

Using persistent homology to reveal hidden information in neural data

Gard Spreemann¹, Benjamin Dunn², Magnus Bakke Botnan¹, and Nils A. Baas¹

¹{gard.spreemann, magnus.botnan, nils.baas}@math.ntnu.no, Department of Mathematical Sciences, Norwegian University of Science and Technology, 7491 Trondheim, Norway

²benjamin.dunn@ntnu.no, Kavli Institute for Systems Neuroscience, Norwegian University of Science and Technology, 7491 Trondheim, Norway

October 23, 2015

Abstract

We propose a method, based on persistent homology, to uncover topological properties of a priori unknown covariates of neuron activity. Our input data consist of spike train measurements of a set of neurons of interest, a candidate list of the known stimuli that govern neuron activity, and the corresponding state of the animal throughout the experiment performed. Using a generalized linear model for neuron activity and simple assumptions on the effects of the external stimuli, we infer away any contribution to the observed spike trains by the candidate stimuli. Persistent homology then reveals useful information about any further, unknown, covariates.

1 Introduction

Due to its apparent simplicity, physical space has long served as a model system for internally generated representations in the brain [32]. In their pioneering work [26], O’Keefe and Dostrovsky discovered in the hippocampi of rats certain neurons, called *place cells*, that seemed to be active at a level far above their baseline when the animal was located in a specific region in space. These regions of elevated activity are known as *place fields*. It has also been demonstrated [23, 31, 30] that neurons tune not only to spatial position, but also to other external covariates, such as for example head direction. Place fields thus do not exist only for an animal’s physical surroundings — we shall rather think of them as regions in an abstract state space, thus generalizing to the more fundamental concept of neural coding. An example of the place fields of three cells recorded in an experiment by Buzsáki et al. [22, 21] can be seen in Figure 1.




Figure 1: The firing activity of three neurons as a rat explores a square box. For the two leftmost neurons the activity is plotted against spatial position (both axes) while the activity for the rightmost one is plotted against head direction (horizontal axis). The regions of elevated activity, the spatial and “head direction state space” place fields, are clearly visible. In the spatial plots, the cells were least active in the blue areas, and most in the red. White areas were not visited by the rat during the experiment. Activity and position data for the plots come from [22, 21].

If a collection of neurons divide the animal’s state space into reasonably nice place fields, it is perhaps not surprising that the activity of these neurons somehow encode information about the state space. A question to pose is then how much of that space can be recovered from just recordings of the electrical activity of the neurons. We will consider this question from a novel point of view in which the state space is partially known, and see what properties of the unknown parts can be glanced from such recordings. This, in turn, will give us information about unknown external covariates of neuron activity by using a generalized linear model (GLM) [20] to infer away the contributions of any known covariates. GLMs have recently been applied in neuroscience as a structured method to uncover dependencies as well as structure in data [27, 33, 28].

It should be pointed out that in order to compute or visualize the place fields themselves, as done for Figure 1, the activity data for the cells must be supplemented by positional and other state space data, e.g. head direction or other external covariates, for the animal. Without prior knowledge of the nature of the covariates, i.e. prior knowledge of the unknown state space, it is of course not apparent what data should be recorded. Since it is therefore too much to ask to know the actual place fields themselves, we take a cue from topology and instead try for only information on how the place fields *fit together*. Indeed, if a set of place fields (which are themselves unknown) intersect, we should be more likely to observe the corresponding cells

firing simultaneously, or nearly so, in time. Figure 2 illustrates the idea. This will be our guide when we build a combinatorial topological space from neuron activity data in the form of firing event time series (so-called *spike trains*). We will, in other words, transform our original time series data — the collection of spike trains — into geometric data, which we will then study with tools from topology.




Figure 2: Place cell cofiring as a proxy for place field intersections. As the animal moves along the indicated path (**left**), we might observe the (highly idealized) firing events (**center**) of the corresponding place cells. The firing events in the leftmost box are indicative of the triple intersection of the red, green and blue place fields, and those in the rightmost box are indicative of the double intersection of the red and blue fields. After the space has been thoroughly explored, the intersection data obtained from the cofiring of the spike trains, define a simplicial complex (**right**) by means of the nerve construction. The simplicial complex and the space covered by place fields have the same homotopy type (in this example, that of a point).

Cofiring is thus a proxy that gives us approximate knowledge of how all the place fields intersect (doubly, triply, etc.). We turn to topology to ask what we can learn from these data. The natural next step, when armed with intersection data, is to construct the *nerve* of the place fields. The unfamiliar reader will find the formal definition in Section 2, but intuitively the process goes as follows: for every cell, draw a point; whenever two place fields intersect, connect the corresponding two points with a line segment; whenever three place fields intersect, draw a filled triangle between the three corresponding points; whenever four place fields intersect, draw a filled tetrahedron between the four corresponding points; and so forth. The abstract combinatorial construction so obtained is one example of a simplicial complex, a kind of topological space. It is a famous theorem (Theorem 2) in algebraic topology that this simplicial complex shares an important property with the space covered by the objects whose intersections we consider, i.e. the state space covered by place fields in our case.

The property that the nerve of the place fields shares (under some assumptions) with the space covered by the place fields is that of *homotopy type*. A reader unfamiliar with elementary topics in algebraic topology may want to see [16], or can think of spaces with the same homotopy type as those that can be continuously deformed into one another *without tearing*. Thus, a square, a disk and a point are of the same homotopy type, while they are of different types from an annulus, which is again of different type from a sphere. This answers the question of what we may hope to recover from the cells' firing information: not the full state space (i.e. not the environment's full geometry in a setting with only spatial place cells), but rather its homotopy type, which again says something about the kind of *holes* the space has. We should for example, in the purely spatial case, at the very least be able to tell whether there is an obstacle, such as an impassable column, somewhere in a box the rat explores (making it, in the eyes of homotopy, an annulus rather than a box). Moreover, the circular nature of a covariate such as head direction also influences the homotopy type of the state space, and is therefore detectable.

We have so far glossed over what it means for cells to cofire. Clearly a single cofiring event is not sufficient to declare place fields as overlapping, as cells may fire spuriously with the animal outside of place fields. Too strict a notion of cofiring is obviously not good either. Moreover, looking to the real-world situation in Figure 1, the *degree* of intersection we should demand to declare two fields as intersecting is not as clear-cut as presented in the idealized setting of Figure 2. *Persistent homology*, which we recap in Section 2, offers a way to consider all degrees of cofiring and intersections simultaneously and as one mathematical object. While not as sharp an invariant as homotopy type, (persistent) homology still captures the number of holes in the space under consideration, their dimensionality, and to some extent also their size.

In summary, the correlations of spike trains allow us to build combinatorial spaces that reflect topological properties of some partially unknown animal state space. In the purely spatial setting, several papers [7, 8, 15] have already demonstrated the feasibility of recovering properties of the animal’s physical environment in this way. We now propose a new method wherein topological properties of a partially unknown state space are uncovered by successively accounting for known covariates using a GLM.

1.1 Contributions

Already in [26] it was clear that spatial position is not the only influence on the firing of certain neurons. Other covariates, confirmed or suspected, include [23, 31, 30]: head direction, theta wave phase, behavioral tasks (eating, biting, sleeping, grooming, etc.) and tactile or olfactory sensory stimuli. In addition to these external influences, neurons are connected to a set of neighboring neurons, and are excited or inhibited by the activity of those neighbors.

In this paper, we consider a general setting of neural data governed by a list of suspected possible covariates (also referred to as stimuli or influences). We further assume that the researcher supplies spike train recordings of the relevant neurons, and, in addition, measurements of the suspected physical covariates throughout the experiment. We then set out to answer the following questions:

Q0: What is the (persistent) homology encoded by the covariates, i.e. what is the (persistent) homology of the corresponding state space?

Q1: Does the list of suspected covariates adequately explain the observations?

Q2: If not, do the covariates missing from the list of candidates encode non-trivial homological information?

The purely spatial version of question Q0 has already been examined by others [7, 8, 15]. Standard spectral methods, for example using Wigner’s semicircle law [34] on the correlations of neurons, may be sufficient to answer question Q1. Still, we will in this paper approach that question using persistent homology, as done in [15]. We believe that question Q2 has not been considered before, and that its answers can be useful to neuroscientists. The technique we outline in this paper seeks to answer that question. Moreover, the method is quite general and may be useful in other applications both inside and outside of neuroscience.

1.2 Outline

Section 2 reviews the necessary theory of persistent homology at the level of elementary applied topology, recapping also some of the major results that we use. We then briefly describe the model we use for neuron activity.

Section 3 describes the process of inferring away the contributions of covariates until the list of suspected such is empty, whereupon information about the homology of any hidden covariates is revealed. Demonstrations of the efficacy of the technique with simulated neuronal data under various conditions follow in Section 4.

Finally, Section 5 sketches some possible applications of our technique outside of neuroscience. We also attempt to place it on a firmer and more general mathematical ground.

1.3 Related work

Curto and Itskov [7] were the first to employ tools from algebraic topology to study neural data in general, and in particular to attempt reconstruction of the homology of an animal’s physical environment from such data.

Dabaghian, Mémoli, Frank, and Carlsson [8] form a model of neuron activity, and use persistent homology very much along the lines sketched above to study qualitative properties of this model. The properties include the time taken to form a homologically correct representation of the environment, and the robustness of this representation with regard to model parameters. Arai, Brandt, and Dabaghian [1] improve on this model, add theta wave phase-precession, and study its effect on the persistent homology learning of the environment.

Giusti, Pastalkova, Curto, and Itskov [15] very recently showed how persistent homology can help determine whether the information encoded in the spike trains of neurons encodes for something *geometric*, or is just random, thus answering our question Q1.

Our work is in part based on the same approach as [15], but seeks also to answer questions Q0 and Q2 within one framework, and should also be applicable in a very general setting not specific to neuroscience.

2 Background

In this section we survey the necessary tools from persistent homology in particular and topological data analysis (TDA) in general, and fix basic notation. Thorough introductions can be found in [13, 10, 11]. Familiarity with elementary algebraic topology is assumed (see [16] for an introduction).

