

Topological Structure of Population Activity in Primary Visual Cortex

Gurjeet Singh, Facundo Mémoli, Tigran Ishkhanov,
Gunnar Carlsson, Guillermo Sapiro and Dario Ringach

November 8, 2007

Abstract

Information in the cortex is widely believed to be represented by the joint activity of neuronal populations. Developing insights into the nature of these representations is a necessary first step in our quest to understanding cortical computation. Here, we show that fundamental questions about neural representation can be cast in terms of the topological structure of population activity. A new method, based on the concept of persistent homology, is introduced and first validated on artificial datasets. The technique is then applied to study the topological structure of neural activity in cell populations of primary visual cortex that were either spontaneously active or driven by natural image sequences. Our analyses confirm that spontaneous activity is highly structured and statistically different from noise. Furthermore, the topological objects derived from spontaneous and driven activity have similar distributions which are dominated by the topology of a circle and the two-sphere. This latter structure, we postulate, corresponds to the representation of orientation and spatial frequency on a spherical surface. Our findings shed new light on the relationship between ongoing and driven activity in primary visual cortex and demonstrates, for the first time, that computational topology offers novel tools to tackle fundamental questions about the representation of information in the nervous system.

Contents

1	Introduction	3
2	Tools from Algebraic Topology	4
3	Results	10
3.1	Validation of the method	10
3.2	Validation dataset # 1: a circle	10
3.3	Validation dataset # 2: a torus	13
4	Empirical data: the topology of spontaneous and driven activity in V1	16
5	Discussion	16
6	Material and Methods	20
6.1	Animal Preparation	20
6.2	Electrophysiology	21
6.3	Visual Stimuli	21
6.4	Creating point clouds	21
6.5	Statistical Significance	22
7	Acknowledgements	22
A	Simplicial Homology	22
A.1	Induced homomorphisms	24
A.2	Building simplicial complexes from the data	24
A.3	Persistent Homology: barcodes	25

1 Introduction

Recent results suggest that spontaneous cortical activity wanders among intrinsic attractor states, and that ongoing activity plays an important role in modulating the responses to an external stimulus (Arieli et al., 1995; Cossart et al., 2003; Ikegaya et al., 2004; Kenet et al., 1997, 1999; Kenet et al., 2003; Lampl et al., 1999; MacLean et al., 2005; Ringach, 2003; Tsodyks et al., 1999). In particular, Kenet et al (2003) reported that population activity in primary visual cortex, as measured with voltage sensitive dyes (Grinvald and Hildesheim, 2004), appears to switch dynamically among states, some of them, corresponding to the orientation maps obtained during stimulation with oriented patterns. After observing a resemblance of spontaneous activity states to those of orientation maps, these researchers confirmed, via post-hoc analysis, that the correlation between the orientation maps and spontaneous activity was beyond what would be expected by chance. However, for about 80% of the time, the cortex spent time in states that were uncorrelated with the orientation states. This result could be caused by a 'background state' of unstructured activity or as structured activity that is not captured by the orientation maps. The methods used in Kenet et al (2003) are not suited for discriminating between these possibilities.

A more principled approach for analyzing the structure of population activity was introduced in a theoretical study by Goldberg et al (2004). These investigators studied the possibility of using of a single real-valued statistic, the correlation coefficient between one of the measured states (the 'reference state') and the remaining ones, to differentiate among the presence of a single background state and the presence of a ring attractor. The basic idea is that the shape of this distribution conveys information about the encoding. To illustrate this point they derived the distribution of correlation coefficients in a case where multiple features are mapped to a high-dimensional unit sphere (a scenario they referred to as a 'combinatorial encoding') and when different variables map into separate manifolds (a scenario they called 'unary encoding'). The shape of the resulting distribution reflected the dimension of the sphere.

Their calculations assume that the target manifold is a perfect sphere, an assumption that is critical in establishing the shape of the resulting distribution. For example, in the case two circular variables mapping to the two-sphere, the distribution of the correlation coefficient between spontaneous activity and one evoked orientation map, is uniform (Goldberg et al (2004), Fig 4g). Deformations of the sphere, even though they may not change its topology, will generate different distributions of the cross-correlation coefficients. Similarly, if some orientations are more heavily represented than others, as found experimentally (Li et al., 2003), the shape of the distribution of the correlation coefficient will be affected and depend on the reference state. In addition, the calculation of the distribution of correlation coefficients can be carried out only if we know the exact shape of the target manifold. In most situations, unfortunately, we lack such information; this is the case for spontaneous activity or natural image stimulation.

An important observation, and the central motivation for pursuing our study, is that basic questions about encoding in neuronal populations can be answered by looking at topological properties of the population activity alone (without making strong assumptions about the metric of the target space). This point is best appreciated by discussing a few specific examples.

Consider the hypothesis that spontaneous cortical activity wanders between states corresponding to orientation maps (Kenet et al., 2003). This would predict that the cloud of dots in high-dimensional space representing the cortical states at different points in time should be topologically equivalent to a circle, which is essentially a topology question. Intuitively, what needs to be verified is that the data points, when connected with their neighbors at an appropriate spatial scale,

generate a single loop consistent with the topology of a circle.

Another example comes from the work of Goldberg et al (2004) who discuss a 'unary' representation of features. In this scenario, different features map to different manifolds with the cortical state of the system dynamically switching between them (Goldberg et al., 2004). Here, the topological invariant of interest is the number of connected components in the population activity, which provides an estimate of the number of different manifolds used in the encoding.

Finally, Bressloff and Cowan proposed a spherical model of orientation and spatial frequency (Bressloff and Cowan, 2003). This interesting proposal, and one that will turn out to be related to our experimental findings, was inspired by the alignment of pinwheels and extreme spatial frequency preferences in imaging experiments (Issa et al., 2000). Testing if cortical activity is consistent with the topology of a sphere is a question that can be answered using the methods introduced here. These examples illustrate how topological analysis can provide basic, qualitative information about the encoding of information in the brain.

Of course, there are inherent challenges in applying topological concepts to real, high-dimensional data, where we only have a finite number of points and, furthermore, they are likely to be noisy. From a mathematical point of view a key problem is how to generalize the notions of classical topology in a rigorous way so that they can be applied to real data sets.

We begin by addressing this issue. Our first goal is to introduce and describe the basic concepts that are central to the method, including the definition of Betti numbers (which formalizes the notion of counting 'holes' of different dimensions), and the tools of persistent homology, barcodes and witness complexes, which allow Betti numbers to be estimated in experimental datasets. Then, we investigate how the method performs under varying amounts of signal to noise ratio (SNR) in simulations where the correct topological structure of the underlying dataset is known. After this validation step, we apply the method to test the idea that spontaneous activity settles into patterns that may resemble those that appear under natural stimulation (Fiser et al., 2004; Kenet et al., 2003). We confirm that ongoing activity is far from noise and highly structured. Furthermore, we provide the first rigorous evidence that the topological structure of ongoing activity is similar to that obtained with natural image stimulation, which is consistent with a two-sphere. This represents the first empirical finding that supports the spherical model of Bressloff and Cowan (2003).

