathcal{C} is locally good $\xrightarrow{(c)}$ \mathcal{C} is connected, and the implications (a)–(c) are not equivalences. Also, it is undecidable to tell whether an arbitrary neural code \mathcal{C} is locally good, NP-hard to tell whether \mathcal{C} is locally great, and the problem remains open for determining whether \mathcal{C} is convex (Question 4.5).

An additional problem suggested by our work pertains to detecting our new local obstructions (those of the second kind). Neural codes have been studied from an algebraic standpoint via neural ideals, which are closely related to the Stanley-Reisner ideal of a code's simplicial complex [4, 5, 7, 11, 13, 16, 18]. Using these algebraic tools, it is possible to find local obstructions that can be detected by homology, which suffices to determine contractibility for small simplicial complexes. It would be interesting to see whether collapsibility of links can be characterized in a similar way, as unlike contractibility the former is decidable.

Next, we revisit the implication, $\text{convex} \Rightarrow \text{locally good}$. The converse is false in general, with the first counterexample a 4-sparse code [17], but is true for 2-sparse codes [15]. We ask, therefore, *Is every 3-sparse locally good code, convex?* Indeed, for \mathcal{C} a 3-sparse code, mandatory codewords have size one or two, so the corresponding links are graphs, where contractibility and collapsibility are equivalent (and thus so are the concepts of locally great and locally good). This problem therefore seems tractable.

Finally, it would be interesting to investigate how the theory of convex codes changes when considering closed rather than open realizations. Such a theory was initiated by Cruz *et al.*, who proved that there are open-convex codes that are not closed-convex and vice-versa [3]. We also know that every closed-convex code is a closed-good-cover code (by definition), while the converse is false [3]. Every closed-convex code also is locally great (Lemma 5.9 and Theorem 5.10 easily generalize to closed sets, as a version of the nerve theorem applies to closed sets in \mathbb{R}^d). Also, every code that is a closed-good-cover code is locally good (by the nerve theorem – see [3]). A related open question posed earlier is whether every locally good code is a closed-good-cover code (Question 3.12). If so, then every locally good code also would be a closed-good-cover code, unifying some of the theory of open-convex and closed-convex codes.

We end by revisiting the neuroscience motivation for our work. Recall that place cells enable the brain to represent an organism’s environment by way of neural codes arising from approximately convex receptive fields. The goal of analyzing which codes are convex, therefore, is aimed at understanding what types of codes allow the brain to represent structured environments. One of the contributions here is to show that while the convex-code decision problem is still open, some related decision problems are hard. Analyzing data from many neurons, therefore, may be computationally challenging. On the other hand, our work reveals the meaning behind the main tool used for precluding convexity, namely, local obstructions: these obstructions exactly characterize when a code can *not* be realized by an (open) good cover.

Acknowledgements. This work was initiated by AC at the 2016 REU in the Department of Mathematics at Texas A&M University, which was supported by the NSF (DMS-1460766). The authors thank Chad Giusti, Laura Matusевич, Kaitlyn Phillipson, Zvi Rosen, Ola Sobieska, Martin Tancer, and Zev Woodstock for insightful comments and discussion. The authors also thank two conscientious referees whose comments improved this work. AS was partially supported by the NSF (DMS-1312473/DMS-1513364) and the Simons Foundation (#521874).

REFERENCES

- [1] A. Björner. Topological methods. In R. L. Graham, M. Grötschel, and L. Lovász, editors, *Handbook of Combinatorics (Vol. 2)*, chapter 34, pages 1819–1872. MIT Press, Amsterdam, 1995.
- [2] Marshall M. Cohen. *A course in simple-homotopy theory*, volume 10. Springer Science & Business Media, 2012.
- [3] Joshua Cruz, Chad Giusti, Vladimir Itskov, and Bill Kronholm. On open and closed convex codes. Available at [arXiv:1609.03502](https://arxiv.org/abs/1609.03502), 2016.
- [4] Carina Curto, Elizabeth Gross, Jack Jeffries, Katie Morrison, Mohamed Omar, Zvi Rosen, Anne Shiu, and Nora Youngs. What makes a neural code convex? *SIAM Journal on Applied Algebra and Geometry*, 1(1):222–238, 2017.
- [5] 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.
- [6] Megan Franke and Samuel Muthiah. Every binary code can be realized by convex sets. *Adv. Appl. Math.*, 99:83–93, 2018.
- [7] Rebecca Garcia, Luis David García Puente, Ryan Kruse, Jessica Liu, Dane Miyata, Ethan Petersen, Kaitlyn Phillipson, and Anne Shiu. Gröbner bases of neural ideals. *Int. J. Algebr. Comput.*, 28(4):553–571, 2018.
- [8] Chad Giusti and Vladimir Itskov. A no-go theorem for one-layer feedforward networks. *Neural Comput.*, 26(11):2527–2540, 2014.