The p -simplices of a simplicial complex K with vertex set V will be written like $[i_0, i_1, \dots, i_p] \subseteq K$, where each $i_0, \dots, i_p \in V$ are implicitly distinct. The geometric realization of K will be denoted $|K|$. Unless otherwise is noted, $H_k(K)$ denotes the k ’th simplicial homology of K , computed with coefficients in $\mathbb{Z}/2\mathbb{Z}$. All simplicial complexes we consider are assumed to be finite.

As hinted at in Section 1, our basic data will be *intersection information*, or a proxy thereof. The following basic construction that lies at the heart of TDA encodes such data as a simplicial complex.

Definition 1. Let X be a topological space. The *nerve* of a collection of sets $\mathcal{U} = \{U_i \subseteq X \mid i \in I\}$ is the simplicial complex $N\mathcal{U}$ defined by

$$[i_0, i_1, \dots, i_p] \subseteq N\mathcal{U} \iff U_{i_0} \cap U_{i_1} \cap \dots \cap U_{i_p} \neq \emptyset.$$

The nerve construction’s importance to TDA stems from its ability to discretely encode the homotopy type of a topological space, as the *nerve theorem* shows. The theorem exists in a more general form, but for the purpose of our paper it suffices to state it for metric spaces.

Theorem 2. Let X be a metric space, and let $\mathcal{U} = \{U_i \mid i \in I\}$ be a finite closed cover of X with the property that for all subsets $J \subseteq I$, the intersection $\bigcap_{j \in J} U_j$ is either empty or contractible. Then X and $|\mathcal{N}\mathcal{U}|$ have the same homotopy type.

A cover satisfying the above condition is said to be *good*. Throughout, we let $B(x; r)$ denote the closed ball of radius r centered at $x \in \mathbb{R}^n$. One aspect of TDA is the study of point clouds — finite point sets in Euclidean space — by applying Theorem 2 to a nerve of closed balls.

Definition 3. Let $P = \{p_1, p_2, \dots, p_N\} \subseteq \mathbb{R}^n$ be a point cloud. The Čech complex of P at scale ε is the nerve

$$\check{C}(P; \varepsilon) = N\{B(p; \varepsilon/2) \mid p \in P\}.$$

We write $[i_0, \dots, i_k] \subseteq \check{C}(P; \varepsilon)$ for the k -simplex corresponding to the intersection $B(p_{i_0}; \varepsilon) \cap \dots \cap B(p_{i_k}; \varepsilon)$.

Since a finite set of closed balls is a good cover of its union, Theorem 2 ensures that the purely combinatorial nerve reflects the homotopy type of that union. Thus, with H_*^{sing} denoting singular homology with coefficients in $\mathbb{Z}/2\mathbb{Z}$,

$$H_*^{\text{sing}} \left(\bigcup_{p \in P} B(p; \varepsilon/2) \right) \cong H_*^{\text{sing}} (|\check{C}(P; \varepsilon)|) \cong H_* (\check{C}(P; \varepsilon)). \quad (1)$$

In equation (1), the first isomorphism is a consequence Theorem 2, while the second is a standard result in algebraic topology. Computing the right hand side amounts to doing linear algebra.

2.1 Persistent homology

If a point cloud is assumed to have been sampled from some unknown subspace X of \mathbb{R}^n , then it is reasonable to apply equation (1) in order to reconstruct $H_*^{\text{sing}}(X)$. This of course begs the question: what is the “correct” scale at which to view the point cloud, or, explicitly in the case of the Čech complex, what (if any) is the “right” ball radius? Persistent homology sidesteps the question by encompassing *all* scales in one unifying construction.

We define the necessary constructions in a rather compact language below. The unfamiliar reader may want to read the wordier survey [12].

Definition 4. A *persistence module* \mathbf{V} (over \mathbb{R}) is a collection of finite-dimensional \mathbf{k} -vector spaces $\{\mathbf{V}(t) \mid t \in \mathbb{R}\}$ and linear maps $\{\mathbf{V}(s \leq t) : \mathbf{V}(s) \rightarrow \mathbf{V}(t) \mid s, t \in \mathbb{R}, s \leq t\}$ satisfying $\mathbf{V}(s \leq s) = \text{id}_{\mathbf{V}(s)}$ and $\mathbf{V}(s' \leq t) \circ \mathbf{V}(s \leq s') = \mathbf{V}(s \leq t)$ for all $s \leq s' \leq t$.

In other words, \mathbf{V} is functor from \mathbb{R} to finite dimensional \mathbf{k} -vector spaces. We define an *interval* $I \neq \emptyset \subseteq \mathbb{R}$ to be a set such that if $s, t \in I$, then $s' \in I$ for all s' satisfying $s \leq s' \leq t$.

Definition 5. The *interval persistence module* χ_I on an interval $I \subseteq \mathbb{R}$ is given by

$$\chi_I(s) = \begin{cases} \mathbf{k} & \text{if } s \in I \\ 0 & \text{otherwise} \end{cases}$$

where any two nontrivial vector spaces are connected by the identity map.

Direct sums of persistence modules are defined point-wise, i.e. $(\mathbf{V} \oplus \mathbf{W})(s) = \mathbf{V}(s) \oplus \mathbf{W}(s)$ and similarly for the maps, and those that can only be written as trivial direct sums are called *indecomposable*. It is easy to see that interval persistence modules are indecomposable. An existence theorem [6] guarantees that there for every \mathbf{V} exists a multiset $B(\mathbf{V})$ of intervals in \mathbb{R} such that

$$\mathbf{V} \cong \bigoplus_{I \in B(\mathbf{V})} \chi_I.$$

That this decomposition is essentially unique is ensured by the Azumaya–Krull–Remak–Schmidt theorem [2]. In particular, (the isomorphism class of) a persistence module \mathbf{V} is completely determined by the multiset $B(\mathbf{V})$. This is the famous *barcode* description of \mathbf{V} . Moreover, an interval in \mathbb{R} can be canonically identified with a point in $(\mathbb{R} \cup \{\pm\infty\})^2$ by means of its endpoints. Hence, a persistence module is completely described by a multiset of points in $(\mathbb{R} \cup \{\pm\infty\})^2$; we shall refer to this description as the *persistence diagram* of \mathbf{V} .

Definition 6. A finite filtration K of simplicial complexes

$$K_{\varepsilon_0} \subseteq K_{\varepsilon_1} \subseteq \cdots \subseteq K_{\varepsilon_N}$$

for $-\infty = \varepsilon_0 < \varepsilon_1 < \cdots < \varepsilon_N$ gives rise to the k 'th *persistent homology module* $PH_k(K)$ of K defined by

$$PH_k(K)(s) = H_k(K_{\max\{\varepsilon_i \mid \varepsilon_i \leq s\}})$$

with the linear maps induced by inclusions. This is commonly known as a *sublevel* persistence module.

We will often refer to $PH_*(K)$ as a persistent homology module without specifying the homology dimension. Observe that as K is a finite filtration of simplicial complexes, there will only be a finite number of intervals in $B(PH_*(K))$.

It is often useful to reduce the information content of persistent homology to an even simpler, if incomplete, descriptor. In line with the k 'th Betti number being defined as the rank of the k 'th homology group, we define the following.

Definition 7. Let \mathbf{V} be a persistence module. Its *Betti curve* is the function $\beta : \mathbb{R} \rightarrow \mathbb{N}$ defined by $\beta(s) = \dim(\mathbf{V}(s))$. When referring to persistent homology modules $PH_*(K)$, we write β_k for the Betti curve of $PH_k(K)$.

There are obviously non-isomorphic persistence modules that give rise to the same Betti curve, so this descriptor discards information. Its advantage is that, as functions, Betti curves are easy to compare, especially if considered as bounded functions $\mathbb{R} \rightarrow \mathbb{R}$. This becomes particularly useful if the persistence diagrams contain a great number of points. See also Sections 2.3 and 2.4.

2.2 Some filtered simplicial complexes

Finite filtrations of simplicial complexes abound in TDA. One natural example is the Čech filtration

$$\check{C}(P; \varepsilon_1) \subseteq \check{C}(P; \varepsilon_2) \subseteq \cdots \subseteq \check{C}(P; \varepsilon_N)$$

of a finite point cloud P given a sequence of radii $\varepsilon_1 < \varepsilon_2 < \cdots < \varepsilon_N$. We write $\check{C}(P)$ for the filtration.

The Čech filtration of a point cloud is of theoretical importance since Theorem 2 applies not only to each of its individual complexes (as in equation (1)), but also functorially so [4] (respecting maps). In many applications, including the one in this paper, the data we are given do not consist of points in Euclidean space, but rather of relationships (such as distances) between points in an unknown metric space. Many simplicial complexes can be built from such data, for example on graphs [17].

Definition 8. Let $G = (V, E)$ be a graph. The *flag complex* or *clique complex* of G is the largest simplicial complex $\text{fl}(G)$ whose 1-skeleton is E .

If a graph $G = (V, E)$ has edge weights $w : E \rightarrow \mathbb{R}$, defining

$$G_\varepsilon = (V, w^{-1}((-\infty, \varepsilon]))$$

gives a filtration of flag complexes

$$V \subseteq \text{fl}(G_{\varepsilon_1}) \subseteq \text{fl}(G_{\varepsilon_2}) \subseteq \cdots \subseteq \text{fl}(G_{\varepsilon_N})$$

for a sequence $\varepsilon_1 < \varepsilon_2 < \cdots < \varepsilon_N$.

The terms “flag complex” and “Čech complex”, and their corresponding notation, will also interchangeably refer to the filtered versions.

Flag complexes can be constructed, and be useful, for any underlying graph. However, they play a special role when that graph is a metric space. Indeed, if $P = \{p_1, \dots, p_N\} \subseteq \mathbb{R}^n$ is a point cloud, then the complete graph G on the vertices $\{1, \dots, N\}$, with edge weights

$$w(i, j) = \|p_i - p_j\|_2$$

gives a (filtration of) flag complex(es) called the *Vietoris–Rips complex* of P . We write $\text{VR}(P)$ for the entire filtration and $\text{VR}(P; \varepsilon)$ for the complex at scale ε in the filtration.