In what follows we will consider a set of points $\{x_t\}$, where $x_t \in \mathbb{R}^d$ is a d-dimensional vector representing the state of population activity at time t. In our particular case the state represents the instantaneous firing rate of a neuronal population obtained with multi-electrode arrays in primary visual cortex. However, we expect the methodology to be applicable to other situations, such as when the cortical state is defined based on voltage-sensitive dye measurements, or based on calcium indicators in two-photon imaging, or as fMRI BOLD signals. In the sequel, we consider this set of points simply as a cloud of dots in \mathbb{R}^d and ignore the fact that such points originate from a time-series.

2 Tools from Algebraic Topology

In Appendix A we provide a technical overview of the main algebraic topological ideas involved in this work. In this section, however, we present these ideas in an intuitive way.

Algebraic Topology is a subbranch of Topology. The motivating insight behind topology is that some geometric problems depend not on the exact shape of the objects involved, but rather on the way they are put together. For example, the square and the circle have many properties

in common: they are both one dimensional objects (from a topological point of view) and both separate the plane into two parts, the part inside and the part outside. Another way of putting it is that topology attempts to understand the global connectivity of an object by considering how the object is connected locally. Objects are assigned classes such that two objects in the same class exhibit the same connectivity. For example, the square and the circle are in the same class, but the sphere and the circle are not.

Algebraic Topology studies properties of objects (in technical terms: *topological spaces*) and maps between them, in particular, it identifies intrinsic properties of objects by transforming them in certain ways and observing which properties do not change. We call these properties **invariants** of the space. The kind of transformations that we will be interested in are called *homotopy equivalences*. We do not define these transformations here, a precise definition can be found in the appendix, but below we present several graphical examples which attempt to convey the idea behind the definition. Informally, Algebraic Topology, because of its use of homotopy equivalences, is sometimes referred to as “rubber sheet geometry” in the sense that it is oblivious to the fact that by stretching the same piece of rubber one can obtain different looking objects, all it cares about is that it was *the same* piece of rubber that was deformed in certain ways to produce differently looking objects. It therefore is concerned with certain intrinsic properties of objects.

When two spaces X and Y are related by such a transformation, we will say that X and Y are (homotopy) *equivalent* and write $X \sim Y$. When no such transformation exists between X and Y we write $X \not\sim Y$. Figure 1 shows several examples of spaces that are and are not homotopy equivalent.

In particular, the notion from Algebraic Topology that we will be using is **Homology**. Homology measures the connectivity of a space in a precise way: It proceeds by associating to a space X a sequence of abstract spaces $H_0(X), H_1(X), H_2(X), \dots$ derived from X which encode this connectivity information. In particular, the *connectivity* information is summarized by the numbers $b_0(X), b_1(X), b_2(X), \dots$ which are computed from the derived spaces $H_0(X), H_1(X), H_2(X), \dots$, respectively. These are called the *Betti numbers* of the space X . It turns out that if we take a space X and apply an admissible transformation to it (an homotopy equivalence) to obtain the space Y (that is, $X \sim Y$) then $b_k(X) = b_k(Y)$ for $k = 0, 1, 2, \dots$, that is, the Betti numbers are invariants of spaces.

Roughly speaking, one can say that $b_0(X)$ tells us the number of disconnected parts of X , $b_1(X)$ counts the number of *essentially* different loops present in the space X , $b_2(X)$ counts the number of voids et cetera. We will often refer to loops, connected components, voids, et cetera as *features* of the space X .

One can understand the way in which the Betti numbers encode the connectivity information of spaces through some examples. In Figure 2 we show the first three Betti numbers of a few standard objects, which together with Figure 1 describe the Betti number signature of many seemingly different spaces.

In practice, however, one does not have access to a full, nice continuous space such as those shown in Figures 1 and 2. One must instead deal with *finitely* many sample points from a space. We will denote by \mathbb{X} such a discrete space. Therefore there are two considerations that one must take into account when implementing these ideas.

Construction of Simplicial Complexes

Starting from \mathbb{X} one must devise/construct a topological space amenable to the computation of the Betti numbers. A natural candidate, which is suitable for studying the connectivity of (points

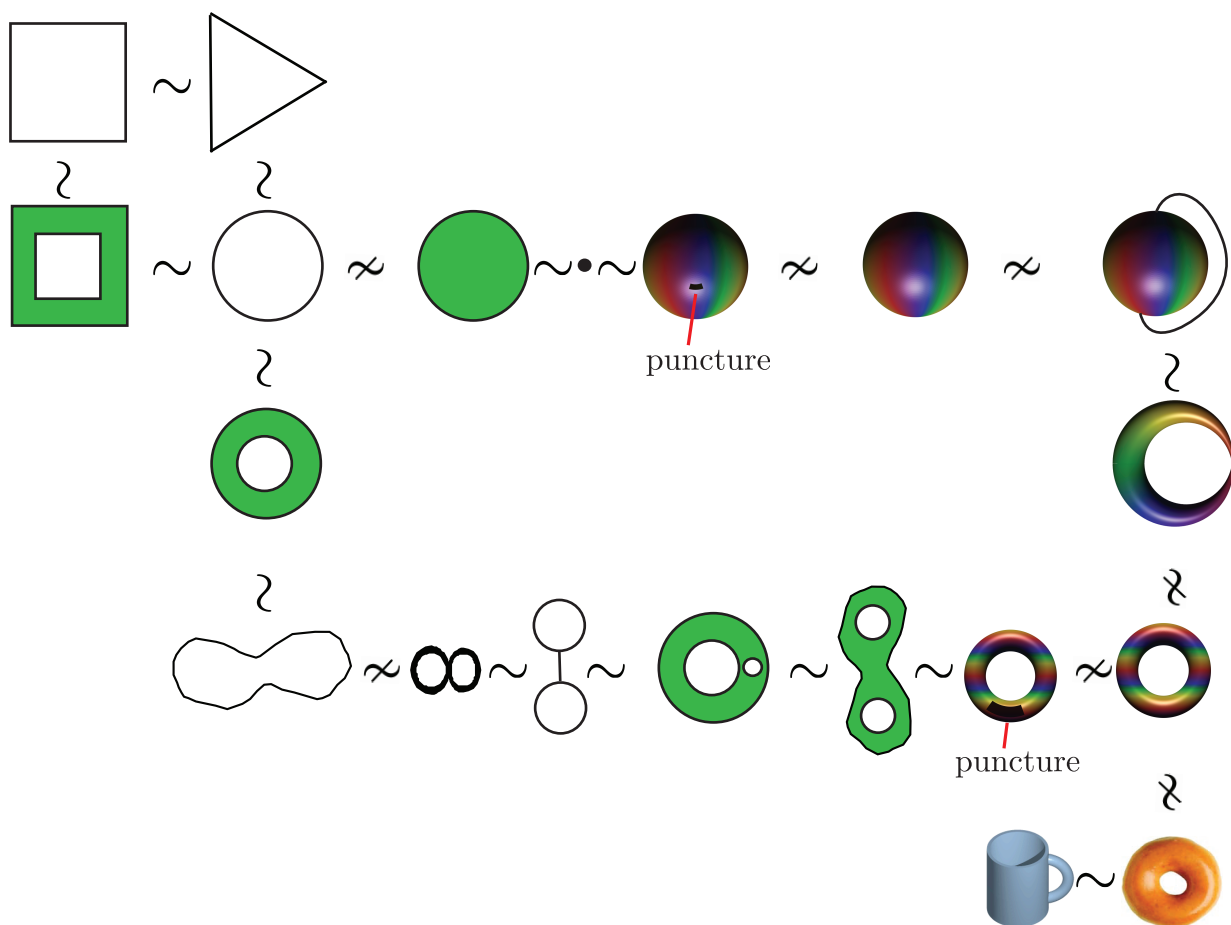


Figure 1: This panel illustrates the notion of homotopy equivalence by showing several topological spaces connected by the symbols \sim when they are equivalent or $\not\sim$ when they are not. The reader should think that all the objects shown are made of an elastic material and is invited to try to visualize the “equivalence” of two spaces by imagining how the deformation of this material takes place.