- [9] Elizabeth Gross, Nida Kazi Obatake, and Nora Youngs. Neural ideals and stimulus space visualization. *Adv. Appl. Math.*, 95:65–95, 2018.
- [10] Allen Hatcher. *Algebraic topology*. Cambridge University Press, 2002.
- [11] Molly Hoch, Samuel Muthiah, and Nida Obatake. On the identification of k -inductively pierced codes using toric ideals. Available at [arXiv:1807.02390](https://arxiv.org/abs/1807.02390).
- [12] Vladimir Itskov, Alex Kunin, and Zvi Rosen. Hyperplane neural codes and the polar complex. Available at [arXiv:1801.02304](https://arxiv.org/abs/1801.02304), 2018.
- [13] Jack Jeffries, Sema Güntürkün, and Jeffrey Sun. Polarization of neural ideals. Available at [arXiv:1706.08559](https://arxiv.org/abs/1706.08559), 2017.
- [14] R. Amzi Jeffs and Isabella Novik. Convex union representability and convex codes. Available at [arXiv:1808.03992](https://arxiv.org/abs/1808.03992), 2018.
- [15] R. Amzi Jeffs, Mohamed Omar, Natchanon Suaysom, Aleina Wachtel, and Nora Youngs. Sparse neural codes and convexity. Available at [arXiv:1511.00283](https://arxiv.org/abs/1511.00283), 2015.
- [16] R. Amzi Jeffs, Mohamed Omar, and Nora Youngs. Homomorphisms preserving neural ideals. *J. Pure Appl. Algebra*, 222:3470–3482, 2018.
- [17] Caitlin Lienkaemper, Anne Shiu, and Zev Woodstock. Obstructions to convexity in neural codes. *Adv. Appl. Math.*, 85:31–59, 2017.
- [18] Angélique Morvant. Strengthening relationships between neural ideals and receptive fields. Available at [arXiv:1803.03204](https://arxiv.org/abs/1803.03204), 2018.
- [19] Raffaella Mulas and Ngoc M. Tran. Characterization and minimal embeddings of connected neural codes. Available at [arXiv:1706.03999](https://arxiv.org/abs/1706.03999), 2017.
- [20] John O’Keefe and Jonathan Dostrovsky. The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res.*, 34(1):171–175, 1971.
- [21] Zvi Rosen and Yan X Zhang. Convex neural codes in dimension 1. Available at [arXiv:1702.06907](https://arxiv.org/abs/1702.06907), 2017.
- [22] T. Benny Rushing. *Topological embeddings*, volume 52. Academic Press, 1973.
- [23] A. H. Stone. Paracompactness and product spaces. *Bull. Amer. Math. Soc.*, 54:977–982, 1948.
- [24] 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.
- [25] Martin Tancer. Recognition of collapsible complexes is NP-complete. *Discrete Comput. Geom.*, 55(1):21–38, 2016.
- [26] Martin Tancer and Dmitry Tonkonog. Nerves of good covers are algorithmically unrecognizable. *SIAM J. Comput.*, 42(4):1697–1719, 2013.
- [27] G. Wegner. *Eigenschaften der Nerven homologisch-einfacher Familien im R^n* . PhD thesis, Universität Göttingen, 1967. In German.
- [28] Gerd Wegner. d -collapsing and nerves of families of convex sets. *Arch. Math.*, 26(1):317–321, 1975.
- [29] André Weil. Sur les théorèmes de de Rham. *Comment. Math. Helv.*, 26:119–145, 1952.
- [30] J.H.C. Whitehead. Simplicial spaces, nuclei and m -groups. *P. Lon. Math. Soc.*, 45:243–327, 1938.
- [31] Robert Williams. Strongly maximal intersection-complete neural codes on grids are convex. *Appl. Math. Comput.*, 336:162–175, 2018.
- [32] EC Zeeman. On the dunce hat. *Topology*, 2(4):341–358, 1963.

