

A No-Go Theorem for One-Layer Feedforward Networks

Chad Giusti

cgiusti@seas.upenn.edu

Vladimir Itskov

vladimir.itskov@psu.edu

Department of Mathematics, University of Nebraska–Lincoln, Lincoln, NE 68588, U.S.A.

It is often hypothesized that a crucial role for recurrent connections in the brain is to constrain the set of possible response patterns, thereby shaping the neural code. This implies the existence of neural codes that cannot arise solely from feedforward processing. We set out to find such codes in the context of one-layer feedforward networks and identified a large class of combinatorial codes that indeed cannot be shaped by the feedforward architecture alone. However, these codes are difficult to distinguish from codes that share the same sets of maximal activity patterns in the presence of subtractive noise. When we coarsened the notion of combinatorial neural code to keep track of only maximal patterns, we found the surprising result that all such codes can in fact be realized by one-layer feedforward networks. This suggests that recurrent or many-layer feedforward architectures are not necessary for shaping the (coarse) combinatorial features of neural codes. In particular, it is not possible to infer a computational role for recurrent connections from the combinatorics of neural response patterns alone. Our proofs use mathematical tools from classical combinatorial topology, such as the nerve lemma and the existence of an inverse nerve. An unexpected corollary of our main result is that any prescribed (finite) homotopy type can be realized by a subset of the form $\mathbb{R}_{\geq 0}^n \setminus \mathcal{P}$, where \mathcal{P} is a polyhedron.

1 Introduction ---

It is often hypothesized that one of the central roles of recurrent connections in the brain is to constrain the set of possible neural response patterns, thereby shaping the neural code (Douglas & Martin, 2007; Luczak, Barthó, & Harris, 2009). This hypothesis is appealing because it provides a concrete computational function for the prevalence of recurrent connectivity in cortical areas. It also implies the existence of neural codes that cannot arise from the structure of feedforward connections alone.

We test this hypothesis by analyzing the neural codes of feedforward networks. Although it is well known that a feedforward network with hidden

layers can approximate any continuous function (Cybenko, 1989; Hornik, 1991), this is not the case for one-layer networks. It is thus reasonable to expect that there exist neural codes that cannot arise in one-layer feedforward networks and are necessarily shaped by a more complex network structure.

For simplicity, we consider combinatorial neural codes (Schneidman, Berry, Segev, & Bialek, 2006; Osborne, Palmer, Lisberger, & Bialek, 2008; Schneidman et al., 2011), which consist of binary neural activity patterns, and disregard details such as exact firing rates or relative spike timing between different neurons. Our first result identifies a large class of combinatorial codes that are not realized by one-layer feedforward networks. However, we also find that in the presence of noise, these codes are difficult to distinguish from other codes that share the same sets of maximal activity patterns. In order to increase robustness to noise, we coarsen the notion of combinatorial code to keep track of only maximal patterns and again seek to find codes that cannot arise in one-layer feedforward networks. To our surprise, there are none.

Our main result is a no-go theorem stating that all coarse combinatorial codes can in fact be realized by one-layer feedforward networks. Our proof of this theorem is constructive and uses mathematical tools from classical combinatorial topology, such as the nerve lemma and the existence of an inverse nerve. An unexpected corollary is that any prescribed (finite) homotopy type can be realized by a subset of the form $\mathbb{R}_{\geq 0}^n \setminus \mathcal{P}$, where \mathcal{P} is a polyhedron.

This no-go theorem implies that recurrent or many-layer feedforward architectures are not necessary for shaping combinatorial features of neural codes. In particular, it is not possible to infer a computational role for recurrent connections solely from the combinatorics of neural response patterns. However, we also show that one-layer feedforward networks that respect Dale's law (Dayan & Abbott, 2001) produce fairly trivial neural codes, possessing just one maximal activity pattern. Thus, one-layer feedforward networks can produce interesting codes only if the projections from each input neuron are allowed to have both positive and negative weights. This suggests that recurrent architecture or hidden layers may be necessary to compensate for the restrictions imposed by Dale's law.

2 Background

A combinatorial neural code of a given population of n neurons is the collection of all possible combinations of neurons that can be simultaneously active. More precisely, a code word is a subset

$$\sigma \subset [n] \stackrel{\text{def}}{=} \{1, \dots, n\}$$




Figure 1: One-layer feedforward network.

of neurons that are simultaneously active (i.e., fire within some small temporal window) at some point of time, while a combinatorial code is the collection of all such subsets

$$\mathcal{C} \subset 2^{[n]} \stackrel{\text{def}}{=} \{\text{subsets of } [n]\}.$$

Note that this notion of neural code describes only the set of possible response patterns of a network (Schneidman et al., 2006, 2011; Osborne et al., 2008). It does not include the “dictionary” of relationships between response patterns and network inputs.