We obviously have the inclusion $\check{C}(P; \varepsilon) \subseteq \text{VR}(P; \varepsilon)$ at any scale ε . Conversely, if the points $p_0, \dots, p_k \in P$ are arranged at the vertices of the standard k -simplex in Euclidean space and scaled so that they are ε apart pairwise, then $[0, 1, \dots, k] \subseteq \check{C}(P; \sqrt{2}\varepsilon)$. The standard simplices are in fact the worst cases [11], and so

$$\check{C}(P; \varepsilon) \subseteq \text{VR}(P; \varepsilon) \subseteq \check{C}(P; \sqrt{2}\varepsilon)$$

for any scale ε . This sandwiching of $\check{C}(P)$ and $\text{VR}(P)$ implies that homological features that persist under the morphism $H_k(\text{VR}(P; \varepsilon)) \rightarrow H_k(\text{VR}(P; \sqrt{2}\varepsilon))$ correspond to features of $H_k(\check{C}(P; \varepsilon))$, and vice versa. The persistent homology of the Vietoris–Rips complex thus contains important information from that of the Čech complex, while using only distance data.

2.3 Comparing persistence modules

We now define a family of metrics on multisets in $(\mathbb{R} \cup \{\pm\infty\})^2$, which in turn pull back to (extended) (pseudo)metrics on persistence modules. Such metrics allow us to compare persistence modules with one another in a quantitative way.

Definition 9. Let A and A' be finite multisets in $(\mathbb{R} \cup \{\pm\infty\})^2$, i.e. persistence diagrams, and let Δ denote the diagonal with countably infinite multiplicity at every point. Write M for the set of all multiset bijections $A \cup \Delta \rightarrow A' \cup \Delta$. For $q = 1, 2, \dots, \infty$, let d_q^E denote the extended Euclidean q -metric on $(\mathbb{R} \cup \{\infty\})^2$. Then for $p = 1, 2, \dots$ define the *Vaserštejn distance*

$$d_{p,q}(A, A') = \inf_{f \in M} \left(\sum_{a \in A} \left(d_q^E(a, f(a)) \right)^p \right)^{1/p}$$

and for $p = \infty$ the *bottleneck distance*

$$d_{\infty,q}(A, A') = \inf_{f \in M} \sup_{a \in A} d_q^E(a, f(a)).$$

The computation of these distances is equivalent to a maximum bipartite matching problem, and so can be performed in $\mathcal{O}(\max(|A|, |A'|)^3)$ time using the (improved) Kuhn–Munkres algorithm [19, 24]. For the purpose of our paper we will only consider $p = 1$ and $q = 2$.

We remark that the bottleneck distance $d_{\infty,q}$ is of great theoretical importance because of its role in *stability theorems* [5]. For example, if P' is an ε -perturbation of a point cloud P in Euclidean space, then $d_{\infty,\infty}(PH_*(\check{C}(P)), PH_*(\check{C}(P'))) \leq \varepsilon$ and $d_{\infty,\infty}(PH_*(VR(P)), PH_*(VR(P'))) \leq \varepsilon$. In particular, persistent homology is *stable* with respect to perturbation of the input data.

2.4 Statistics of persistence modules and Betti curves

Later in the paper, we will build filtered simplicial complexes from measurements, and we may wonder how persistent homology depends on the statistical properties of the measurements.

The survey of Kahle [18] covers much of what is known about some special cases of random simplicial complexes, namely those that are flag complexes of Erdős–Rényi random graphs¹ and those that are Vietoris–Rips complexes of random points from Euclidean space (the latter will hereafter be referred to as *random geometric complexes*). While most results known are asymptotic in the number of vertices in the complex, and thus of little direct relevance to our setting, some qualitative conclusions can be drawn also about the finite case assuming the the size of the vertex set is not too small.

The first qualitative observation is that non-bounding cycles, i.e. homology generators, are likely to occur only in ER complexes when the edge probability parameter is in a certain range, and this range becomes narrower as the homology dimension grows. Within the “allowed range”, however, a large Betti number may occur. As the homology dimension increases, the allowed range shrinks while the peak Betti number within the range grows. Figure 3 illustrates the behavior. The second observation is that in the random *geometric* setting, large Betti numbers

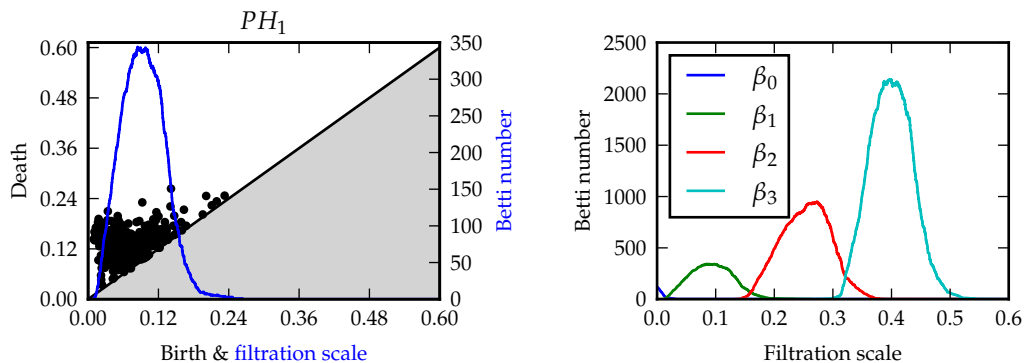


Figure 3: Persistent homology of a realization of an Erdős–Rényi complex with 120 vertices. Increasing homology dimension results in an increase in peak Betti number and a narrowing of the edge probability range where non-vanishing Betti numbers are likely. This behavior is a signature of ER complexes.

¹An (n, p) -Erdős–Rényi random graph has n nodes, and each possible edge appears independently with probability p . We will refer to flag complexes of such graphs as *ER complexes* with parameters n and p .

are much less likely. Intuitively, the triangle inequality puts heavy constraints on where points may be placed in space for the cycles they are a part of not to become bounding at comparatively small filtration scales. Highly persistent non-bounding cycles are thus geometrically fragile and very sensitive to point coordinates, and thus unlikely to occur by chance. As homology dimension increases, the peak Betti number goes down, as one would expect from intuition, since high-dimensional cycles are even more susceptible to being filled in by the triangle inequality. Figure 4 illustrates the behavior.




Figure 4: Persistent homology of a random geometric complex, the Vietoris–Rips complex of 120 points sampled uniformly at random from the unit cube in \mathbb{R}^4 .

The takeaway for us is that the Betti curves of ER complexes behave substantially different from geometric complexes and, of course, from those of complexes arising from structured data. Betti curves can therefore act as an indicator of the absence of meaningful information; contrast Figure 3 and Figure 4. Experiments show that around 100 vertices seems to be more than sufficient for this indicator to be robust.

2.4.1 Quantitatively testing for randomness

If G is a complete weighted graph on V , we write $\text{sh}(G)$ for (a realization of) the complete graph on V having the edge weights of G randomly shuffled. Except in degenerate cases $\text{fl}(\text{sh}(G))$ should be a good realization of an ER random complex. The discussion above then suggests two ad hoc measures for how consistent a flag complex is with an ER random complex by comparing $PH_*(\text{fl}(G))$ and $PH_*(\text{fl}(\text{sh}(G)))$.

For various p and q we define

$$\delta_k(G) = d_{p,q}(PH_k(\text{fl}(G)), PH_k(\text{fl}(\text{sh}(G)))) ,$$

i.e. the persistence module metric between $PH(\text{fl}(G))$ and its shuffled version. We further let β_k and β'_k denote the Betti curve of $PH_k(\text{fl}(G))$ and of $PH_k(\text{fl}(\text{sh}(G)))$, respectively, and define the ratio

$$\Delta_k(G) = \frac{\max_s \beta_k(s)}{\max_s \beta'_k(s)}$$

whenever it exists. $\Delta_k(G)$ thus compares the peak Betti number of $\text{fl}(G)$ to that of its shuffled version. If, for as large a k as is computationally feasible, $\delta_k(G)$ is close to zero and $\Delta_k(G)$ is close to one, we have an indication that $\text{fl}(G)$ resembles an ER random complex.

2.5 Model for neuron activity

We need a model for neuron activity for two reasons. Most importantly, the process of inferring away the contribution that specific covariates have on the neuron activity, as sketched in Section 1 and detailed in Section 3, requires such a model. In addition, we prefer to work with synthetic data so that we are in complete control of all “experimental” parameters when testing the feasibility of our technique.

It has recently been shown [29, 28, 9] that statistical physics’ *kinetic Ising model*, a simple generalized linear model (GLM), does a good job modelling networks of neurons. In its original language, the model governs the discrete time evolution of a set of spins under the influence of both an external driving field and the neighboring spins (as defined either by a discretization of Euclidean space in the original applications, or in general by a weighted graph). In general, the model may be defined as follows.

Definition 10. A set of ± 1 -valued random variables $S_i(k)$, with $1 \leq i \leq N$ and $k = 1, 2, \dots$, are said to obey the (discrete time) *kinetic Ising model* with *couplings* $J \in \mathbb{R}^{N \times N}$ and *external fields* $E_1, \dots, E_n : \mathbb{N} \rightarrow \mathbb{R}$ if their conditional probabilities are

$$P(S_i(k+1) = s_i(k+1) \mid S_1(k) = s_1(k), \dots, S_N(k) = s_N(k)) \\ = \frac{\exp\left(s_i(k+1) \left(E_i(k) + \sum_{j=1}^N J_{i,j} s_j(k)\right)\right)}{2 \cosh\left(E_i(k) + \sum_{j=1}^N J_{i,j} s_j(k)\right)}.$$

For ease of notation, we set $s_1(0) = \dots = s_N(0) = -1$. We refer to the expression $F_i(k) = E_i(k) + \sum_{j=1}^N J_{i,j} s_j(k)$ as the system’s *Hamiltonian* (at time step k).

In the neuroscience interpretation of the model, the probability of neuron i firing or not at time step k , as signified by $S_i(k)$ taking the value $+1$ or -1 respectively, is governed by the external field $E_i(k)$ and by whether or not the neighboring neurons $\{j \mid J_{i,j} \neq 0\}$ fired during time step $k-1$. We will in Section 3 make vital assumptions about the external fields E_1, \dots, E_N .

The function $x \mapsto \exp(x) / \cosh(x)$ is a sigmoidal, and Figure 5 shows the firing probability of a cell (i.e. the probability of $s = 1$ in Definition 10) as a function of the Hamiltonian.