Figure 2: Five standard spaces together with their Betti number signature. From left to right: a point, a circle, a (void) torus, a Klein bottle and a (void) sphere. For the case of the torus we the figure beneath shows three loops. The red ones are *essential* in that they cannot be shrunk to a point whereas the green loop can in fact be deformed to a point without any obstruction; this is actually picked by the Betti numbers, in particular by the fact that $b_1 = 2$. The Klein bottle has the same Betti numbers as the torus. This is a reflection of the fact Betti numbers do not uniquely characterize the underlying object. For the case of the sphere, the loops shown (and actually all loops on the sphere) can be contracted to points, what is signaled by the fact that $b_1 = 0$. Both the sphere and the torus have $b_2 = 1$, this is due to the fact both surfaces enclose a part of space (a void).




Figure 3: From left to right, simplices of dimensions 0 (points), 1 (edges), 2 (triangles) and 3 tetrahedra.

in) a discrete space is a *graph*. One such construction can be carried out as follows. Pick and fix a number $\epsilon > 0$. Draw an edge between two points in your dataset \mathbb{X} whenever their distance is smaller than or equal ϵ . Figure 7 (a) shows such a construction for a simple case of points sampled near the boundary of a circle. It turns out that this new space, call it $R_\epsilon^{(1)}(\mathbb{X})$ is only useful for computing b_0 . The *building blocks* we assembled together were 1-dimensional (we drew edges), this is the reason for the supraindex (1) in $R_\epsilon^{(1)}(\mathbb{X})$.

The next step is the notion of a triangulation. Say that now, in addition to edges, we consider using 2-dimensional building blocks shaped as triangles. Now we can construct a new space, call it $R_\epsilon^{(2)}(\mathbb{X})$ (richer than our graph) from the dataset \mathbb{X} which will contain $R_\epsilon^{(1)}(\mathbb{X})$ and also all (filled) triangles formed by any three points whose pairwise distances are smaller than or equal ϵ . Figure 7 (b) shows this construction for the simple case of points sampled near the boundary of a circle. Now, it turns out, that since we have included 2-dimensional building blocks in our derived space, we can compute b_0 and b_1 , that is, we can use $R_\epsilon^{(2)}(\mathbb{X})$ for also finding the number of essential loops present in that space.

This construction can of course be generalized to give rise to $R_\epsilon^{(d)}(\mathbb{X})$ for any $d \geq 0$, for example, in the case $d = 3$, we would have to consider drawing a new class of building blocks, which will be 3-dimensional (tetrahedra). Clearly, it is enough to consider $d \leq n$ since for ϵ large enough one will be left with only a single simplex of dimension $n - 1$. We let $R_\epsilon(\mathbb{X})$ denote the Rips complex one obtains for $d = n - 1$ which clearly contains all $R_\epsilon^{(d)}(\mathbb{X})$.

In Topology, there is a specialized term for referring to building blocks of different dimensions, we call them **simplices**. A point is a 0-simplex, an edge is a 1-simplex, a triangle is a 2-simplex, etc. Figure 3 shows a depiction of different simplices. Also, the derived space $R_\epsilon(\mathbb{X})$ receives the name of **Rips complex**. In Algebraic Topology, a *simplicial complex* is a space which one obtains by assembling together *simplices* while respecting certain rules. Figure 4 shows example of a valid construction. For a precise definition see the appendix.

One last note before moving on to another important consideration is that in all our constructions there is an implicit notion of *metric*, that is, we said for example that we will draw an edge between “two points whenever their distance is smaller than..”. This notion of distance is application dependent and can be decided upon depending on the nature of the dataset.

Multiscale Computations

When we compute the Betti numbers of $R_\epsilon(\mathbb{X})$ there is an implicit choice for the value ϵ . This scalar can be thought of as fixing a *scale* for the computations. Larger values of ϵ will produce a Rips complex with more simplices. When $\epsilon = 0$ the Rips complex is just a set of points, with no connections between them, therefore the underlying topology, as signaled by the Betti numbers, is




Figure 4: A (valid) simplicial complex. Edges are shown in black, triangles in orange and tetrahedra in green.

very simple: $(b_0, b_1, \dots) = (n, 0, 0, \dots)$ where n is the number of points in X . When ϵ equals the maximal distance between any two points in \mathbb{X} , all possible simplices will be in the Rips complex $R_\epsilon(\mathbb{X})$. In technical terms one says that $R_\epsilon(\mathbb{X})$ is a **blob**. It follows from the theory that its Betti numbers signature is again very simple.¹ One has $(b_0, b_1, b_2, \dots) = (1, 0, 0, \dots)$. However, if there is any real feature underlying the dataset, one would hope to recover it for some intermediate value of ϵ . Therefore, in practice one computes the Betti number signature for all values of ϵ at once. Also, it is important that when doing this computation one keeps track of the *life* of the features captured by each Betti number. In other words, one will be computing the Betti numbers across scales in a *persistent* way. This leads to the definition of the **Betti barcodes**. Without entering into technicalities which are deferred to the appendix, Figure 6 shows an example computation of the barcodes corresponding to (b_0, b_1, b_2) for points sampled on the surface of a torus.

For each Betti number one plots the *birth* and *death* of individual topological structures as ϵ goes from zero to infinity. Each horizontal line segment in the 'barcodes' of Figure 6 (a) represents a different structure (a high-dimensional hole) with their ends indicating the points (corresponding to different spatial scales) where it was created (left) and destroyed (right). Different intervals are displaced vertically to allow their visualization; beyond this, the vertical axis has no significance. The first Betti number, b_0 , which represents the number of connected components will tend to one as ϵ increases, while b_j for $j > 0$ will go to zero for large values of ϵ .

The horizontal axes will therefore always represent, from left to right, increasing values of the scale parameter ϵ . Longer lines are expected from features that are more prominent while short ones will arise from small details, and usually termed *topological noise*.

The reader can observe how for a range of values of the persistent parameter, the signature is correct. In order to compute the Betti numbers in this persistent fashion one needs to keep track of the topological features while going from $R_\epsilon(\mathbb{X})$ to $R_{\epsilon'}(\mathbb{X})$ for $\epsilon' > \epsilon$. This is made possible thanks to a fundamental property enjoyed by homological constructions termed **functoriality**, see the appendix. All computations of barcodes in the paper were performed using the software PLEX freely available from <http://math.stanford.edu/comptop/programs/>.

A final point is that frequently the Rips complex gets too large: too many simplices may constitute it for a range of values of ϵ . Other useful constructions of simplicial complexes have been devised for achieving computational efficiency, with **Witness complexes** being the umbrella

¹In fact, it is clear that we will have only one connected component and thus $b_0 = 1$.