We investigate what combinatorial codes can arise as response patterns in a one-layer feedforward network—that is, a collection of uncoupled neurons (perceptrons) (Rosenblatt, 1958), driven by an input layer of neurons (see Figure 1). The firing rates, $x_i(t) \geq 0$ of such neurons can be approximated by the equations

$$x_i = \phi_i \left(\sum_{a=1}^m U_{ia} y_a - \theta_i \right), \quad i = 1, \dots, n, \tag{2.1}$$

where $\theta_i > 0$ are the neuronal thresholds, U_{ia} are the effective strengths of the feedforward connections, $y_a(t) \geq 0$ are the firing rates of the neurons in the input layer, and the transfer functions $\phi_i : \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$ are monotone increasing.

For a nonnegative firing rate vector $x \in \mathbb{R}_{\geq 0}^n$, we denote the subset of coactive neurons as

$$\text{supp}(x) = \{i \in [n] \mid x_i > 0\}.$$

Given a particular choice of the transfer functions ϕ_i , the combinatorial neural code $\mathcal{C}(U, \theta) \subseteq 2^{[n]}$ of the network described by equation 2.1 is the collection of all possible subsets of neurons that can be coactivated by nonnegative firing rate inputs $y \in \mathbb{R}_{\geq 0}^m$.

$$\mathcal{C}(U, \theta) = \{\sigma = \text{supp}(x(y)) \mid y \in \mathbb{R}_{\geq 0}^m\}.$$

Our definition of $\mathcal{C}(U, \theta)$ captures all possible response patterns of the network, without any constraints on the network inputs. While it may be more realistic to assume that the inputs y are themselves constrained to a subset $Y \subset \mathbb{R}_{\geq 0}^m$, our primary interest is in combinatorial codes that are shaped by the structure of the feedforward network, as opposed to codes whose structure is largely inherited from the previous layer. As an extreme example, consider a relay network where the output layer is a copy of the input later and each input neuron relays its activity to exactly one output neuron (i.e., U is the identity matrix). Then any code could in principle be realized by an appropriate choice of Y , even though the network itself plays no role in shaping the code. Because we are interested in networks whose computational function is to shape the structure of the code, we consider the opposite extreme, where input firing rates are allowed to range over all nonnegative values and the code $\mathcal{C}(U, \theta)$ is completely determined by the feedforward connections U and the thresholds θ .

It is well known that a collection of n perceptrons described by equation 2.1 can be thought of as an arrangement of n hyperplanes in the nonnegative orthant $\mathbb{R}_{\geq 0}^m$, where m is the size of the input layer. Without loss of generality, we can assume that the monotone increasing functions $\phi_i(t)$ satisfy the conditions

$$\phi_i(t) = 0, \text{ if } t \leq 0, \quad \phi_i(t) > 0 \text{ if } t > 0. \tag{2.2}$$

This condition implies that given an input $y \in \mathbb{R}_{\geq 0}^m$,

$$x_i > 0 \iff \sum_{a=1}^m U_{ia}y_a - \theta_i > 0.$$

The combinatorial code of the network, $\mathcal{C}(U, \theta)$, can thus be identified with the list of the regions into which the above hyperplanes partition the positive orthant $\mathbb{R}_{\geq 0}^m$, as in Figure 2. More precisely,¹

$$\mathcal{C}(U, \theta) = \{ \sigma \subseteq [n] \mid (\cap_{i \in \sigma} H_i^+) \cap \cap_{j \notin \sigma} H_j^- \neq \emptyset \}, \tag{2.3}$$

where

$$H_i^+ = \{ y \in \mathbb{R}_{\geq 0}^m \mid (Uy)_i > \theta_i \}, \quad \text{and}$$

$$H_i^- = \{ y \in \mathbb{R}_{\geq 0}^m \mid (Uy)_i \leq \theta_i \}.$$

¹Here, by convention: $\cap_{i \in \emptyset} H_i^+ = \cap_{j \notin [n]} H_j^- = \mathbb{R}_{\geq 0}^m$.




Figure 2: Code words in $\mathcal{C}(U, \theta)$ correspond to regions in the positive orthant $\mathbb{R}_{\geq 0}^m$.

3 Results

We now describe our main results, with proofs deferred until the end.