Figure 5: Firing probability as a function of the Hamiltonian in the GLM.

Figure 6 shows the place field of a cell whose firing is generated by the GLM. In the terminology of Definition 10, the external field is in this case predominantly a spatial Gaussian, together with a smaller head direction tuning Gaussian. The cells are randomly, but sparsely, coupled.




Figure 6: Example of the place field of one cell in a population of 100 whose activities are governed by the GLM. Colors have qualitatively the same meaning as in Figure 1.

3 Uncovering hidden information

We suppose that a neuroscientist provides a list of *candidate stimuli*, such as for example spatial preference, head direction preference, theta phase preference, and so forth. We further suppose that the influence of such stimuli can be described by certain functions of distances on some simple manifolds. We will specifically consider

- boxes $(0, 1)^d$ of any dimension d , with the Euclidean metric,
- boxes with any number number of d -disks removed so long as these do not disconnect the manifold,
- circles, with their usual Riemannian metric,
- spheres, with their usual Riemannian metric,
- products of the above, with the product metric.

If the candidate list contains L stimuli with corresponding manifolds M_1, \dots, M_L from the list above, then a complete description of the relevant state of the animal at any time is assumed to be a point in the product $M = M_1 \times \dots \times M_L$. An experiment is then performed wherein the neurons are recorded together with samples of the trajectory $\alpha : \mathbb{R} \rightarrow M$ of the animal through its state space.

The neuronal activity recorded is temporally binned with width δt (in this paper $\delta t \approx 10$ ms). Thus the primary input data to our method consist of a spike train for each neuron, i.e. N binary vectors of the form

$$s_i = (s_i(1), s_i(2), \dots, s_i(T)) \in \{-1, 1\}^T,$$

together with samples of the state space trajectory

$$\alpha(1), \alpha(2), \dots, \alpha(T) \in M.$$

As a concrete example, we might imagine a hypothetical situation where the researcher believes that only spatial position and head direction govern neuron activity. In this case, an experiment might be performed wherein a rat explores a box, while its position and the direction of its head are recorded at the same regular time instances as neuron activity. Supposing that 100

neurons are recorded over the course of 10 minutes², the data we are given then consist of 100 binary spike train vectors of length 60000 together with 60000 samples of the rat state as points in the state space $(0, 1)^2 \times \mathbb{S}^1$. Our goal, as outlined in Section 1.1, is then to determine whether the two suspected stimuli — spatial position and head direction — describe the observed activity, and, most importantly, if not, what the homological properties of any remaining unknown stimuli are.

3.1 Spike train order complex

Having obtained the necessary spike trains, we compute a measure of the degree of cofiring between all pairs of neurons. The Pearson correlation of the corresponding spike trains is a sensible choice, and for $s \in \mathbb{R}^T$ and $s' \in \mathbb{R}^{T'}$ we thus define

$$\text{corr}(s, s') = \frac{\sum_{k=1}^{\min(T, T')} (s(k) - \bar{s})(s'(k) - \bar{s}')}{\sqrt{\sum_{k=1}^{\min(T, T')} (s(k) - \bar{s})^2} \sqrt{\sum_{k=1}^{\min(T, T')} (s'(k) - \bar{s}')^2}}, \quad (2)$$

where

$$\bar{s} = \frac{1}{\min(T, T')} \sum_{k=1}^{\min(T, T')} s(k).$$

To reduce the Pearson correlation's sensitivity to errors in the binning process and slight systematic timing errors in measurements, we will average over a small number of time bins. For a vector $s = (s(1), \dots, s(T)) \in \mathbb{R}^T$, we will write its i 'th left shift as

$$s[i] = (s(1+i), \dots, s(T)) \in \mathbb{R}^{T-i}.$$

Then for $\tau \geq 0$, define

$$\overline{\text{corr}}_{\tau}(s, s') = \frac{1}{\tau + 1} \max \left(\sum_{i=0}^{\tau} \text{corr}(s[i], s'), \sum_{i=0}^{\tau} \text{corr}(s, s'[i]) \right).$$

Finally, to follow the convention that simplicial complexes are filtered in order of *increasing* distance, we define for a set of spike trains s_1, \dots, s_N the *spike train distances*

$$D_{\tau}(s_i, s_j) = \max_{k, l \in \{1, \dots, N\}} \overline{\text{corr}}_{\tau}(s_k, s_l) - \overline{\text{corr}}_{\tau}(s_i, s_j).$$

We will throughout this paper write $D = D_0$ and only briefly discuss specific choices of the averaging time scale τ in Section 4.

While D_{τ} is not a metric when $\tau \neq 0$ (clearly the triangle inequality need not hold), it does provide some semblance of a closeness notion for spike trains, and should be a useful indirect measure of how much place fields intersect. An obvious next step is therefore to take the flag complex of the graph with vertices $1, \dots, N$ corresponding to the neurons, and whose edge weights are $w(i, j) = D(s_i, s_j)$. However, as pointed out in [15], the chemical and biological processes that go into a neuron building its action potential and firing, and the physical processes that constitute the measurement of that event, are of a highly non-linear nature. What we observe, then, is merely a highly transformed version of the real cofiring relationship of neurons.

²This is a realistically sized data set according to computational neuroscience data sharing website <http://crcns.org>.

In other words, if we think of $C(i, j)$ as a true low-level measure of the cofiring relationship of neurons i and j , then $D(s_i, s_j) = f(C(i, j))$ for some unknown function $f : \mathbb{R} \rightarrow \mathbb{R}$. On biological grounds it is reasonable to assume that the only thing that is known, and can ever be known, about f is that it is monotonic.

Following [15], we rescale our data so as not to prescribe meaning to the actual values of the distance D , since it may well be the case that no such meaning exists.

Definition 11. Let $G = (E, V)$ be the complete graph on the vertices $V = \{1, \dots, n\}$ with edge weights w . Let $\varphi : \{1, \dots, \binom{n}{2}\} \rightarrow E$ be a bijection that sorts the weights

$$w(\varphi(1)) < w(\varphi(2)) < \dots$$

(breaking ties arbitrarily), and let \tilde{G} be the complete graph on V with the edge weights

$$\tilde{w}(i, j) = \frac{\varphi^{-1}(i, j)}{\binom{n}{2}}.$$

The *order complex* of G is the flag complex $\text{OC}(G) = \text{fl}(\tilde{G})$.

If we take the order complex of the spike trains and their distances, we discard all information except precisely that which survives the unknown monotone transformation f , namely the ordering of the edge weights.

3.2 Removing the contribution of stimuli to reveal hidden information

The likelihood of data

$$\mathbf{s} = \{s_i(k) \mid 1 \leq i \leq N, 1 \leq k \leq T\}$$

observed under the GLM is

$$L(\mathbf{s}) = \prod_{i=1}^N \prod_{k=1}^{T-1} \frac{\exp\left(s_i(k+1) \left(E_i(k) + \sum_{j=1}^N J_{i,j} s_j(k)\right)\right)}{2 \cosh\left(E_i(k) + \sum_{j=1}^N J_{i,j} s_j(k)\right)}.$$

We now assume that the external field part of the Hamiltonian decomposes into a sum of Gaussians on the various factors of the state space M corresponding to candidate covariates.

As before, $M = M_1 \times \dots \times M_L$, where each M_i is assumed to be one of the manifolds listed earlier (Section 5 discusses a generalization). Projections onto the factors are written $\text{pr}_i : M \rightarrow M_i$, and we denote by d_i the metric on M_i . We define the Gaussians $V_{l,q} : M_l \rightarrow \mathbb{R}$ by

$$V_{l,q}(x) = \exp\left(-\frac{\left(d_l(x, c_{l,q})\right)^2}{2\sigma_{l,q}^2}\right)$$

for $1 \leq l \leq L$ and $1 \leq q \leq Q_l$, and assume that the external part of the Hamiltonian for neuron i at time step k can be written (recall that α is the animal's path through its state space)

$$E_i(k) = \sum_{l=1}^L \sum_{q=1}^{Q_l} A_{i,l,q} \left(V_{l,q} \circ \text{pr}_l \circ \alpha\right)(k)$$

for $c_{l,q} \in M_l$, $A_{i,l,q} \in \mathbb{R}$ and $\sigma_{l,q} \in \mathbb{R}$.

We may use the language of Section 1 if for each i and l there is only one q for which $A_{i,l,q} \neq 0$. Fix an i and an l and let q be the only index for which $A_{i,l,q} \neq 0$. Then $c_{l,q}$ is the center of a place field corresponding to covariate number l , $\sigma_{l,q}$ is a measure of its width, while $A_{i,l,q}$ specifies the peak strength with which it influences the firing of neuron i . In other words, neuron i has a place field specified by $c_{l,q}$ and $\sigma_{l,q}$. When the above relationship between the indices of $A_{\bullet,\bullet,\bullet}$ does not hold, we allow each covariate to govern place cell activity through a linear combination of Gaussian fields, which will be necessary for the inference process described next.

The log-likelihood of the observed data is

$$\log(L(\mathbf{s})) = \sum_{i=1}^N \sum_{k=1}^{T-1} \left[s_i(k+1) \left(\sum_{l=1}^L \sum_{q=1}^{Q_l} A_{i,l,q} \left(V_{l,q} \circ \text{pr}_l \circ \alpha \right) (k) + \sum_{j=1}^N J_{i,j} s_j(k) \right) - \log \left(2 \cosh \left(\sum_{l=1}^L \sum_{q=1}^{Q_l} A_{i,l,q} \left(V_{l,q} \circ \text{pr}_l \circ \alpha \right) (k) + \sum_{j=1}^N J_{i,j} s_j(k) \right) \right) \right], \quad (3)$$

and it is an easy calculus exercise to show that

$$\begin{cases} \mathbb{R}^{N \sum_{l=1}^L Q_l} \rightarrow \mathbb{R} \\ (A_{\bullet,\bullet,\bullet}, J_{\bullet,\bullet}) \mapsto \log(L(\mathbf{s})) \end{cases}$$

is a convex function. We can therefore perform likelihood maximization by means of convex optimization to infer the $A_{i,l,q}$'s and $J_{i,j}$'s that best fit the observed data. With these coefficients in hand, we can selectively remove the (expected) contribution of each stimulus on the candidate list. The residual data after removal no longer consists of binary spike trains, but instead real-valued time series. Low and high values correspond to improbable and probable firing events, respectively, while values near zero reflect lack of knowledge of the probability of either outcome.