Figure 5: This figure explains how to read barcodes using a dataset of points \mathbb{Y} regularly sampled from the surface of a torus. We endowed the points with Euclidean metric and used a Rips simplicial complex construction. The barcode resulting from running persistent Betti numbers computations is displayed on the right for (b_0, b_1, b_2) . The top subfigure represents b_0 , the one in the middle b_1 and the bottom one b_2 . The horizontal axes always represent the persistent parameter ϵ , which increases from left to right. If one wants to know what are the Betti numbers of $R_{\epsilon_0}(\mathbb{Y})$ for a certain value of ϵ_0 , one needs to slice the three barcodes at that value and count how many horizontal lines in each subfigure are intersected by the vertical slicing line. This is illustrated here with a vertical red line in the barcodes; counting gives $(b_0, b_1, b_2) = (1, 2, 1)$. Long lines in the barcodes are to be interpreted as salient/persistent features of the dataset. For a particular range of the persistence parameter, indicated by the shaded area, there is a stable signature $(b_0, b_1, b_2) = (1, 2, 1)$ indicative of a the torus. In assessing the significance of these results one needs to evaluate that probability that such structures would result from the null hypothesis of random data.

term for them. A detailed explanation of the construction can be found in the appendix, the only concept that the reader should retain is that witness complexes can be used to deal with larger datasets and that they achieve efficiency by splitting the dataset \mathbb{X} into two disjoint sets, a small set, the *landmark set* and then the set of *witnesses*. The former set is used to justify the addition of a small number of simplices which (generally) approximate well the underlying structure of the dataset. Figure 7 shows both the Rips and Witness complex constructions (only 1 and 2-dimensional simplices are shown) for the same set of points sampled around the bounday of a circle. The set of landmarks is frequently built using the so called *max-min* procedure described in the appendix.

3 Results

3.1 Validation of the method

Our first step was to validate the methodology using artificial datasets where the 'ground truth' (in this case the topological structure) is known. To approximate the statistical properties of the point cloud in neural data, we simulated population of Poisson spiking neurons tuned to various parameters. The objective of the simulations was to evaluate the minimum number of neurons and mean firing rates required to recover the correct topological structure of different objects, such as a circle and a torus.

3.2 Validation dataset # 1: a circle

Our first simulation is designed to simply verify that one is indeed able to recover the structure of a circle from the simulation of a population of cells with homogeneous tuning curves, where the mean spike count per time window is given by $\lambda(\theta) = r_{max} \exp(\kappa \cos \theta) / \exp(\kappa)$, and the spike counts are

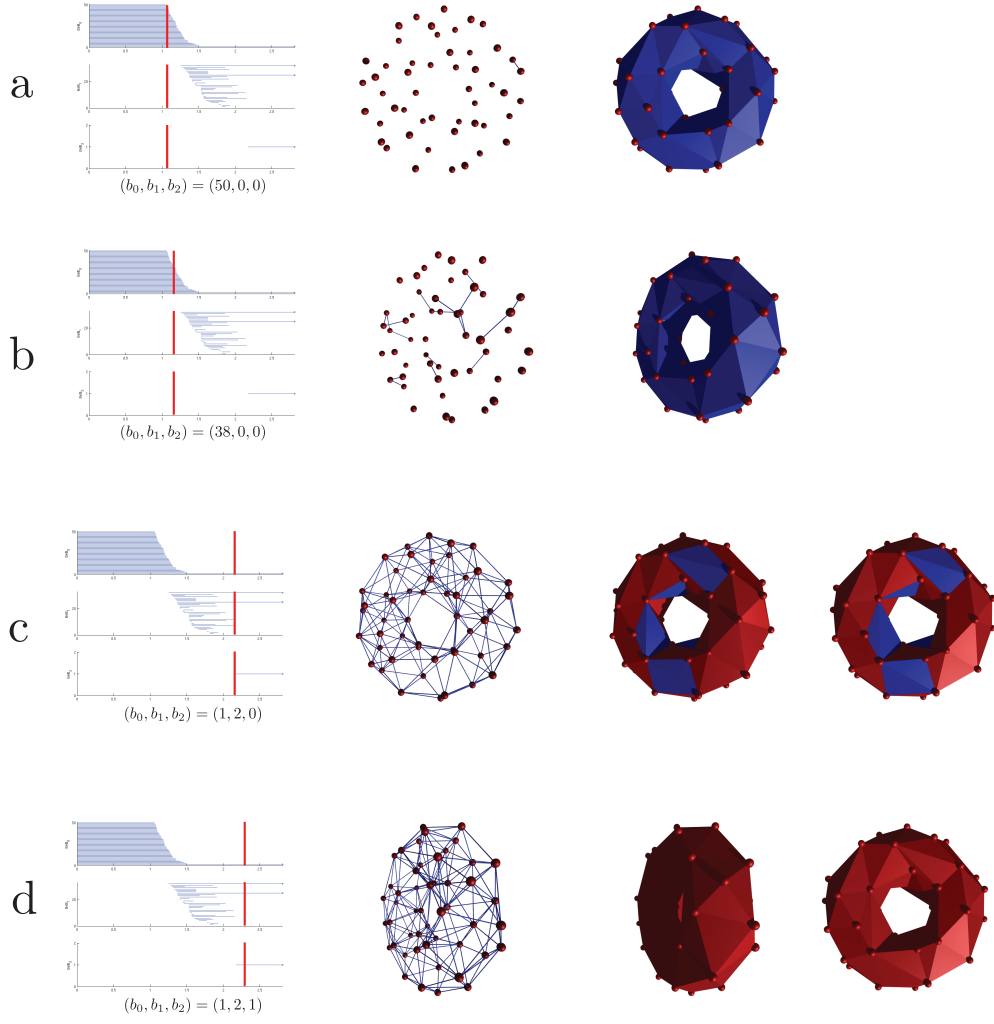


Figure 6: *Reading Barcodes*. This figure shows the Betti barcode computed (only first three Betti numbers are displayed) for the Rips complex on a set of 51 points randomly sampled from the surface of a torus. Subfigures (a) to (d) show the barcode sliced at different values of ϵ with different Betti signatures. The three dimensional renderings on each of (a), (b), (c) and (d) show $R_\epsilon^{(1)}$ and $R_\epsilon^{(2)}$ for the corresponding value of ϵ . Triangles in red are the ones that already belong to the simplicial complex and those in blue are the ones that will belong to the complex at a later value of ϵ . Note for example that in (a), only one edge has been drawn and therefore the graph $R_\epsilon^{(1)}$ has many components, no loops and no voids $((b_0, b_1, b_2) = (50, 0, 0))$; (b) shows more edges and still no loops nor voids; In (c) we already have only one component, and the fact that the red vertical line intersects two horizontal bars in the signature corresponding to b_1 reveals that we have two essential loops, whereas the same line does not intersect any horizontal line corresponding to b_2 , what means that there is no void. This can be confirmed by looking at $R_\epsilon^{(2)}$ and checking that there are some punctures (missing triangles). Finally, (d) shows a slice for which the signature is correct.