First, observe that not every combinatorial neural code $\mathcal{C} \subset 2^{[n]}$ can be encoded by a one-layer feedforward network. A code \mathcal{C} is called convex (Curto, Itskov, Veliz-Cuba, & Youngs, 2013) if there exists a collection of convex subsets $\{V_i\}_{i=1}^n$ of $X \subseteq \mathbb{R}^d$ such that²

$$\mathcal{C} = \{\sigma \subseteq [n] \mid (\bigcap_{i \in \sigma} V_i) \cap \bigcap_{j \notin \sigma} (X \setminus V_j) \neq \emptyset\}. \tag{3.1}$$

Note that choosing $V_i = H_i^+$ and $X = \mathbb{R}_{\geq 0}^m$ yields equation 2.3; thus, any feedforward code $\mathcal{C}(U, \theta)$ is convex. Therefore, nonconvex codes cannot be encoded by a one-layer feedforward network.

Perhaps the simplest example of a nonconvex code is (Curto et al., 2013, code B3)

$$\mathcal{C} = 2^{[3]} \setminus \{\{1, 2, 3\}, \{2, 3\}, \{1\}\} = \{\{2\}, \{3\}, \{1, 2\}, \{1, 3\}\}. \tag{3.2}$$

We generalize this example by considering combinatorial codes that contain subcodes that are obstructions to convexity. A subcode of a combinatorial code $\mathcal{C} \subset 2^{[n]}$ is a collection of patterns obtained by restricting code words in \mathcal{C} to neurons inside a given subset $\sigma \subset [n]$:

$$\mathcal{C}(\sigma) \stackrel{\text{def}}{=} \{v = \sigma \cap \mu \mid \mu \in \mathcal{C}\} \subseteq 2^\sigma.$$

²Here, by convention: $\bigcap_{i \in \emptyset} V_i = \bigcap_{j \notin [n]} (X \setminus V_j) = X$.

Any subcode of a convex code is convex. Thus, possessing a nonconvex subcode is an obstruction to being realizable as a feedforward code $\mathcal{C}(U, \theta)$.

The following theorem is a generalization of example 3.2 and is a corollary of the more general theorem 5.

Theorem 1. *Assume a combinatorial code \mathcal{C} possesses a subcode $\mathcal{C}(\sigma) \subset 2^\sigma$ that satisfies the following conditions:*

- i. $\sigma \notin \mathcal{C}(\sigma)$,
- ii. *there exists a nonempty subset $v \subset \sigma$ with $|v| \leq |\sigma| - 2$ such that $v \notin \mathcal{C}(\sigma)$,*
- iii. *for every subset $\sigma' \subset \sigma$ with $|\sigma'| = |\sigma| - 1$, if $\sigma' \supset v$ then $\sigma' \in \mathcal{C}(\sigma)$.*

The code \mathcal{C} is nonconvex and cannot be realized by a one-layer feedforward network.

While the subcodes in theorem 1 provide a large repertoire of codes that can be used to rule out a feedforward code, these codes are not robust to noise. Specifically, a nonconvex code may differ from a convex code by as little as one neuron's participation in a single code word.³ Since some amount of noise or imprecision of neuronal firing is common in the brain, the notion of combinatorial code is too fine to meaningfully discern whether patterns of activity could be generated by a one-layer feedforward network. However, a coarser notion that still preserves combinatorial features of the neural code might potentially not suffer this deficiency.

We say that a code word $\sigma \in \mathcal{C}$ is a maximal pattern of the code \mathcal{C} if it is not contained in any larger pattern and denote the set of all maximal patterns of \mathcal{C} as

$$\max(\mathcal{C}) = \{\sigma \in \mathcal{C} \mid \sigma' \supsetneq \sigma \Rightarrow \sigma' \notin \mathcal{C}\}.$$

We refer to $\max(\mathcal{C})$ as a coarse combinatorial code. The coarse code is more robust than the full combinatorial code \mathcal{C} to subtractive noise that results in missing one or more neurons' participation in a code word.⁴

Moreover, this notion of code preserves core combinatorial features of sparse codes. Experimentally observed neural activity is often sparse (Hromádka et al., 2008; Barth & Poulet, 2012), that is, the number of coactive neurons in a code word σ is bounded as $|\sigma| \leq sn$, with the fraction of active neurons often as low as $s < 0.1$. Such combinatorial codes tend to consist almost exclusively of maximal patterns and are thus well approximated by their corresponding coarse codes.