Throughout this paper all the inference will be performed with $25^2 = 625$ spatial basis functions, and 25 circular basis functions, both with means uniformly distributed. For example, if the state space is $M = M_1 \times M_2$ with $M_1 = (0,1)^2$ and $M_2 = S^1$, then $Q_1 = 25^2$ and $Q_2 = 25$, and with the $c_{1,q}$'s forming a regular grid on $(0,1)^2$ and the $c_{2,q}$'s uniformly distributed on S^1 .

The assumptions on the external fields above thus allow us to remove the contributions of specific stimuli towards the firing activity, suggesting the following algorithm:

1. Perform an experiment as described above, yielding spike trains s_1, \dots, s_N and samples of the state space path α .
2. Compute the spike train distances $D(s_i, s_j)$ for $1 \leq i, j \leq N$, and let them be the edge weights on a complete graph G on N vertices.
3. Compute $PH_*(OC(G))$.
4. Is $PH_*(OC(G))$ consistent with the persistent homology of an ER random complex per Section 2.4.1? If so, go to step 4a. If not, go to step 5.
 - (a) There is nothing more to learn from our method. We are done.
5. Is the list of candidate stimuli exhausted? If so, go to step 5a. If not, go to step 6.
 - (a) $PH_*(OC(G))$ may reveal information about the homology of the state space corresponding to unknown covariate(s). We are done.

6. Using the inference process described above, remove the contributions of a selected stimulus. This yields real-valued time-series that from now on replace the spike trains s_1, \dots, s_N . Go to step 2.

Figure 7 summarizes the above procedure.

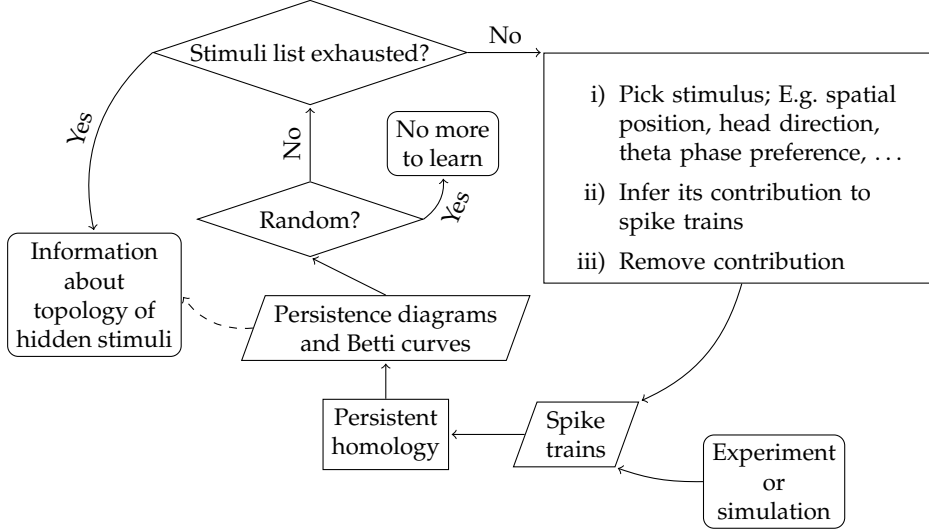


Figure 7: Our main contribution summarized.

As a concrete illustration, we return to the example from earlier: assume that place cell activity is in fact governed only by spatial position and head direction, but suppose that head direction tuning is unknown to the researchers. The candidate covariate list therefore consists only of spatial position. For an experiment conducted with the animal exploring a box, the state space is $(0, 1)^2$. After having performed steps 2 through 6 once, the second iteration leads us to step 5a; there are no more stimuli to account for, yet we see persistent homology inconsistent with random data. Examination of the persistence diagrams of $PH_*(OC(G))$ reveals that the unknown covariates correspond to a state space with the same homology as a circle. Together with other evidence, this may lead the researchers to suspect head direction tuning as a hidden covariate. New experiments may then be performed to investigate this, and our method may be applied also to the new data.

We finally point out that it is *not* essential to our method that we necessarily capture the correct homology of the full state space. It is the state space after the removal of known stimuli that matters (and even then, the weaker information of whether the remaining state space is homologically trivial or not may be useful).

4 Computational results

We test the efficacy of our method using synthetically generated data in order to be in complete control of all “experimental parameters”, and because publicly available real data often come from experiments without a topological focus (for example the spatial environment tends to be homologically trivial). For the results presented here, neuron activity was simulated from the same GLM as used for the inference process by appropriate selection of peak field strength coefficients $A_{\bullet,\bullet,\bullet}$, centers $c_{\bullet,\bullet}$ and widths $\sigma_{\bullet,\bullet}$. In addition, a constant negative term (typically

–1) was added to the external fields to make the overall firing rate low outside of place fields, as is the case for many real cells. This term essentially just lowers the noise floor of our data, and should be of no deeper significance to us.

Even if one of the experiments below does not call for spatial or head direction tuning, the state space will always consist of at least a factor for the physical environment and a factor \mathbb{S}^1 for the head direction. If the physical environment is denoted B , then the $B \times \mathbb{S}^1$ factor of the state space is explored as follows: If at some time step the animal’s state is $(x, y, \theta) \in B \times \mathbb{S}^1$, then the next state is found by choosing a θ' randomly and uniformly within 0.02 of θ in \mathbb{S}^1 . If $(x + 5 \cdot 10^{-4} \cos \theta', y + 5 \cdot 10^{-4} \sin \theta', \theta')$ is within $B \times \mathbb{S}^1$, this point is the next state. If not, new angles are drawn until the new state is valid. We note that qualitatively similar results are obtained if this slightly realistic random walk is replaced by ordinary Brownian motion or uniform random sampling.

Except in Section 4.6, we do not average the correlation measure, i.e. we work with the spike train distance D_0 .

Some additional computational experiments have been relegated to Appendix A.

4.1 Recovering spatial homology

While the central point of this paper is the uncovering of the homological properties of unknown stimuli underlying neural activity, we begin by providing an example of how persistent homology of the order complex of spike train distances can recover the homology of a spatial environment. This is analogous to some of the results presented in [8].

State space is now the unit square punctured by four disks of radius 0.15. The disks are centered at $(0.27, 0.27)$, $(0.27, 0.72)$, $(0.72, 0.27)$ and $(0.72, 0.72)$. We refer to the punctured box as B below. For technical reasons, data were generated with the state space having an extra circle factor, but every cell had its activity coefficients corresponding to this factor set to zero. Thus *only spatial tuning* is present in this example.

In the notation of Section 3.2, $N = Q_1 = Q_2 = 100$, $M_1 = B$, $M_2 = \mathbb{S}^1$, $L = 2$, and $J_{i,j} = 0$ for all $1 \leq i, j \leq N$. The $c_{1,q}$ ’s form a regular grid on B , and

$$A_{i,1,q} = \begin{cases} 2 & \text{if } i = q \\ 0 & \text{otherwise,} \end{cases}$$

while $A_{i,2,q} = 0$ for all i, q . Figure 8a shows the spatial tuning of a single neuron.

Figure 8b shows that we correctly recover the homology of B , the only part of M detected by the neural activity. In other words, this example illustrates how we can detect and even count the obstructions in a space from neuron activity alone.

4.2 Proof of concept

The simplest possible setting wherein our inference scheme, laid out in Section 3.2, is useful, is perhaps one where both spatial and head direction tuning govern neuron activity, but where the researcher believes only one of those to be real.

For this computation, the spatial component of the state space is a unit square punctured in its center by a single disk of radius 0.2, denoted by B . The head direction component is \mathbb{S}^1 .

In the notation of Section 3.2, $N = Q_1 = Q_2 = 100$, $M_1 = B$, $M_2 = \mathbb{S}^1$, $L = 2$, and $J_{i,j} = 0$ for




Figure 8: Results of the experiment in Section 4.1. **a:** The activity of a single cell as a function of spatial position. **b:** Persistent homology of the order complex. All higher dimensions of homology are trivial or close to trivial.

all $1 \leq i, j \leq N$. The $c_{1,q}$'s form a regular grid on B , and

$$A_{i,1,q} = \begin{cases} 2 & \text{if } i = q \\ 0 & \text{otherwise,} \end{cases}$$

Similarly, the $c_{2,q}$'s are uniformly spread out over S^1 . To avoid artificially coupling head direction and spatial tuning through the ordering of their place field centers, we let σ be a random permutation of $\{1, \dots, N\}$ and then let

$$A_{i,2,q} = \begin{cases} 2 & \text{if } i = \sigma(q) \\ 0 & \text{otherwise,} \end{cases}$$

for all i, q .

We reiterate that we assume that the researcher is *unaware of head direction tuning* as a real influence on place cell activity in this example; he believes spatial position is the only relevant stimulus. After conducting an experiment, the researcher sees the neurons' spatial dependence exemplified in Figure 9a. The head direction dependence in Figure 9b is *not* known to the researcher.

Persistent homology of the order complex of the correlations observed can be seen in Figure 10. Note that we do *not* observe the actual homology of the state space M , which is a thickened torus. This is not entirely satisfactory, but at least the observed persistence diagrams indicate there is homologically nontrivial information present in the neuron activity.

Satisfied that the persistence diagrams are consistent with his hypothesis about the relevant covariates, the researcher proceeds to the next step, namely removing the effect of the spatial covariate (the only one he is aware of). We maximize the expression in equation (3) in terms of the $A_{i,l,1}$'s using L-BFGS-B³, and subtract from the observed spike trains the expected activities predicted from only spatial influence.

The result of the removal on spatial activity tuning can be seen in Figure 11. It is as expected, and obviously does not provide new information to the researcher. The persistence diagram in

³As provided by SciPy.




Figure 9: Activity of a single neuron in the experiment from Section 4.2. **a:** Spatial dependence. This is the only view of activity visible to the researcher. **b:** Head direction dependence. The researcher is *not* privy to this information, as he is unaware that head direction preference is real.




Figure 10: Observed persistent homology of the order complex of the correlations for the experiment from Section 4.2.