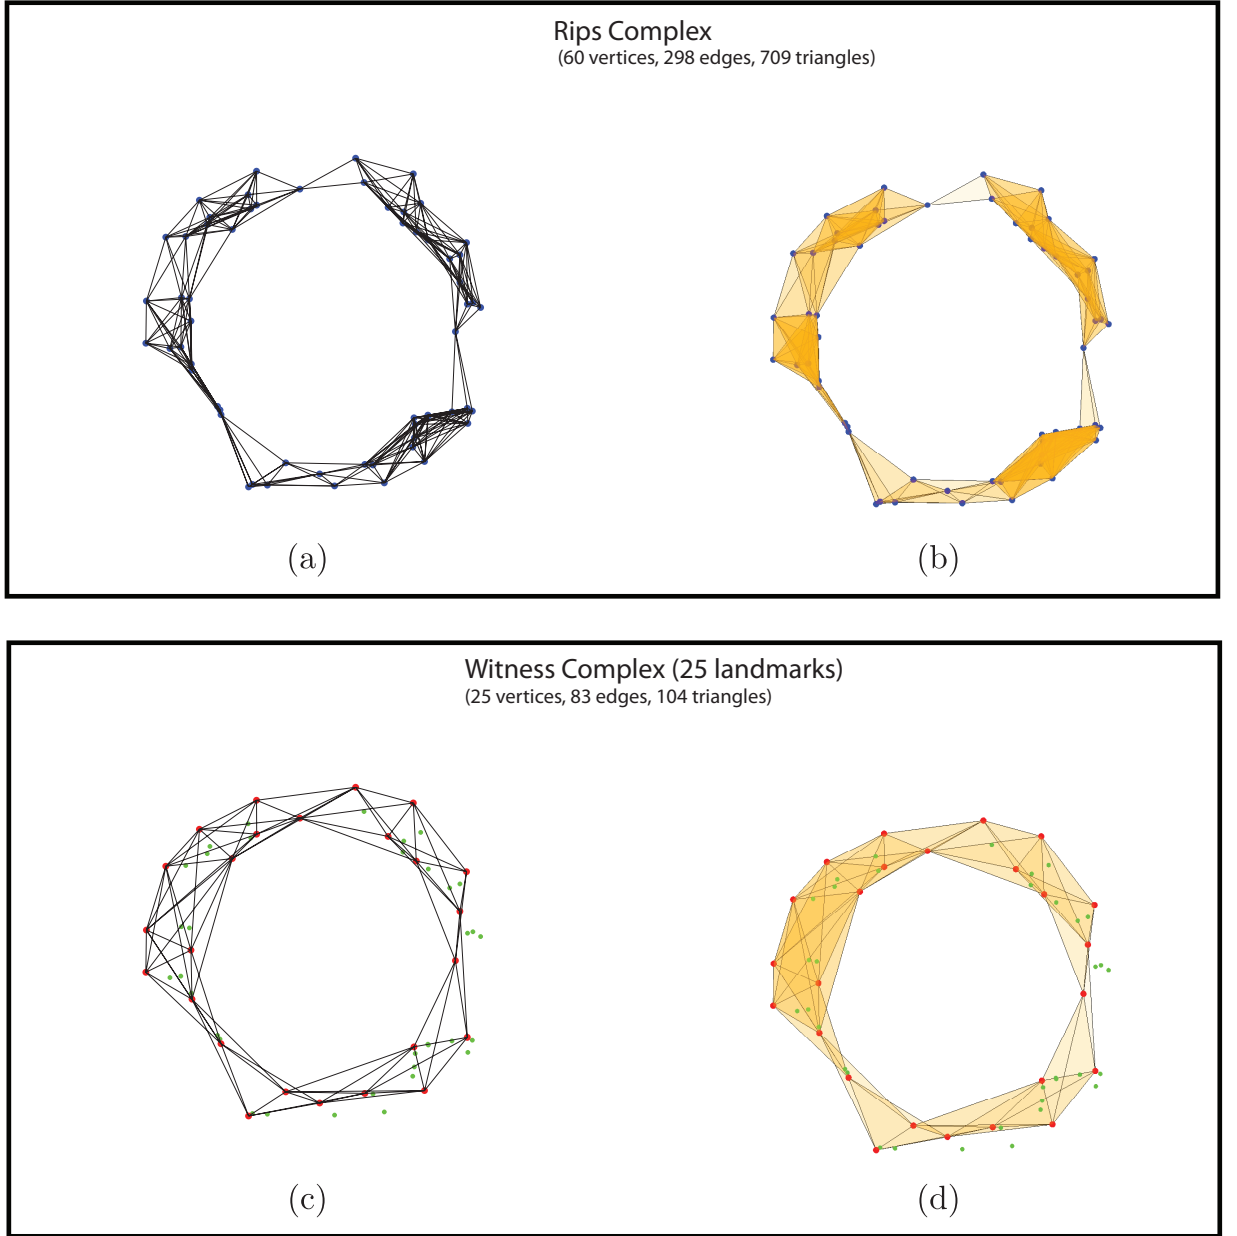


Figure 7: In this panel, \mathbb{Y} is a fixed set of points sampled around the boundary of a circle. (a) and (b) show $R_\epsilon^{(1)}(\mathbb{X})$ and $R_\epsilon^{(2)}(\mathbb{Y})$, respectively. (c) and (d) show the corresponding Witness Simplicial complexes for a comparable value of ϵ . In general by using Witness complexes one can obtain simplicial complexes that faithfully represent the underlying structure of a dataset by using much fewer simplices than a Rips construction would yield.

Poisson distributed. The total number of cells in the population will be denoted by N . For any given population, the preferred orientations will be equally spaced, such that the tuning curve of the i -th cell is given by $\lambda_i(\theta) = r_{max} \exp(\kappa \cos(\theta - 2\pi i/N)) / \exp(\kappa)$, for $i = 0, \dots, N-1$. We chose a value $\kappa = 2$ which represents the average tuning as observed experimentally (Ringach et al., 2002). Here, r_{max} represents the mean number of spike counts in a time-bin. Fig 2a illustrates the tuning curves for $\kappa = 2$ and $N = 3$. In all cases we simulated 100 presentations of 18 orientations equally spaced around the circle. Thus, in all situations there are a total of 1800 points in the data set which can be arranged in a data matrix of spike counts (Fig 2b). This represents our point cloud. We then calculated the maximal length of the signature $(b_0, b_1, b_2) = (1, 1, 0)$ using 25 landmarks positioned by the max-min procedure (Fig 2c).

These simulations were performed for different values of (r_{max}, N) to evaluate the performance of the algorithm at various levels of SNR and population size. For each pair of values we generated datasets G_1, \dots, G_{10} using the procedure described above. Using each dataset G_i , we generated shuffled control data $B_1^{(i)}, \dots, B_{100}^{(i)}$ by randomly shuffling the elements of the data matrix. For example, the barcode of the shuffled data in Fig 2c is shown in Fig 2d. If the shuffling operation destroys structure in the data one expects the length of intervals for Betti numbers higher than zero to decrease, as seen by comparing the barcodes in Fig 2c,d.

For each G_i and $B_k^{(i)}$ we compute the length of the longest intervals in which we observe the topological signature $(b_0, b_1, b_2) = (1, 1, 0)$ and obtain numbers l_i and $l_k^{(i)}$ respectively. Then, we calculated $p(r_{max}, N) := \frac{1}{10} \sum_i \frac{\#\{k | l_i < l_k^{(i)}\}}{100}$. The reason for generating G_1, \dots, G_{10} and taking the average over in the preceding equation is to get an estimate of the average probability over different experiments.