³For example, adding either $v = \{1\}$ or $\sigma = \{1, 2, 3\}$ to the code in equation 3.2 would turn it into a convex code. Similarly, removing $\sigma' = \{1, 3\}$ from that code would also turn it into a convex code.

⁴Biologically, the subtractive noise may be a result of either a neuron's failure to spike or a spike misalignment, where the neuron's spike occurs too early or too late as compared to the spikes of other neurons in the code word.

This motivates us to investigate the possibility of ruling out a one-layer feedforward network based on the coarse combinatorial code $\max(C)$. Our central result is that this cannot be done. Surprisingly, one-layer feedforward networks may encode any prescribed set of maximal patterns.

Theorem 2 (No-Go Theorem). *For every collection $\mathcal{M} \subset 2^{[n]}$ of maximal patterns, there exists a one-layer feedforward network of perceptrons (see equation 2.1) with*

$$\max(C(U, \theta)) = \mathcal{M}.$$

In particular, it is not possible to infer a computational function for recurrent connections, or hidden layers, in a network from observations of the coarse combinatorial code alone. The proof of theorem 2 is constructive and uses tools from combinatorial topology. An unexpected corollary is the following topological fact:

Corollary 1. *For any finite abstract simplicial complex Δ , there exists a set of the form $\mathbb{R}_{\geq 0}^m \setminus \mathcal{P}$, for \mathcal{P} a polyhedron, which is homotopy equivalent to Δ .⁵*

This is perhaps counterintuitive because if the polyhedron is fully contained within $\mathbb{R}_{> 0}^m$, the complement has the homotopy type of a sphere. The richness in the topology emerges from the intersection of the polyhedron with the boundary of the positive orthant.

An important caveat for interpreting our no-go theorem is that many types of neurons in the brain conform to a strong constraint called Dale's law (Dayan & Abbott, 2001). Dale's law states that neurons have either purely excitatory or purely inhibitory synapses onto other neurons. More formally, a one-layer feedforward network (equation 2.1) respects Dale's law if there is a partition of the columns of the synaptic matrix into two families, $U = [U_+ | U_-]$, so that all entries of U_+ are nonnegative and those of U_- are nonpositive. It turns out that one-layer networks that respect Dale's law are capable of producing only an extremely restricted class of coarse combinatorial codes:

Proposition 1. *Suppose that the one-layer feedforward network (equation 2.1) respects Dale's law. Then the combinatorial code of this network has exactly one maximal pattern $\sigma_{\max} \in C(U, \theta)$, that is,*

$$\max(C(U, \theta)) = \{\sigma_{\max}\}.$$

This property, in particular, excludes most known sparse neural codes, such as place field codes or orientation tuning codes.

It is worth noting, however, that a one-layer feedforward network without the Dale's law constraint can be thought of as an approximation of

⁵A collection of sets Δ is a simplicial complex if it is closed under inclusion, that is, if $\sigma \in \Delta$ and $\nu \subset \sigma$, then $\nu \in \Delta$. See section 5 for details.

a two-layer feedforward network that respects Dale’s law. For example, assuming $\phi_i(y) = [y]_+ \stackrel{\text{def}}{=} \max(0, y)$, that the first layer synapses U^1 are excitatory while the second layer synapses U^2 respect Dale’s law, and moreover that the first layer has zero thresholds, one obtains the following equation for the firing rates x of the output neurons:

$$x = [U^2[U^1y]_+ - \theta]_+ = [U^2U^1y - \theta]_+ = [Uy - \theta]_+.$$

Observe that the resulting matrix $U = U^2U^1$ no longer has to respect Dale’s law, even though the component matrices U^1 and U^2 do so. Thus, two-layer feedforward networks that obey Dale’s law are capable of encoding any prescribed simplicial complex, that is, in the presence of a layer of inhibitory neurons, a feedforward network can encode topologically interesting stimuli.

4 Proofs of the Main Results

In order to understand the coarse neural codes $\max(\mathcal{C})$, it is convenient to consider the maximal possible code with the same $\max(\mathcal{C})$. This can be thought of as a completion of the code \mathcal{C} obtained by adding all the subsets of $\max(\mathcal{C})$; this results in a new code $\Delta(\mathcal{C}) \supseteq \mathcal{C}$,

$$\Delta(\mathcal{C}) = \{v \subseteq \sigma \mid \sigma \in \mathcal{C}\} = \{v \subseteq \sigma \mid \sigma \in \max(\mathcal{C})\}. \tag{4.1}$$

This collection of sets is closed under inclusion: $v \subset \sigma \in \mathcal{C}$ implies that $v \in \mathcal{C}$. A collection of sets with this property is called an abstract simplicial complex.