Figure 11: Average activity of a single neuron in the experiment from Section 4.2 *after removal of spatial tuning*. **a:** Spatial dependence. This is the only view of activity visible to the researcher. **b:** Head direction dependence. The researcher is *not* privy to this information, as he is unaware that head direction preference is real.




Figure 12: Observed persistent homology of the order complex of the correlations for the experiment from Section 4.2 *after removal of spatial tuning*.

Figure 12, however, shows that there is still homologically non-trivial information contained in the observed data. This should hopefully lead the researcher to suspect that there are further, hidden, influences on neuron activity, and, most importantly, that this/these *are of a circular nature*.

Guided by this, the researcher might consider head direction tuning. He therefore sets up a new experiment where also this is tracked, so that also its influence may be removed from the data. Doing so results in the activity plot in Figure 13.




Figure 13: Average activity of a single neuron in the experiment from Section 4.2 *after removal of both spatial and head direction tuning*. **a:** Spatial dependence. **b:** Head direction dependence.

The persistence diagrams in Figure 14 now lack clear suggestions of non-trivial homology. Moreover, the Betti curves are highly indicative of an ER random complex (compare Figure 3). To verify this last claim, we also permuted the neuron correlations randomly; Figure 15 shows the result. This should serve as a firm indication that all stimuli underlying the neuron activity in the data have been accounted for.

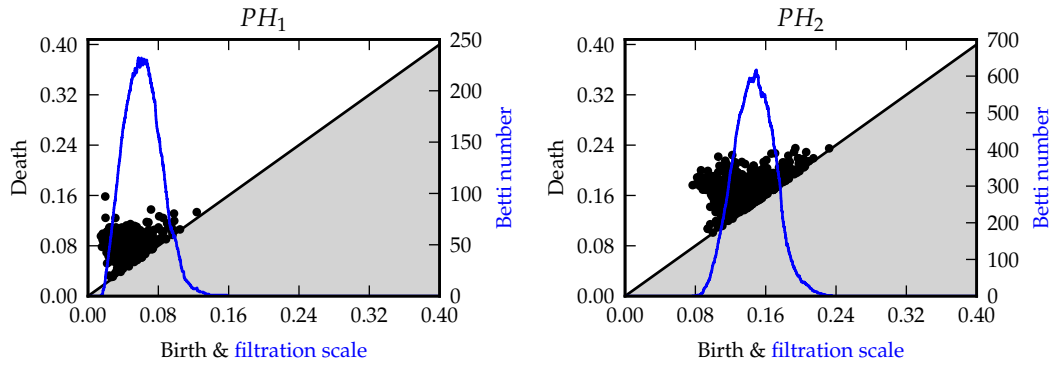


Figure 14: Observed persistent homology of the order complex of the correlations for the experiment from Section 4.2 after removal of both spatial and head direction tuning.

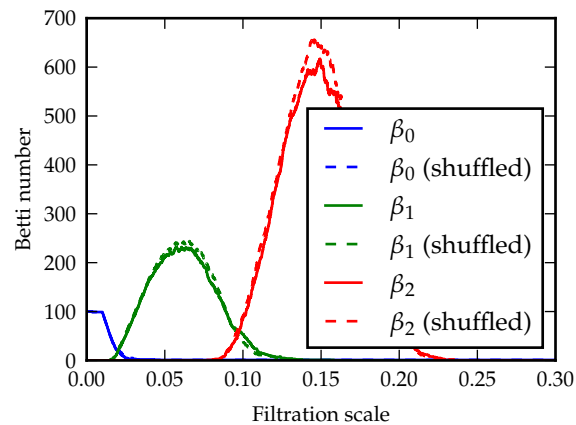


Figure 15: After all relevant covariates have been removed from the data in the experiment from Section 4.2, the Betti curves of the order complex of the correlations are consistent with those of an ER random complex (here formed by randomly shuffling the correlations). C.f. Figure 3.

4.2.1 Alternative scenario

One may also want to consider an alternative hypothetical scenario wherein head direction tuning is the only suspected stimulus. For readability reasons we do not include that scenario in full here. The interesting part is the persistent homology after the removal of head direction tuning. Figure 16 shows that we recover the correct PH_1 also in this case.

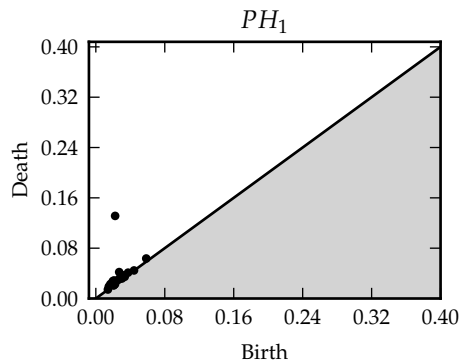


Figure 16: Persistent homology of the order complex after removing the effect if head direction tuning in an alternative version of the experiment in Section 4.2 (see Section 4.2.1).

4.3 Effect of couplings

In the preceding experiments, cells were never coupled. To illustrate that our method also copes with such “internal stimuli”, we repeated the experiment from Section 4.2 with the change that every cell is given a weak but random coupling to a every other cell. Specifically, we kept all simulation parameters as before, but let every $J_{i,j}$ be drawn independently and uniformly from $[-0.1, 0.1]$. The couplings are thus weak compared to the external stimuli (which peak at 2 in the centers of fields), but should nevertheless introduce noise to the data.

Figure 17 shows the results before any covariate removal. Observe that the random couplings introduce significant noise in the spatial dependence of the activity compared to that in Figure 9. Figure 18 shows that we are able to carry out the same procedure as in Section 4.2 also in the presence of couplings.

4.4 Effect of theta wave phase preference

We simplistically model theta wave phase preference as each neuron preferentially firing near a randomly chosen phase of a 7 Hz sinusoidal wave in time. The experimental parameters are the same as in Section 4.2, except now $L = 3$, and the state space gains an extra factor $M_3 = S^1$. The $c_{3,q}$'s are uniformly spread out over S^1 , and for a random permutation⁴ σ of $\{1, \dots, N\}$ the field strength coefficients are

$$A_{i,3,q} = \begin{cases} 2 & \text{if } i = \sigma(q) \\ 0 & \text{otherwise,} \end{cases}$$

⁴Present for the same reason as for the head direction tuning in Section 4.2.

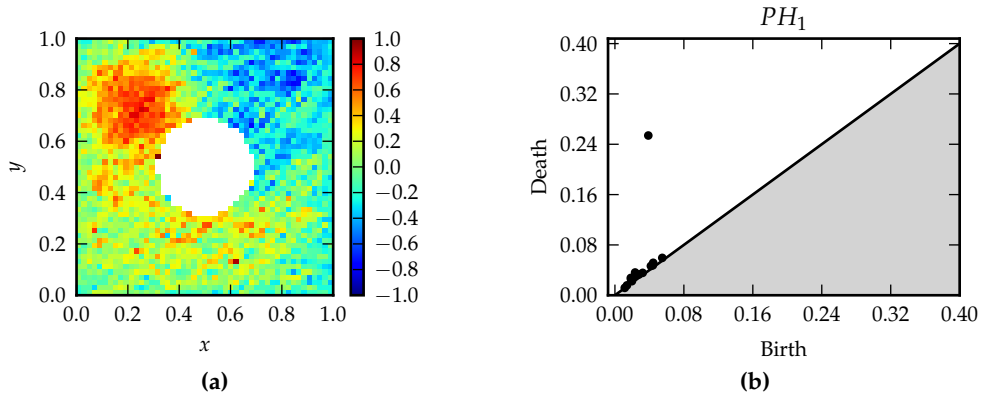


Figure 17: Prior to removal of any covariates in the experiment from Section 4.3. **a:** Spatial dependence for the activity of a single neuron. **b:** PH_1 persistence diagram.

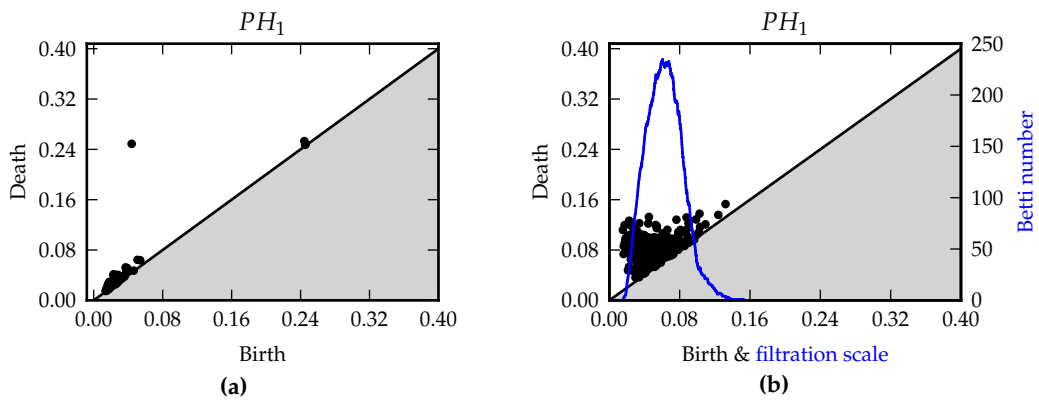


Figure 18: Removing external and internal covariates in the experiment from Section 4.3. **a:** Persistence diagram after spatial dependence and internal couplings are removed. **b:** Persistence diagram and Betti curve after all external (space and head) and internal (couplings) influences are removed.

for all i, q .

Theta phase preference is thus, as far as topology is concerned, precisely the same as head direction preference, except that the M_3 factor of state space is explored by always moving forward in time (modulo $1/14$) instead of by a random walk. We therefore expect that theta phase preference will contribute to homology in the same way as head direction tuning. Figure 19 confirms this. Again it should be pointed out that we are not observing homology consistent with the three dimensional torus $S^1 \times S^1 \times S^1$. While this may seem unsatisfactory, it is a quite natural effect of one covariate suppressing the expression of the homology of the others, i.e. one of the radii of the torus dominating over the others. This illustrates well why the inference process and sequential removing of covariates really is necessary; the fact that there are *three* circular factors in the state space cannot be glanced directly from the observed data.