APPENDIX A. REMOVING INTERIORS OF FACES

Here we prove Lemma 3.7, which states that for a contractible simplicial complex Σ , removing the relative interiors of some faces with contractible link, yields a space that is still contractible.

We will use the following notation. For a face τ of a simplicial complex Σ , the *open star* of τ is:

$$\text{St}_\tau(\Sigma) := \bigcup_{\tau \subseteq \sigma \in \Sigma} \text{int}(\sigma) \subseteq |\Sigma|.$$

(This union is in fact a disjoint union.) If v is a vertex of Σ , we denote $\text{St}_{\{v\}}(\Sigma)$ also by $\text{St}_v(\Sigma)$. We denote the barycentric subdivision of Σ by Σ' , that is, the faces of Σ' are flags $\tau_1 \subsetneq \tau_2 \subsetneq \cdots \subsetneq \tau_k$ of nonempty

faces in Σ . Equivalently, Σ' is the order complex of the face poset of Σ . Finally, if τ is a face of Σ , we let $v(\tau)$ denote the corresponding vertex in the barycentric subdivision Σ' .

Lemma A.1. *Let Σ be a simplicial complex, and let $\tau_1, \tau_2, \dots, \tau_k$ be nonempty faces of Σ . Then the intersection $\bigcap_i \text{St}_{v(\tau_i)}(\Sigma')$ is nonempty if and only if the faces τ_i are totally ordered by inclusion.*

Proof. If $\bigcap_i \text{St}_{v(\tau_i)}(\Sigma') \neq \emptyset$, let $x \in \bigcap_i \text{St}_{v(\tau_i)}(\Sigma')$. There is a unique face $\sigma' \in \Sigma'$ with $x \in \text{int}(\sigma')$. Then $\text{int}(\sigma') \subseteq \text{St}_{v(\tau_i)}(\Sigma')$ for all $i \in \{1, 2, \dots, k\}$. Thus, $v(\tau_1), v(\tau_2), \dots, v(\tau_k)$ are vertices of σ' , and so form a face in Σ' . This means that the τ_i 's, after reordering, form a flag in Σ , i.e., are totally ordered by inclusion.

Conversely, if, after reordering, $\tau_1 \subseteq \tau_2 \subseteq \dots \subseteq \tau_k$, then this flag corresponds to a face σ' of Σ' with vertex set $v(\tau_1), v(\tau_2), \dots, v(\tau_k)$. Now $\text{int}(\sigma') \subseteq \text{St}_{v(\tau_i)}(\Sigma')$ for all $i \in \{1, 2, \dots, k\}$, so $\bigcap_i \text{St}_{v(\tau_i)}(\Sigma') \neq \emptyset$. \square

Lemma A.2. *Let Γ be a collection of faces of a simplicial complex Σ . Let $K = \bigcup_{\sigma \in \Gamma} \text{int}(\sigma) \subseteq |\Sigma|$. Then for any $\tau \in \Gamma$, the intersection $K \cap \text{St}_{v(\tau)}(\Sigma')$ is open in K and contractible. Moreover, the collection*

$$\mathcal{U} = \{K \cap \text{St}_{v(\tau)}(\Sigma') \mid \tau \in \Gamma\}$$

is a good cover of K , and the nerve $\mathcal{N}(\mathcal{U})$ is homotopy equivalent to K .

Proof. The open star $\text{St}_{v(\tau)}(\Sigma')$ is open in $|\Sigma|$ and thus $K \cap \text{St}_{v(\tau)}(\Sigma')$ is open in K . There is a deformation retraction of $\text{St}_{v(\tau)}(\Sigma')$ to $v(\tau)$ along straight lines. Indeed, if σ is a face of Σ with $\tau \subseteq \sigma$, then for any $x \in \text{int}(\sigma)$ the segment from $v(\tau)$ to x , excluding $v(\tau)$, is entirely within $\text{int}(\sigma)$. Thus the deformation retraction induces a deformation retraction of $K \cap \text{St}_{v(\tau)}(\Sigma')$ to $v(\tau)$. So, $K \cap \text{St}_{v(\tau)}(\Sigma')$ is contractible.

To show that \mathcal{U} is a cover of K , let $x \in K$. Then $x \in \text{int}(\sigma)$ for some $\sigma \in \Gamma$, and so $x \in \text{int}(\sigma') \subseteq \text{St}_{v(\sigma)}(\Sigma')$ for some face σ' of Σ' for which $v(\sigma)$ is a vertex. Hence, x is in the set $K \cap \text{St}_{v(\sigma)}(\Sigma')$ from \mathcal{U} .