4.1 Convex Codes. It is well known that every abstract simplicial complex, that is, a code that satisfies $\mathcal{C} = \Delta(\mathcal{C})$, is a convex code (Wegner, 1967; Tancer, 2013). The following result shows that convex codes that are not simplicial complexes also have strong restrictions.

Theorem 3. *Assume that $\mathcal{C} \subset 2^{[n]}$ is a convex code of the form 3.2 and that a code word $v \in \Delta(\mathcal{C})$ violates the simplicial complex property, that is, $v \notin \mathcal{C}$. Then the localized complex,*

$$\Delta(\mathcal{C})|_v \stackrel{\text{def}}{=} \{\tau \subseteq [n] \setminus v \mid (\tau \cup v) \in \Delta(\mathcal{C})\} \subseteq 2^{[n] \setminus v},$$

is contractible.

Proof. Denote by $V_v = \bigcap_{j \in v} V_j$. It is easy to see that $\Delta(\mathcal{C})|_v$ is the nerve of the cover of V_v by the convex sets $\tilde{V}_j \stackrel{\text{def}}{=} V_j \cap V_v$. Moreover, since $v \notin \mathcal{C}$, the sets \tilde{V}_j cover V_v , in other words, $V_v = \bigcup_{j \notin v} \tilde{V}_j$. Therefore we can use the nerve

lemma (see Hatcher, 2002, Corollary 4G.3, p. 460) and conclude that the simplicial complex $\Delta(\mathcal{C})_{|v}$ is homotopy equivalent to the set V_v and thus contractible.

We now give the proof of theorem 1 as a corollary of theorem 3:

Proof of Theorem 1. Given the conditions i to iii of theorem 1, it is easy to see that the maximal patterns of the localization $\Delta(\mathcal{C}(\sigma))_{|v}$ can be described as

$$\max(\Delta(\mathcal{C}(\sigma))_{|v}) = \{\tau \subset (\sigma \setminus v) \mid |\tau| = |\sigma| - |v| - 1\}.$$

Therefore, the simplicial complex $\Delta(\mathcal{C}(\sigma))_{|v}$ can be identified with the boundary of the $(|\sigma| - |v| - 1)$ -dimensional simplex $\sigma \setminus v$, and thus is not contractible. Therefore, $\mathcal{C}(\sigma)$ is not a convex code and thus neither is \mathcal{C} .

4.2 Proof of the No-Go Theorem 3. We prove the no-go theorem in two steps. First, we find an explicit construction of a feedforward network that encodes the nerve $\mathcal{N}(\Delta)$ of any simplicial complex Δ . We then use a result by Grünbaum (1970) showing the existence of an inverse nerve.

Recall that a face $\alpha \in \Delta$ of a simplicial complex is maximal if it is not contained in any larger face of Δ . Let $\Delta \subset 2^{[m]}$ be an abstract simplicial complex on m vertices and $\{\alpha_1, \dots, \alpha_n\} = \max(\Delta)$ be the maximal sets of Δ . The nerve of Δ is another abstract simplicial complex $\mathcal{N}(\Delta) \subseteq 2^{[n]}$ on n vertices, such that for any nonempty $v \in [n]$

$$v \in \mathcal{N}(\Delta) \iff \bigcap_{i \in v} \alpha_i \neq \emptyset.$$

Proposition 2. *Let $\Delta \subseteq 2^{[m]}$ be a simplicial complex with m vertices and n maximal faces $\{\alpha_1, \dots, \alpha_n\}$. Then there exists a strictly feedforward network (U, θ) such that*

$$\Delta(\mathcal{C}(U, \theta)) = \mathcal{N}(\Delta).$$

Moreover, the network parameters U and θ for such a network can be constructed with

$$U_{i,a} = \begin{cases} -n & \text{if } a \notin \alpha_i, \\ 1 & \text{if } a \in \alpha_i, \end{cases} \quad \text{and } \theta_i = \frac{1}{2}. \tag{4.2}$$

Proof. Given the choice, equation 4.2, the regions $H_i^+ \subset \mathbb{R}_{\geq 0}^m$ can be described by the inequalities $y_a \geq 0$ and

$$\sum_{a \in \alpha_i} y_a - n \sum_{a \notin \alpha_i} y_a > \frac{1}{2}. \tag{4.3}$$

Note that equation 2.3 implies that

$$\Delta(\mathcal{C}(U, \theta)) = \{\sigma \subseteq [n] \mid \bigcap_{i \in \sigma} H_i^+ \neq \emptyset\} \cup \emptyset.$$