Figure 19: Observations from the experiment in Section 4.4 before removal of any covariates. **a:** Activity of a single neuron as a function of theta wave phase. **b:** Persistent homology of the order complex of the spike train distances.

Figure 20 shows that we obtain the expected results when left with only theta phase preference and when all covariates are removed.

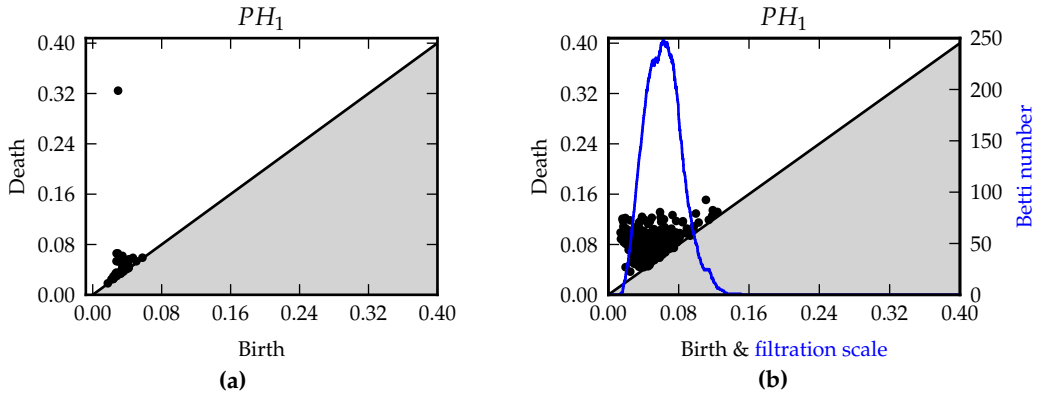


Figure 20: Persistent homology of the order complex of the residual spike trains in the experiment in Section 4.4 after removing some external covariates. **a:** After removing spatial and head direction tuning. **b:** After removing all covariates, including theta phase preference.

4.5 Sensitivity to field strength variations

The preceding experiments illustrate the efficacy of our technique for fixed choices of peak field strengths. We now wish to demonstrate that our method is robust over a wide range of field strengths. To this end, we repeat the experiment from Section 4.2, but now for varying choices of peak strength for both the spatial and head direction fields.

For conducting this experiment, we need a numerical measure of the outcome that is more succinct than a persistence diagram or a Betti curve. Since we generate data with head direction and spatial tuning enabled, we hope to see the same outcome as in the experiment from Section 4.2:

- Before removal of any covariates, we should see a prominent PH_1 generator, i.e. one that stands out in the persistence diagram by having a longer lifetime than most. At the same time, the persistence diagrams and their associated Betti curves should be inconsistent with those of ER random complexes.
- This should remain the case when inferring away one of either head direction or spatial tuning.
- After inferring away both head direction and spatial tuning, the persistence diagrams and Betti curves should be highly consistent with those of ER random graphs.

As a very crude measure of the presence of “highly persistent” or “prominent” PH_1 generators, we simply consider the ratio of the lifetime of the most persistent one to the lifetime of the second-most persistent one. We refer to the ratio as ρ_1 throughout this section. Note that our aim is that ρ_1 measures how easy a single prominent generator is to distinguish from noise. In the case of multiple highly persistent PH_1 generators, such as for the persistent homology of a torus, it reports a false negative. However, since the experiment in Section 4.2 failed to capture this true homology of the state space, we feel confident this measure is sufficient. We also verified this by manual inspection of several of the persistence diagrams computed.

To compare with random ER complexes, we use both δ_1 and Δ_1 . Recall that the closer $\delta_1(G)$ is to zero and $\Delta_1(G)$ is to one, the more consistent the flag complex built on the data in G is with an ER complex.

Figure 21 summarizes the relevant measures of success. Observe that we meet our criteria for expected outcome for parameters that lie outside the blue band in Figure 21a and that at the same time are not too weak in either the spatial or the head direction tuning strength (see e.g. Figure 21g). The latter requirement is no surprise, as we with weak field strengths observe neurons that barely respond to external stimuli. Further investigation reveals that as the head direction field strength is made weaker (i.e. as we approach the band of failure from above in Figure 21a), the most persistent PH_1 lives for a shorter and shorter time, until it becomes indistinguishable from the noise near the diagonal. For even weaker head direction field strengths, i.e. below the band, we are essentially in the domain of activity governed entirely by the spatial fields.

Note also that when we in Figure 21a are within the blue band, Figures 21d and 21g show that the observed data are *not* consistent with an ER complex. Inferring away the spatial influence is therefore something one may still wish to do, whereupon one uncovers homology strongly consistent with a circle everywhere except for with weak head direction tuning, as seen in Figure 21b.




Figure 21: Summary of the behavior of our method over varying field strengths, as detailed in Section 4.5. In all the plots, the horizontal axis denotes the peak spatial field strength and the vertical axis denotes the peak head direction field strength. Arrows in the color bar indicate that the colors are clipped above or below a value. For a spike train distance graph G obtained using the given peak field strengths, the **first row** shows $\rho_1(G)$, the **second row** shows $\Delta_1(G)$, and the **third row** shows $\delta_1(G)$. In the **first column** G comes from the original observations, in the **second column** the spatial influence is removed, and in the **third column** both spatial and head direction influence are removed.

4.6 Homology from internal couplings

In the preceding experiments, internal couplings have been absent, or, as in Section 4.3, not themselves been the focus of our attention. We now illustrate that our method is also capable of detecting the homology of (the flag complex of) the graph defining the neighbor relations of the neurons.

We generate data with only spatial fields and internal couplings. Specifically, $N = Q_1 = 100$, $M = M_1 = (0, 1)^2$, the $c'_{1,q}$ s form a regular grid on $(0, 1)^2$, and the peak field strengths are

$$A_{i,1,q} = \begin{cases} 1 & \text{if } i = q \\ 0 & \text{otherwise} \end{cases}$$

(for technical reasons, just as in Section 4.2, the explored state space still contains an extra S^1 factor, but the corresponding field strengths are set to zero). The symmetric matrix J describes a circle on all N nodes with edges in both directions with weights 2. The indices defining the edges of the circle are chosen randomly to avoid an unnatural coupling to the spatial fields through the ordering.




Figure 22: Persistent homology of the order complex of the observed spike train distances D_τ in the experiment from Section 4.6. **a:** $\tau = 1$. **b:** $\tau = 10$.

Figure 22a shows an unexpected result consistent with (at least) *two* circles. Manual inspection of the spike train distances reveals that this is an artifact of the GLM being coupled across only two consecutive time steps. When we use a spike train distance without temporal average (i.e. D_0), we are unable to resolve any coupling interactions across an even number of neurons. The 1-skeleton of the order complex of the spike train distances thus breaks into *two* circles, corresponding to the even and odd parity edges of the neighborhood graph $(\{1, \dots, N\}, \{(i, j) \mid J_{i,j} \neq 0 \text{ or } J_{j,i} \neq 0\})$. This undesired behavior vanishes when the correlation time average is greater, for example $\tau = 10$, as is shown in Figure 22b.

5 Discussion and sketches of a general framework

We believe that the core aspects of the method presented in this paper are applicable outside of neuroscience. In a general setting, we imagine a point cloud $C = \{p_1, \dots, p_N\}$ of points on a manifold M . The exact assumptions on M have not been worked out, but we believe that a

compact, connected, homogenous Riemannian manifold without boundary suffices.⁵ Write the metric of M as d . The points themselves, and their pairwise distances, are a priori unknown. Estimates of the latter are obtained by performing a random walk on M while observing N Bernoulli processes S_1, \dots, S_N (producing time series s_1, \dots, s_N). If we denote by $X(k)$ the random variable of the random walk's position at time step k , the vital assumption is that the parameter (the "success" probability) of Bernoulli process S_i at time step k is

$$P(S_i(k) = 1 \mid X(k) = x) = f(d(x, p_i)) \quad (4)$$

for some monotonically decreasing sufficiently integrable (unknown) function $f : \mathbb{R}^+ \rightarrow [0, 1]$. If the random walk has progressed long enough that the distribution of $X(k)$ is close to uniform, then we are in a setting where we believe our methods are applicable.

The question to ask is whether the persistent homology of the flag complex (or order complex) of the graph with edge weights $D_\tau(s_i, s_j)$ for $1 \leq i, j \leq N$ closely approximates the persistent homology of $\text{VR}(C)$.

In the above setup, define $g_k : M \times M \rightarrow \mathbb{R}$ by

$$g_k(y, z) = \int_M f(d(y, x))f(d(z, x)) \, d\rho_k(x)$$

with ρ_k the probability density function for the random walk at time step k . The function g_k arises naturally as the probability

$$P(S_i(k) = 1 \cap S_j(k) = 1) = g_k(p_i, p_j)$$

and is the essential part of the estimated Pearson correlation $\text{corr}(s_i, s_j)$ from equation (2). In the case that $X(k)$ is fully uniform, i.e. $d\rho_k = dx$, proving correctness of our recovered persistent homology amounts to showing the existence of a monotonically decreasing $h : \mathbb{R}^+ \rightarrow \mathbb{R}$ that makes

$$\begin{array}{ccc} M \times M & \xrightarrow{g} & \mathbb{R} \\ d \downarrow & \nearrow h & \\ \mathbb{R}^+ & & \end{array}$$

commute. In the Euclidean situation ($M = \mathbb{R}^n$ and dropping the assumptions on M) this is an easy calculus exercise, but the details need to be worked out for the more general situations.

The actual persistent homology of (say, the Vietoris–Rips complex of) C , i.e. the level of question Q0, may be of interest in applications outside of neuroscience. While there are perhaps not many situations in which observations of the Bernoulli processes are available, but the values on the right hand side of equation (4) are *not*, one example might be sensor network coverage problems [14] with sensors of severely reduced capabilities. Instead of being able to sense the strength of their neighbors, we imagine that the sensors carry simple devices which trigger with a frequency monotonically related to the distances to other sensors.

Work should be done to give rigorous bounds on the persistence modules based on the statistical properties of the Pearson correlations $\text{corr}(s_i, s_j)$, especially in the settings when the random walk distribution is not yet truly uniform.