If $\tau_1, \tau_2, \dots, \tau_k$ are faces in Γ such that $\bigcap_i (K \cap \text{St}_{v(\tau_i)}(\Sigma')) \neq \emptyset$, then the faces τ_i are totally ordered by inclusion by Lemma A.1 and thus correspond to a face τ' of Σ' . It is straightforward to check that the intersection $\bigcap_i (K \cap \text{St}_{v(\tau_i)}(\Sigma'))$ is equal to $K \cap \text{St}_{v(\tau')}(\Sigma')$, which is contractible, since there is a straight-line deformation retraction to the barycenter of τ' , which does not leave K . Thus, \mathcal{U} is a good cover, and so the nerve theorem (Proposition 2.13) implies that the nerve $\mathcal{N}(\mathcal{U})$ is homotopy equivalent to K . \square

Proposition A.3 (Lemma 3.7, restated). *Let Σ be a contractible simplicial complex, and let $\sigma_1, \sigma_2, \dots, \sigma_k$ be faces of Σ such that, for $i \in \{1, 2, \dots, k\}$, the link $\text{Lk}_{\sigma_i}(\Sigma)$ is contractible. Then*

$$|\Sigma| \setminus (\text{int}(\sigma_1) \cup \text{int}(\sigma_2) \cup \dots \cup \text{int}(\sigma_k))$$

is contractible.

Proof. Let Γ be the collection of all faces of Σ with the exception of $\sigma_1, \sigma_2, \dots, \sigma_k$. Let $K = \bigcup_{\sigma \in \Gamma} \text{int}(\sigma) = |\Sigma| \setminus (\text{int}(\sigma_1) \cup \text{int}(\sigma_2) \cup \dots \cup \text{int}(\sigma_k))$, and let $\mathcal{U} = \{K \cap \text{St}_{v(\tau)}(\Sigma') \mid \tau \in \Gamma\}$. By Lemma A.2, the collection \mathcal{U} is a good cover of K , and the nerve $\mathcal{N}(\mathcal{U})$ is homotopy equivalent to K .

Pick an inclusion-maximal face among $\sigma_1, \sigma_2, \dots, \sigma_k$, say, σ_1 is not contained in any other σ_i . Let $\hat{\Gamma} = \Gamma \cup \{\sigma_1\}$, and let $\hat{K} = \bigcup_{\sigma \in \hat{\Gamma}} \text{int}(\sigma)$. Let $\hat{\mathcal{U}} = \{K \cap \text{St}_{v(\tau)}(\Sigma') \mid \tau \in \hat{\Gamma}\}$. Then, again by Lemma A.2, $\hat{\mathcal{U}}$ is a good cover of \hat{K} , and the nerve $\mathcal{N}(\hat{\mathcal{U}})$ is homotopy equivalent to \hat{K} .

To recap, $\mathcal{N}(\mathcal{U}) \simeq K$ and $\mathcal{N}(\hat{\mathcal{U}}) \simeq \hat{K}$. Hence, the following claim would imply that $K \simeq \hat{K}$:

Claim: $\mathcal{N}(\mathcal{U}) \simeq \mathcal{N}(\hat{\mathcal{U}})$.

Moreover, repeating this argument an additional $k - 1$ times (applying it next to \widehat{K} , and so on) would imply that $K \simeq |\Sigma|$, and thus, as desired, K is contractible. Hence, we need only prove the Claim.

We first view the two nerves $\mathcal{N}(\mathcal{U})$ and $\mathcal{N}(\widehat{\mathcal{U}})$ as (abstract) simplicial complexes. Let $\tau_1, \tau_2, \dots, \tau_\ell \in \Gamma$. Then $\bigcap_i (K \cap \text{St}_{v(\tau_i)}(\Sigma')) \neq \emptyset$ if and only if $\bigcap_i (\widehat{K} \cap \text{St}_{v(\tau_i)}(\Sigma')) \neq \emptyset$. This is because the open vertex stars $\text{St}_{v(\tau_i)}(\Sigma')$, if they intersect, cannot only intersect in a subset of $\text{int}(\sigma_1)$, since the τ_i 's are totally ordered by inclusion by Lemma A.1, so the corresponding face τ' of Σ' would have $\text{int}(\tau')$ contained in the intersection $\bigcap_i \text{St}_{v(\tau_i)}(\Sigma')$, which would imply that one of the τ_i 's is equal to σ_1 , and this would contradict the fact that the τ_i 's are in Γ . Thus the simplicial complexes $\mathcal{N}(\mathcal{U})$ and $\mathcal{N}(\widehat{\mathcal{U}})$ differ only in the vertex w corresponding to the set $\widehat{K} \cap \text{St}_{v(\sigma_1)}(\Sigma')$ and the faces incident to w .