Thus, to show that $\Delta(\mathcal{C}(U, \theta)) \supseteq \mathcal{N}(\Delta)$, we need to prove that for any nonempty $\sigma \in \mathcal{N}(\Delta)$,

$$\bigcap_{i \in \sigma} H_i^+ \neq \emptyset. \tag{4.4}$$

For any nonempty $\alpha \subseteq [m]$, we define a vector $y^\alpha \in \mathbb{R}_{\geq 0}^m$ as

$$y_a^\alpha \stackrel{\text{def}}{=} \begin{cases} 1 & \text{if } a \in \alpha \\ 0 & \text{if } a \notin \alpha. \end{cases}$$

Let $\sigma \subseteq [n]$ and define $\bigcap_{i \in \sigma} \alpha_i \stackrel{\text{def}}{=} \alpha(\sigma) \subseteq [m]$. By the definition of the nerve, if a nonempty $\sigma \in \mathcal{N}(\Delta)$, then $\alpha(\sigma) \neq \emptyset$ and plugging in $y^{\alpha(\sigma)}$, inequality 4.3 becomes $|\alpha_i| > \frac{1}{2}$. Thus, $y^{\alpha(\sigma)} \in \bigcap_{i \in \sigma} H_i^+$ and property 4.4 holds.

In order to show that $\Delta(\mathcal{C}(U, \theta)) \subseteq \mathcal{N}(\Delta)$, we need to prove that

$$\text{for any } \tau \notin \mathcal{N}(\Delta), \quad \bigcap_{i \in \tau} H_i^+ = \emptyset. \tag{4.5}$$

Assume the converse: for some $\tau \notin \mathcal{N}(\Delta)$, there exists $y \in \bigcap_{i \in \tau} H_i^+$. From the definition of the nerve, we obtain $\alpha(\tau) = \emptyset$ and thus for every $a \in [m]$,

$$|\{j \in \tau \mid a \in \alpha_j\}| \leq |\tau| - 1, \text{ and } |\{j \in \tau \mid a \notin \alpha_j\}| \geq 1. \tag{4.6}$$

Summing the inequalities in equation 4.3, over all $j \in \tau$, we obtain

$$\sum_{j \in \tau} \sum_{a \in \alpha_j} y_a - n \sum_{j \in \tau} \sum_{a \notin \alpha_j} y_a > \frac{|\tau|}{2}.$$

Thus, using equation 4.6, we obtain

$$\begin{aligned} & \sum_{a=1}^m y_a (|\tau| - 1 - n) \\ & \geq \sum_{a=1}^m y_a (|\{j \in \tau \mid a \in \alpha_j\}| - n |\{j \in \tau \mid a \notin \alpha_j\}|) > \frac{|\tau|}{2} > 0. \end{aligned}$$

Taking into account that $y_a \geq 0$ for all a , we obtain that $|\tau| \geq n + 1$, which is a contradiction. Thus, condition 4.5 holds.

To finish the proof of theorem 3 we use the following classical result of Grünbaum (1970). (For completeness, we also give Grünbaum’s explicit construction of the inverse nerve in the appendix.)

Theorem 4 (Grünbaum, 1970). *For every abstract simplicial complex Δ , there exists a simplicial complex $\tilde{\Delta}$, such that $\mathcal{N}(\tilde{\Delta}) = \Delta$.*

Proof of Theorem 2. Given a set $\mathcal{M} \subset 2^{[n]}$ of maximal patterns, choose an abstract simplicial complex $\tilde{\Delta}$ such that $\mathcal{N}(\tilde{\Delta}) = \Delta(\mathcal{M})$. By proposition 6, prescription 4.2 yields a one-layer network network such that

$$\Delta(\mathcal{C}(U, \theta)) = \mathcal{N}(\tilde{\Delta}) = \Delta(\mathcal{M}) \Rightarrow \max(\mathcal{C}(U, \theta)) = \mathcal{M}.$$

Proof of Corollary 1. Given an abstract simplicial complex $\Delta \subseteq 2^{[n]}$, consider a one-layer neural network (U, θ) from theorem 2, so that $\max(\mathcal{C}(U, \theta)) = \max(\Delta)$. The code words in $\mathcal{C}(U, \theta)$ are identified with nonempty polyhedra in the positive orthant as in equation 2.3, with the empty code word $\sigma = \emptyset$ identified with the polyhedron $\mathcal{P} \stackrel{\text{def}}{=} \bigcap_{i \in [n]} H_i^-$.