A plot of $p(r_{max}, N)$ is shown in the pseudo-color image of Fig 3a. The dashed line shows the approximate boundary for detecting a circle at a significance level of $p < 0.05$. This analysis makes it evident that there is a trade-off between the number of cells and mean spike counts per time bin that is necessary to detect the circle at a confidence level of $p < 0.05$. The larger the number of neurons, the smaller the spike rates can be and still allow for the reliable estimation of the underlying topology. For five cells, for example, one would need average spike counts $\simeq 4.5$ spikes per time-bin; for 10 cells, on the other hand, the rates can be as low as 1.5 spikes per time-bin.

3.3 Validation dataset # 2: a torus

Next, we investigated the ability of the technique to recover the structure of a torus when we simulated a population of cells with tuning curves over two circular variables (θ, ϕ) , given by $\lambda(\theta, \phi) := r_{max} \exp(\kappa_\theta \cos \theta + \kappa_\phi \cos \phi) / \exp(\kappa_\theta + \kappa_\phi)$. As before this represents the mean spike counts per bin and the spike counts were Poisson distributed. The total number of cells in the population will be denoted by N . For any given population, the centers of the tuning curves, (θ_i, ϕ_i) , were chosen randomly inside the rectangle $[0, \pi] \times [0, 2\pi]$, such that the tuning curve of the i -th cell was given by

$$\lambda_i(\theta, \phi) := r_{max} \exp(\kappa_\theta \cos(\theta - \theta_i) + \kappa_\phi \cos(\phi - \phi_i)) / \exp(\kappa_\theta + \kappa_\phi)$$

for $i = 1, \dots, N$. For this simulation we used values of $\kappa_\theta = 2$ and $\kappa_\phi = 1.5$. Here, r_{max} represents the mean number of spike counts in a time-bin. In all cases we simulated 100 presentations of 400 orientations and phases equally spaced around the circle: our set of stimuli were pairs $(\pi k/20, \pi l/20)$, where $k, l = 0, 1, \dots, 19$. Thus, in all situations there are a total of 40,000 points in the data set.

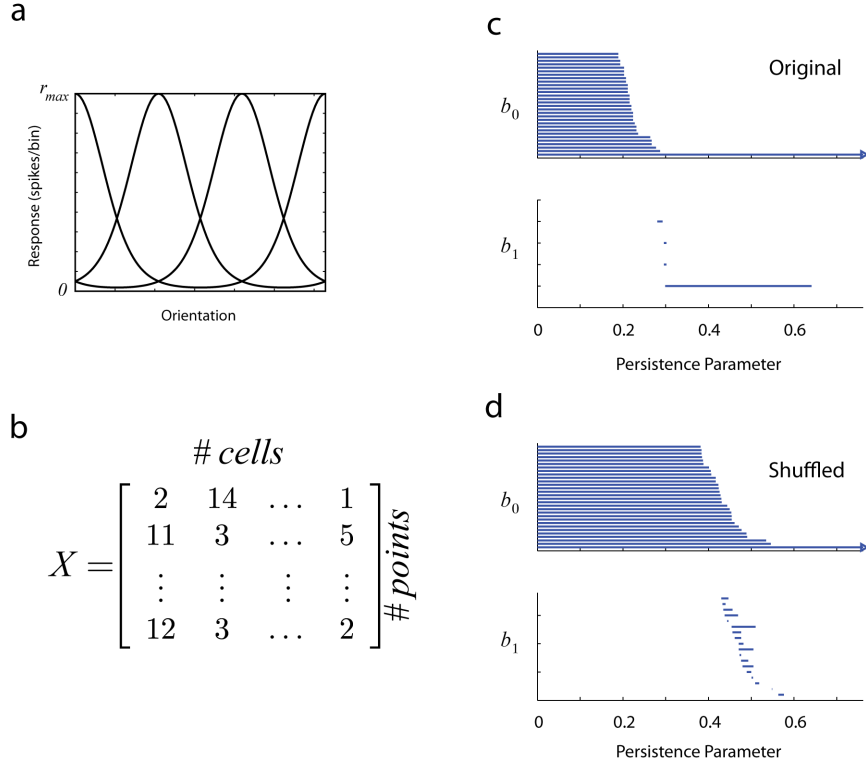


Figure 8: Validation of the techniques using Montecarlo simulations in the case of a circle. (a) We assume an initial population of Poisson-spiking neurons tuned for orientation. (b) The simulated response of this population to the presentation of different orientations is collected into a data matrix (or point cloud). (c) Analysis of the simulated data shows a long interval with a signature of $(b_0, b_1) = (1, 1)$, which correctly identifies the circle. (d) We also compute the bar codes by shifting the relative positions of the columns (data shuffling). In this case, the statistical distributions of spike counts for each axis remain unchanged, but their relationship is destroyed. By computing the distribution of maximal b_1 lengths under this null hypothesis, we can evaluate the likelihood that our data was generated by the null hypothesis.

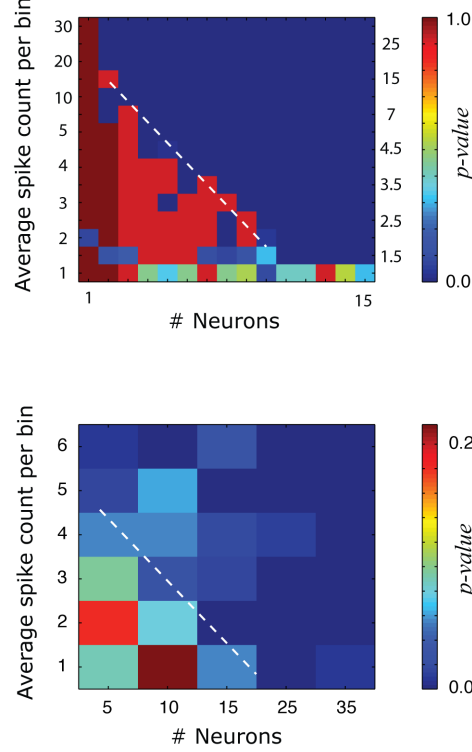


Figure 9: Recovering the structure of data consistent with a circle and a torus. The panels illustrate the capability of the algorithm in recovering the structure of a circle and torus for different values of the total number of cells and their maximal firing rates. Small p -values are regions where the algorithm correctly identified the circle and the torus and the likelihood of obtaining such data by chance is low. There is a tradeoff between number of neurons and maximal spike rates. The more neurons the smaller the firing rates can be to be able to detect the structures at the same level of significance.

We then calculated the maximal length of the signature $(b_0, b_1, b_2) = (1, 2, 1)$ using 50 max-min landmarks. To evaluate the statistical significance of the result we computed the same statistic for 50 random permutations of all the elements within the data matrix.

Following the analogous procedure for the circle, for each pair (r_{max}, N) , we generated a few datasets G_1, \dots, g_5 . Using each G_i , we generated shuffled datasets $B_1^{(1)}, \dots, B_{50}^{(i)}$. For each G_i and $B_k^{(i)}$ we computed the length of the longest intervals in which we observed the signature $(b_0, b_1, b_2) = (1, 2, 1)$, and obtain numbers l_i and $l_k^{(i)}$ respectively. Then, we computed $p(r_{max}, N) = \frac{1}{5} \sum_i \frac{\#\{k | l_i < l_k^{(i)}\}}{50}$.