Let Δ_1 denote the join $\text{Lk}_w(\mathcal{N}(\widehat{\mathcal{U}})) * \{w\}$, and let $\Delta_2 = \mathcal{N}(\mathcal{U})$. Then Δ_1 is a cone over the vertex w and thus is contractible; also, $\Delta_1 \cap \Delta_2 = \text{Lk}_w(\mathcal{N}(\widehat{\mathcal{U}}))$. So, to complete the proof it suffices to show that $\text{Lk}_w(\mathcal{N}(\widehat{\mathcal{U}}))$ is contractible, as [1, Lemma 10.3] will imply that $\Delta_1 \cup \Delta_2 \simeq \Delta_2$, that is, $\mathcal{N}(\widehat{\mathcal{U}}) \simeq \mathcal{N}(\mathcal{U})$.

Let P be the poset of all faces in Γ that are properly contained in σ_1 , and denote its order complex by $\Delta(P)$. We will show that $\text{Lk}_w(\mathcal{N}(\widehat{\mathcal{U}}))$ is isomorphic to the join $\Delta(P) * (\text{Lk}_{\sigma_1}(\Sigma))'$. A face of $\text{Lk}_w(\mathcal{N}(\widehat{\mathcal{U}}))$ corresponds to faces $\tau_1, \tau_2, \dots, \tau_k \in \Gamma$ such that

$$\widehat{K} \cap \text{St}_{v(\sigma_1)}(\Sigma') \cap \bigcap_i \widehat{K} \cap \text{St}_{v(\tau_i)}(\Sigma') \neq \emptyset.$$

Up to reordering the τ_i 's, this implies, by Lemma A.1, that $\tau_1 \subseteq \dots \subseteq \tau_\ell \subseteq \sigma_1 \subseteq \tau_{\ell+1} \subseteq \dots \subseteq \tau_k$. The flag of faces $\tau_1 \subseteq \dots \subseteq \tau_\ell$ determines a face of $\Delta(P)$, while $\tau_{\ell+1} \subseteq \dots \subseteq \tau_k$ determines a face of $(\text{Lk}_{\sigma_1}(\Sigma))'$. Conversely, a face of $\Delta(P) * (\text{Lk}_{\sigma_1}(\Sigma))'$ corresponds to a flag of faces $\tau_1 \subseteq \dots \subseteq \tau_\ell \subseteq \sigma_1$ with $\tau_1, \dots, \tau_\ell \in \Gamma$ and a flag $\sigma_1 \subseteq \tau_{\ell+1} \subseteq \dots \subseteq \tau_k$ with $\tau_{\ell+1}, \dots, \tau_k \in \Sigma$. We chose σ_1 as an inclusion-maximal face of Σ not in Γ , so $\tau_{\ell+1}, \dots, \tau_k \in \Gamma$. Thus we have inclusions $\tau_1 \subseteq \dots \subseteq \tau_\ell \subseteq \sigma_1 \subseteq \tau_{\ell+1} \subseteq \dots \subseteq \tau_k$ with $\tau_1, \dots, \tau_k \in \Gamma$, which precisely corresponds to a face of $\text{Lk}_w(\mathcal{N}(\widehat{\mathcal{U}}))$ by Lemma A.1. Hence, $\text{Lk}_w(\mathcal{N}(\widehat{\mathcal{U}}))$ is isomorphic to $\Delta(P) * (\text{Lk}_{\sigma_1}(\Sigma))'$.

By hypothesis, $(\text{Lk}_{\sigma_1}(\Sigma))'$ is contractible, and thus the join $\Delta(P) * (\text{Lk}_{\sigma_1}(\Sigma))'$ is contractible. This join, we saw, is isomorphic to $\text{Lk}_w(\mathcal{N}(\widehat{\mathcal{U}}))$, so $\text{Lk}_w(\mathcal{N}(\widehat{\mathcal{U}}))$ is contractible. Hence, as explained above, [1, Lemma 10.3] implies that $\mathcal{N}(\mathcal{U}) \simeq \mathcal{N}(\widehat{\mathcal{U}})$, proving the Claim. \square

APPENDIX B. TWO DEFINITIONS OF COLLAPSIBLE

Here we prove that the definition of ‘‘collapsible’’ given earlier (Definition 5.1(2)) is equivalent to Definition B.1 below, which is also used in the literature [1, 2]. The proof was conveyed to us by Martin Tancer.