Now observe that the region $Y = \mathbb{R}_{\geq 0}^n \setminus \mathcal{P}$ is covered by the convex sets $\{H_i^+\}_{i \in [n]}$, and for any nonempty $\sigma \subseteq [n]$, the intersection $\bigcap_{i \in \sigma} H_i^+$ is nonempty if and only if $\sigma \in \Delta(\mathcal{C}(U, \theta)) = \Delta$. Thus, Δ is the nerve of the cover of Y by the convex sets $\{H_i^+\}$. Therefore, by the nerve lemma, Y is homotopy equivalent to Δ .⁶

4.3 Networks constrained by Dale’s Law.

Proof of Proposition 1. Suppose each neuron in the input layer is either excitatory or inhibitory, that is, one can reorder the input neurons so that $U = [U_+ | U_-]$. Denote by $\sigma_{\max} \subseteq \{1, \dots, n\}$ the set of all neurons in the output layer that are either “on” in the absence of external drive or receive at least one excitatory connection:

$$\sigma_{\max} = \{i \mid -\theta_i > 0\} \cup \{i \mid \exists j \text{ with } U_{ij} > 0\}.$$

Because $U = [U_+ | U_-]$, setting the firing rates of the excitatory neurons sufficiently high and the firing rates of the inhibitory neurons to zero yields a firing rate vector $x \in \mathbb{R}_{\geq 0}^n$ with support $\sigma_{\max} \in \mathcal{C}(U, \theta)$.⁷ Moreover, any element of the code $\mathcal{C}(U, \theta)$ must be a subset of σ_{\max} .

⁶See e.g., Hatcher, 2002, corollary 4G.3 p. 460.

⁷Setting $y_j > \max_{\{i: U_{ij} > 0\}} \left\{ \frac{|\theta_i|}{U_{ij}} \right\}$ is sufficient.

5 Conclusion

Motivated by the idea that recurrent or many-layer feedforward architecture may be necessary to shape the neural code, we set out to find combinatorial codes that cannot arise from one-layer feedforward networks. Although we found a large class of such examples, the nonconvex codes, we also observed that in the presence of noise, they would be virtually indistinguishable from other codes having the same maximal patterns. When we considered coarse combinatorial codes, which keep track of only the maximal patterns, we found that there do not exist any codes that cannot be realized by a one-layer feedforward network.⁸ On the contrary, any coarse code in principle can be learned by a one-layer network. Our results suggest that recurrent connections, or multiple layers, may not be necessary for shaping the neural code.

Appendix

For completeness, we include the construction due to Grünbaum (1970) of a simplicial complex $\tilde{\Delta}$ in the collection $\mathcal{N}^{-1}(\Delta)$ for any abstract simplicial complex Δ .

Suppose $\Delta \subseteq 2^{[n]}$ is an abstract simplicial complex with maximal faces $\max(\Delta) = \{\alpha_1, \dots, \alpha_k\}$. Write $I(\Delta)$ for a collection of vertices of Δ that can be written as an intersection of its maximal faces,

$$I(\Delta) = \left\{ v \in 2^{[n]} \mid \exists \sigma \subseteq [k] \text{ with } \{v\} = \bigcap_{i \in \sigma} \alpha_i \right\}.$$

Now, choose the vertices of $\tilde{\Delta}$ to be the maximal faces of Δ along with vertices of Δ , which do not appear as intersections of maximal faces:

$$V = \max(\Delta) \cup (2^{[n]} \setminus I(\Delta)).$$

The maximal faces $\max(\tilde{\Delta}) = \{\beta_i\}_{i \in [n]}$, are chosen in one-to-one correspondence with the vertices of Δ , with face β_i supported on elements of V that

⁸Note that enumerating (counting) coarse codes is the same as enumerating simplicial complexes or monotone functions and is often called Dedekind's problem. Finding the exact formula (or an algorithm) for such a number is a long-standing problem in combinatorics that is still unsolved. An asymptotic estimate of the $\log_2 M(n)$, where $M(n)$ is the number of simplicial complexes on n vertices is the following (Kleitman & Markowsky, 1975):

$$\binom{n}{\lfloor n/2 \rfloor} \leq \log_2 M(n) \leq \binom{n}{\lfloor n/2 \rfloor} \left(1 + O\left(\frac{\log n}{n}\right) \right).$$

contain vertex i :

$$\beta_i = \begin{cases} \{\alpha_j \in \max(\Delta) \mid i \in \alpha_j\} & \text{if } i \in I(\Delta) \\ \{\alpha_j \in \max(\Delta) \mid i \in \alpha_j\} \cup \{i\} & \text{if } i \notin I(\Delta). \end{cases}$$

By construction, the only maximal intersections of the elements of $\max(\mathcal{N}(\tilde{\Delta}))$ are the elements of V corresponding to elements of $\max(\Delta)$, thus $\mathcal{N}(\tilde{\Delta}) = \Delta$ and $\tilde{\Delta} \in \mathcal{N}^{-1}(\Delta)$.