At the level of questions Q1 and Q2 we encounter other systems described by a GLM (such as in [3]) or, in principle, by other models that accommodate an inference process like that in Section 3.2. Our method may be able to shed light on the homological properties of parts of the

⁵Some of the state spaces considered in Section 4 are of course not covered by these assumptions. Numerical evidence still strongly suggests that our method works well in practice also in the situations considered.

external influences on such a system, or, as in the experiment in Section 4.6, unknown internal couplings.

Future applications of our method on actual data should demonstrate its usefulness in identifying hidden topological information in data, both inside and outside of neuroscience.

6 Acknowledgments

This paper is a product of a cooperative project where also Yasser Roudi (Kavli Institute for Systems Neuroscience, Norwegian University of Science and Technology) takes part.

GS would like to thank Geir-Arne Fuglstad for valuable discussions on matters of probability.

All persistent homology computations were performed using Nanda's *Perseus* software [25].

A Additional computational results

Some additional computational results have been relegated here for the sake of readability.

A.1 Comparison with multi-dimensional scaling

One may well ask whether the use of persistent homology actually contributes anything useful to our proposed method. To investigate this, we considered whether multi-dimensional scaling (MDS) could successfully, i.e. near-isometrically, embed the "distances" from the various stages of the experiment in Section 4.2 in low-dimensional Euclidean space.

Figure 23 shows that MDS provides the same useful information as persistent homology before removal of covariates and after removal of only spatial tuning. After removing head direction tuning, however, MDS is unable to detect the annular nature of the remaining spatial covariate (compare in particular Figure 23c and Figure 16). This illustrates the essential role played by persistent homology in our work.

We point out that embedding in \mathbb{R}^3 does not seem to qualitatively improve the situation.

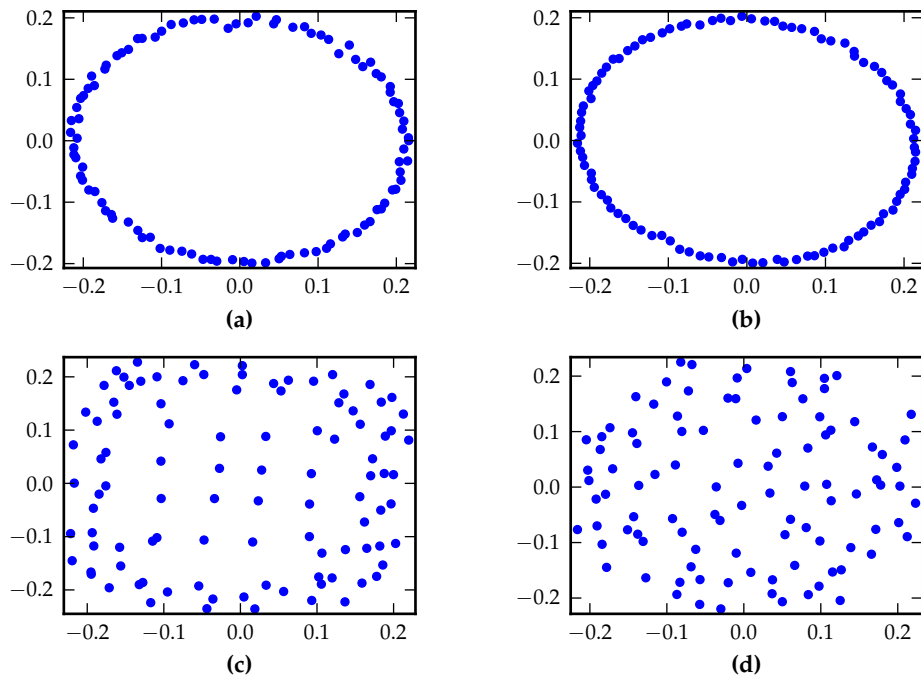


Figure 23: MDS embedding of the neuron distances from the experiment in Section 4.2. **a:** Original data. **b:** After removal of spatial tuning. **c:** After removal of head direction tuning. **d:** After removal of both covariates.

References

- [1] Mamiko Arai, Vicky Brandt, and Yuri Dabaghian. “The Effects of Theta Precession on Spatial Learning and Simplicial Complex Dynamics in a Topological Model of the Hippocampal Spatial Map”. In: *PLoS computational biology* 10.6 (2014).
- [2] Gorô Azumaya. “Corrections and supplementaries to my paper concerning Krull–Remak–Schmidt’s theorem”. In: *Nagoya Mathematical Journal* 1 (1950), pp. 117–124.
- [3] Stanislav Borysov, Yasser Roudi, and Alexander Balatsky. “US stock market interaction network as learned by the Boltzmann Machine”. In: *ArXiv e-prints* (2015). arXiv:1504.02280 [q-fin.ST].
- [4] Frédéric Chazal and Steve Yann Oudot. “Towards persistence-based reconstruction in Euclidean spaces”. In: *Proceedings of the twenty-fourth annual symposium on Computational geometry*. ACM, 2008, pp. 232–241.
- [5] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. “Stability of persistence diagrams”. In: *Discrete & Computational Geometry* 37.1 (2007), pp. 103–120.
- [6] William Crawley-Boevey. “Decomposition of pointwise finite-dimensional persistence modules”. In: *Journal of Algebra and Its Applications* 14.5 (2015).
- [7] Carina Curto and Vladimir Itskov. “Cell groups reveal structure of stimulus space”. In: *PLoS Computational Biology* 4.10 (2008). DOI: 10.1371/journal.pcbi.1000205.
- [8] Yuri Dabaghian, Facundo Mémoli, L. Frank, and Gunnar Carlsson. “A Topological Paradigm for Hippocampal Spatial Map Formation Using Persistent Homology”. In: *PLoS computational biology* 8.8 (2012).
- [9] Benjamin Dunn, Maria Mørreaunet, and Yasser Roudi. “Correlations and functional connections in a population of grid cells”. In: *PLoS computational biology* 11.2 (2015).
- [10] Herbert Edelsbrunner. *A Short Course in Computational Geometry and Topology*. Springer, 2014.
- [11] Herbert Edelsbrunner and John Harer. *Computational Topology: An Introduction*. American Mathematical Society, 2010. ISBN: 978-0-8218-4925-5.
- [12] Robert Ghrist. “Barcodes: the persistent topology of data”. In: *Bulletin of the American Mathematical Society* 45.1 (2008), pp. 61–75.
- [13] Robert Ghrist. *Elementary Applied Topology*. 2014.
- [14] Robert Ghrist and Abubakr Muhammad. “Coverage and hole-detection in sensor networks via homology”. In: *Proceedings of the 4th international symposium on Information processing in sensor networks*. IEEE Press, 2005.
- [15] Chad Giusti, Eva Pastalkova, Carina Curto, and Vladimir Itskov. “Clique topology reveals intrinsic geometric structure in neural correlations”. In: *Proceedings of the National Academy of Sciences of the United States of America* (2015). DOI: 10.1073/pnas.1506407112.
- [16] Allen Hatcher. *Algebraic Topology*. Cambridge University Press, 2002. ISBN: 0521-79540-0.
- [17] Jakob Jonsson. *Simplicial Complexes of Graphs*. Vol. 1928. Lecture Notes in Mathematics. Springer, 2008.
- [18] Matthew Kahle. “Topology of random simplicial complexes: a survey”. In: *Contemporary Mathematics* 620 (2014), pp. 201–222.
- [19] Harold Kuhn. “The Hungarian method for the assignment problem”. In: *Naval research logistics quarterly* 2.1-2 (1955), pp. 83–97.

- [20] Peter McCullagh and John A. Nelder. *Generalized linear models*. Vol. 37. Monographs on Statistics and Applied Probability. CRC press, 1989.
- [21] Kenji Mizuseki, Anton Sirota, Eva Pastalkova, and György Buzsáki. “Theta oscillations provide temporal windows for local circuit computation in the entorhinal-hippocampal loop”. In: *Neuron* 64.2 (2009), pp. 267–280.
- [22] Kenji Mizuseki, Anton Sirota, Eva Pastalkova, Kamran Diba, and György Buzsáki. *Multiple single unit recordings from different rat hippocampal and entorhinal regions while the animals were performing multiple behavioral tasks*. CRCNS.org. DOI: 10.6080/K09G5JRZ.
- [23] Robert Muller. “A quarter of a century of place cells”. In: *Neuron* 17.5 (1996), pp. 813–822. DOI: 10.1016/S0896-6273(00)80214-7.
- [24] James Munkres. “Algorithms for the assignment and transportation problems”. In: *Journal of the Society for Industrial & Applied Mathematics* 5.1 (1957), pp. 32–38.
- [25] Vidit Nanda. *Perseus*. URL: <http://www.sas.upenn.edu/~vnanda/perseus>.
- [26] John O’Keefe and Jonathan Dostrovsky. “The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat”. In: *Brain research* 34.1 (1971), pp. 171–175.
- [27] Jonathan W. Pillow et al. “Spatio-temporal correlations and visual signalling in a complete neuronal population”. In: *Nature* 454.7207 (2008), pp. 995–999.
- [28] Yasser Roudi, Benjamin Dunn, and John Hertz. “Multi-neuronal activity and functional connectivity in cell assemblies”. In: *Current opinion in neurobiology* 32 (2015), pp. 38–44.
- [29] Yasser Roudi and John Hertz. “Mean field theory for nonequilibrium network reconstruction”. In: *Physical review letters* 106.4 (2011).
- [30] Alon Rubin, Michael Yartsev, and Nachum Ulanovsky. “Encoding of head direction by hippocampal place cells in bats”. In: *The Journal of Neuroscience* 34.3 (2014), pp. 1067–1080.
- [31] Etienne Save, Ludek Nerad, and Bruno Poucet. “Contribution of multiple sensory information to place field stability in hippocampal place cells”. In: *Hippocampus* 10.1 (2000), pp. 64–76.
- [32] Edward C. Tolman. “Cognitive maps in rats and men”. In: *Psychological review* 55.4 (1948), pp. 189–208.
- [33] Wilson Truccolo, Uri T. Eden, Matthew R. Fellows, John P. Donoghue, and Emery N. Brown. “A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects”. In: *Journal of neurophysiology* 93.2 (2005), pp. 1074–1089.
- [34] Eugene Wigner. “On the distribution of the roots of certain symmetric matrices”. In: *Annals of Mathematics* (1958), pp. 325–327.