The result showing the behavior of $p(r_{max}, N)$ is shown in the pseudo-color image of Fig 3b. The dashed line shows the approximate boundary for detecting a torus at a significance level of $p < 0.05$. As for the case of the circle, the larger the number of neurons, the smaller the spike rates can be and still allow for the reliable estimation of the underlying topology. For five cells, for example, one would need average spike counts $\simeq 4.5$ spikes per time-bin, while for 15 cells, on the other hand, the rates can be as low as 1.5 spikes per time-bin.

4 Empirical data: the topology of spontaneous and driven activity in V1

We now turn to the analysis of empirical data. As described in the Methods these data were obtained from multi-electrode recordings in primary visual cortex of macaque V1. Fig 4a illustrates one insertion sequence into V1 using a 10x10 electrode array. Using these electrode arrays we recorded the population activity in two experimental conditions. In the spontaneous activity condition, population activity was measured when there was no visual stimulus present on the screen (the eyes were occluded). In the natural image stimulation condition, we presented a video sequence obtained by sampling different movie clips and the evoked spike trains from the neuronal population were recorded (Fig 4b). A total of 20 to 30 minutes of data were collected in both conditions. These data were split into 10 second segments and spike trains binned as described in Section "Methods". The goal was to characterize the distribution of the 'topological signatures' of the data segments in both conditions.

Fig 5a illustrates the various topological signatures observed labeled by the first three Betti numbers (b_0, b_1, b_2) . On top of each triplet of Betti numbers an object consistent with each signature is shown. The distributions of topological signatures for both experimental conditions are shown in the histograms of Fig 5b, where the x-axis represents the same ordering of signatures as depicted in Fig 5a. The left column represents distributions for the spontaneous condition, while the right column represents the distributions for the natural image stimulation condition. Each row represents a different threshold for the length of the signature (in the barcode) as a fraction of the covering radius of the data. Larger thresholds represent instances where the signature was 'long-lived' and likely to represent a salient feature of the data. We emphasize that all the topological features shown are statistically different from noise, as Montecarlo simulations using shuffled data show that the probability of obtaining segments of b_1 or b_2 longer than 0.3 by chance (which is the smallest threshold used) was less than 0.005.

Interestingly, at high threshold values, there are two main signatures that dominate: a circle with a signature of $(1,1,0)$ and a sphere with a signature of $(1,0,1)$. This holds for both experimental conditions. Thus, both spontaneous and natural image stimulation have similar topological distributions. However, the relative frequency of observation of these signatures appears to be different, as the likelihood of observing the signature of a circle is higher with natural image stimulation rather than in the spontaneous condition. At a lower threshold of 0.3 we observe a more diverse distribution of topological signatures for the spontaneous activity, while the distribution of driven activity remains dominated by the signatures of a circle and a sphere (Fig 5b, bottom panels).

5 Discussion

We have seen that computational topology can help address basic questions about the encoding of information by neuronal populations. The result of the analysis is a topological characterization of the activity, which provides qualitative information about its structure, such as the number of clusters and loops in the activity patterns. A key concept we used, originating in the work of Edelsbrunner and colleagues (Edelsbrunner, 1998; Edelsbrunner et al., 2000), was that of 'persistent homology', where 'holes' of different dimensions are tracked as the 'spatial scale' of analysis changes. Structures that are present across a substantial range of the persistence parameter are likely to be real features of the data. We showed how Montecarlo simulations can be used to test if the structures

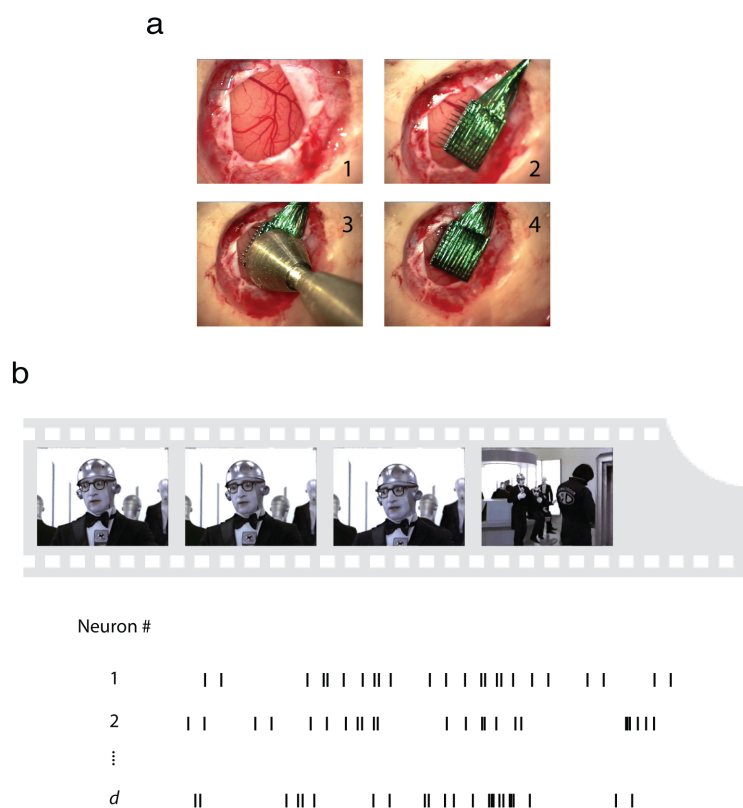



Figure 10: The experiment. (a) Multi-electrode arrays are inserted into primary visual cortex. The panels show an insertion sequence. For a more detailed description of the method see (Nauhaus and Ringach, 2007). (b) Natural image sequences, sampled from commercial movies, were used to stimulate all receptive fields of neurons isolated by the array. In the spontaneous condition, activity was recorded while both eyes were occluded.

a



b




Figure 11: Experimental results. (a) Ordering of topological signatures observed in our experiments. Each triplet (b_0, b_1, b_2) is shown along an illustration of objects consistent with the signature. (b) Distribution of topological signatures in the spontaneous and natural image stimulation conditions. Each row correspond to signatures with a minimum length (denoted as the threshold) expressed as a fraction of the covering radius of the data cloud.

observed could have resulted from the null hypothesis of a set of randomly spiking neurons.

First, through the use of computer simulations we demonstrated that the method works as expected, by recovering the underlying structure of the data in artificial datasets where the topology of the data cloud was under our control. These simulations were done incorporating the variability one would expect from real neuronal data. This was achieved by using Poisson distributed spike counts with mean firing rates, population sizes and record lengths, comparable to those in the actual experiments.

We then explored the structure of population activity when primary visual cortex was spontaneously active and when it was driven by natural image sequences. It was found that the structure was similar in both cases, consistent with prior results suggesting that natural stimulation modulates ongoing activity only weakly (Arieli et al., 1995; Fiser et al., 2004; Grinvald et al., 2003; Tsodyks et al., 1999).