Example. Let $\Delta \subseteq 2^{[6]}$ have maximal faces $\max(\Delta) = \{\{1, 2, 3\}, \{2, 3, 4\}, \text{ and } \{2, 5, 6\}\}$. The set $I(\Delta)$ consists of the single vertex 2, which is the only vertex that appears as an intersection of maximal faces. The complex $\tilde{\Delta}$ thus has vertices $\{\{1, 2, 3\}, \{2, 3, 4\}, \{2, 5, 6\}, 1, 3, 4, 5, 6\}$ and maximal faces $\{\{1, \{1, 2, 3\}\}, \{\{1, 2, 3\}, \{2, 3, 4\}\}, \{\{1, 2, 3\}, \{2, 5, 6\}\}, \{3, \{1, 2, 3\}, \{2, 3, 4\}\}, \{4, \{2, 3, 4\}\}, \{5, \{2, 5, 6\}\}, \{6, \{2, 5, 6\}\}$.

Acknowledgments

This work was supported by NSF DMS-1122519. V.I. is thankful to Carly Klivans for pointing to the reference (Grünbaum, 1970); the authors are also grateful to Carina Curto for comments on the manuscript.

References

- Barth, A. L., & Poulet, J. F. (2012). Experimental evidence for sparse firing in the neocortex. *Trends Neurosci.*, 35(6), 345–355.
- Curto, C., Itskov, V., Veliz-Cuba, A., & Youngs, N. (2013). The neural ring: An algebraic tool for analyzing the intrinsic structure of neural codes. *Bull Math Biol*, 75, 1571–1611.
- Cybenko, G. (1989). Approximations by superpositions of sigmoidal functions. *Mathematics of Control, Signals, and Systems*, 2(4), 303–314.
- Dayan, P., & Abbott, L. F. (2001). *Theoretical neuroscience*. Cambridge, MA: MIT Press.
- Douglas, R. J., & Martin, K. A. C. (2007). Recurrent neuronal circuits in the neocortex. *Current Biology*, 17(13), 496–500.
- Grünbaum, B. (1970). Nerves of simplicial complexes. *Aequationes Math.*, 4, 63–73.
- Hatcher, A. (2002). *Algebraic topology*. Cambridge: Cambridge University Press.
- Hornik, K. (1991). Approximation capabilities of multilayer feedforward networks. *Neural Networks*, 4 (2), 251–257.
- Hromádka, T., Deweese, M. R., & Zador, A. M. (2008). Sparse representation of sounds in the unanesthetized auditory cortex. *PLoS Biol*, 6(1).
- Kleitman, D., & Markowsky, G. (1975). On Dedekind's problem: The number of isotone Boolean functions. II. *Trans. Amer. Math. Soc.*, 213, 373–390.
- Luczak, A., Barthó, P., Harris, K. D. (2009). Spontaneous events outline the realm of possible sensory responses in neocortical populations. *Neuron*, 62(3), 413–425.

- Osborne, L. C., Palmer, S. E., Lisberger, S. G., & Bialek, W. (2008). The neural basis for combinatorial coding in a cortical population response. *J. Neurosci.*, *28*(50), 13522–13531.
- Rosenblatt, F. (1958). The perceptron: A probabilistic model for information storage and organization in the brain. *Psychol. Rev.*, *65*(6), 386–408.
- Schneidman, E., Berry, M. J., Segev, R., & Bialek, W. (2006). Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, *440*(7087), 1007–1012.
- Schneidman, E., Puchalla, J. L., Segev, R., Harris, R. A., Bialek, W., & Berry, M. J. (2011). Synergy from silence in a combinatorial neural code. *J. Neurosci.*, *31*(44), 15732–15741.
- Tancer, M. (2013). Intersection patterns of convex sets via simplicial complexes: A survey. In J. Pach, (Ed.), *Thirty essays on geometric graph theory* (pp. 521–540). New York: Springer.
- Wegner, G. (1967). *Eigenschaften der Nerven homologisch-einfacher Familien im R^n* . Doctoral dissertation, Universität Göttingen.

Received February 3, 2014; accepted May 19, 2014.