Recall that for a simplicial complex Δ , an *elementary collapse* induced by a pair (σ, τ) , where σ is a face contained in a unique facet τ , is the simplicial complex $\Delta' = \Delta \setminus \{\nu \in \Delta \mid \sigma \subseteq \nu\}$.

Definition B.1. A **collapse** of Δ to Δ' is a sequence of elementary collapses induced by pairs (σ, τ) with $\dim \sigma = \dim \tau - 1$ starting with Δ and ending with Δ' . A simplicial complex is **collapsible** if it collapses to a point (via some sequence).

Proposition B.2. *The definitions of ‘‘collapsible’’, in Definitions 5.1(2) and B.1, respectively, are equivalent.*

Proof. Let Δ be a simplicial complex. We need only show that every elementary collapse $\Delta \rightarrow \Delta'$ as in Definition 5.1(2), i.e., induced by some (σ, τ) with $\sigma \subsetneq \tau$ and τ a facet, can be obtained by a sequence of

elementary collapses as in Definition B.1. We proceed by induction on $(\dim \tau - \dim \sigma)$. The base case, when $\dim \tau - \dim \sigma = 1$, is already an elementary collapse under either definition.

For the inductive step, assume that $\dim \tau - \dim \sigma \geq 2$ and that every elementary collapse (of any simplicial complex) induced by pairs (σ', τ') with $(\dim \tau' - \dim \sigma') < (\dim \tau - \dim \sigma)$, can be obtained by a sequence of elementary collapses as in Definition B.1.

Let $z \in \tau \setminus \sigma$. Then $\sigma \cup \{z\} \subsetneq \tau$, and also τ is the unique facet containing $\sigma \cup \{z\}$. Therefore, by the inductive hypothesis, the collapse $\Delta \rightarrow \tilde{\Delta}$ induced by $(\sigma \cup \{z\}, \tau)$ can be obtained by a sequence of elementary collapses as in Definition B.1. Thus, we need only show that we can apply such elementary collapses to $\tilde{\Delta}$ to obtain Δ' (the outcome of the (σ, τ) -collapse).

Note that σ and $\tau \setminus \{z\}$ are both faces of $\tilde{\Delta}$ (only faces of Δ containing $\sigma \cup \{z\}$ were removed). We claim that the unique facet of $\tilde{\Delta}$ that contains σ is $\tau \setminus \{z\}$. Indeed, first, $\tau \setminus \{z\}$ contains σ (by our choice of σ, τ , and z), and, second, $\tau \setminus \{z\}$ is maximal in $\tilde{\Delta}$ (otherwise τ would *not* have been maximal in Δ).

Thus, by the inductive hypothesis, the collapse $\tilde{\Delta} \rightarrow \hat{\Delta}$ induced by $(\sigma, \tau \setminus \{z\})$ can be obtained by a sequence of elementary collapses as in Definition B.1. So, to complete the proof, we need only show that $\hat{\Delta} = \Delta'$. Indeed, $\tilde{\Delta}$ was obtained by removing faces of Δ that contain $\sigma \cup \{z\}$, and then $\hat{\Delta}$ was obtained by removing faces that contain σ ; hence, $\hat{\Delta} = \Delta \setminus \{\nu \in \Delta \mid \sigma \subseteq \nu\} = \Delta'$. \square

AARON CHEN, DEPARTMENT OF MATHEMATICS, CORNELL UNIVERSITY, ITHACA NY 14853

Current address: University of Chicago, Chicago IL 60637

E-mail address: aaronchen@uchicago.edu

FLORIAN FRICK, DEPARTMENT OF MATHEMATICS, CORNELL UNIVERSITY, ITHACA NY 14853

Current address: Department of Mathematical Sciences, Carnegie Mellon University, Pittsburgh PA 15213

E-mail address: frick@cmu.edu

ANNE SHIU, DEPARTMENT OF MATHEMATICS, TEXAS A&M UNIVERSITY, COLLEGE STATION TX 77843

E-mail address: annejls@math.tamu.edu