Our results go beyond prior investigations by providing the first rigorous study of the topological structure of population activity. We showed that both the data for spontaneous and driven conditions were similarly distributed, with the signatures of the circle and the sphere dominating the results. We think one possible interpretation of the data is that there might be a single underlying object that is insufficiently sampled, or sampled with biases, during our 10 sec time intervals that are used to construct our data clouds. This variability/bias in the sampling results in a distribution of topological signatures, rather than a single answer. Specifically, it is easy to see that all the geometries represented in Fig 5a can be obtained from a torus by application of a small number of operations such as removing a single point, collapsing a pair of points to a single point, and collapsing a non contractible closed curve to a point. For example, the geometry with signature $(1, 3, 1)$ can be obtained by collapsing a pair of points to a single point, the one with signature $(1, 1, 1)$ can be obtained by collapsing a closed curve to a single point, and the geometry with signature $(1, 1, 2)$ can be obtained by performing two collapses of closed curves to points. The same remarks apply to a Klein bottle, i.e. all the geometries can be obtained by a small number of applications of these operations.

Despite these caveats, at high thresholds, our data are consistent with a single spherical object, which we postulate, represents a mapping of orientation and spatial frequency as suggested by Bressloff and Cowan (2003). In this model, orientation is mapped to the azimuth and spatial frequency with elevation and the poles representing zero and infinite spatial frequencies. These findings are consistent with the data of Kenet et al (2003) and, in fact, they may explain why 80% of the time the cortical state is uncorrelated with the orientation maps. If the proposed spherical model is correct, and the state wandered out of the equator near the poles, one may expect the activity to be uncorrelated with the orientation states lying near the equator. A strong prediction of our study is that if one were to repeat the Kenet et al experiment by imaging both orientation and spatial frequency columns, the result will be that when activity is uncorrelated with the orientation maps, they will be correlated with the spatial frequency maps.

There are many areas where topological analyses of neural activity can help guide further research. First, the tools can be used to test specific hypotheses, such as 'is the activity consistent with a single loop?' Second, it provides a rigorous tool to study the phenomenon of cortical 'songs', where repeated patterns of activity have been interpreted as attractors (Cossart et al., 2003; Deneve et al., 2001; Ikegaya et al., 2004; Latham and Nirenberg, 2004; Tsodyks, 1999). The statistical analyses of these recurring patterns is a delicate matter, and it has been suggested that the patterns may not be present at all (Mokeichev et al., 2007; Oram et al., 1999; Richmond et al., 1999; Wiener

and Richmond, 2003). The presence of distinct stable fixed points, or line attractors, is something that could be tested with our methods, as they would show up as different connected components in the analysis. Third, topological analysis may be suitable for exploring the basic structure of population activity in situations where we have no prior information, or specific hypotheses, about the structure of the stimulus or the encoding. The encoding of odors and object shapes, are good examples that may fall within these categories (Carrasco and Ridout, 1993; Edelman, 1998; Fdez Galan et al., 2004; Feldman and Richards, 1998; Kayaert et al., 2005; Mazor and Laurent, 2005; Tanaka et al., 1991; Wehr and Laurent, 1996). Fourth, understanding the topological structure of population activity may help in the design of better decoding methods for use in brain-machine interfaces (Andersen et al., 2004; Donoghue, 2002; Jazayeri and Movshon, 2006; Nicolelis, 2003; Nicolelis and Chapin, 2002; Ohnishi et al., 2007; Santhanam et al., 2006; Serruya et al., 2003; Shoham et al., 2005). For example, if one knew that the activity of a population of neurons in a high dimensional space was equivalent to that of a circle, one could collapse the entire activity to single number (such as the distance from a reference location).

6 Material and Methods

6.1 Animal Preparation

Experiments were approved by the UCLA Animal Research Committee and were performed following the National Institutes of Health’s Guidelines for the Care and Use of Mammals in neuroscience. Old-world monkeys (*Macaca fascicularis*, 3-5kg) were used. Initially, animals were sedated with acepromazine (30-60 g/kg) and anesthetized with ketamine (5-20mg/kg, im). Initial surgery was then performed under 1.5-2.5% isoflurane. Two intravenous lines were put in place for the continuous infusion of drugs. A urethral catheter was inserted to collect and monitor urine output. An endotracheal tube was inserted to allow for artificial respiration. Pupils were dilated with ophthalmic atropine, and the eyes protected with ophthalmic Tobradex (Alcon Laboratories, Texas) and custom-made gas permeable contact lenses.

At the completion of this initial surgery, the animal was transferred to a stereotaxic frame. At this point, anesthesia was switched to a combination of sufentanil (0.15 g/kg/h) and propofol (2-6 mg/kg/h). After monitoring the anesthetic plane for about 10-20 minutes we proceeded to perform a craniotomy over primary visual cortex. Only after the completion of all surgical procedures, including the insertion of the electrode array, the animal was paralyzed (Pavulon, 0.1 mg/kg/h).

To ensure a proper level of anesthesia throughout the duration of the experiment, rectal temperature, heart rate, noninvasive blood pressure, end-tidal CO₂, SpO₂, and EEG were continually monitored by an HP Virida 24C neonatal monitor. Urine output and specific gravity were measured every 4-5 h to ensure adequate hydration. Drugs were administered in balanced physiological solution at a rate to maintain a fluid volume of 5-10 ml/kg/h. Rectal temperature was maintained by a self-regulating heating pad at 37.5C. Expired CO₂ was maintained between 4.5 and 5.5% by adjusting the stroke volume and ventilation rate. The maximal pressure developed during the respiration cycle was monitored to verify that there was no incremental blocking of the airway. A broad spectrum antibiotic (bicillin, 50,000 IU/kg) and anti-inflammatory steroid (dexamethasone, 0.5 mg/kg) were given at the beginning of the experiment and every other day.

6.2 Electrophysiology

The database considered in this study was obtained using micro-machined electrode arrays (Cyberkinetics, Salt Lake City, UT) consisting of a square grid of either 10x10 or 5x5 electrodes of 1- or 1.5-mm in length. The distance between neighboring electrodes was either of for the 10x10 arrays and for the 5x5 arrays. The receptive fields of neurons from the arrays overlapped significantly (only those at opposite ends of the array were non-overlapping). Thus, our recordings come from populations whose receptive fields are responding to the same area of visual space.

Spike sorting was performed offline using principal component analysis on the waveform shapes with software developed in our laboratory. Stimuli were generated on a Silicon Graphics O2 and displayed on monitor at a refresh rate of 100 Hz and a typical screen distance of 80 cm. The mean luminance was 56 cd/m². A Photo Research Model 703-PC spectro-radiometer was used for calibration. The eyes were initially refracted by direct ophthalmoscopy to bring the retinal image into focus for a stimulus roughly 80 cm from the eyes. Once neural responses were isolated, we measured spatial frequency tuning curves and maximized the response at high spatial frequencies by changing external lenses in steps of 0.25 D. This procedure was performed independently for both eyes. All recordings originate from eccentricities of 2-7 deg.

6.3 Visual Stimuli

In the spontaneous condition the eyes were occluded. The stimuli in the evoked condition were image sequences generated by digitally sampling commercially available videotapes in VHS/NTSC format. Images had a spatial resolution of 320 240 pixels and were sampled at a temporal rate of 15 Hz. The selected movies included both man-made and natural landscape scenes. Six segments of 30-s duration were used, making a total of 24 minutes of video. The movies were compressed using Silicon Graphics' MVC2 compression scheme (proprietary) and stored on a disk. A Silicon Graphics O2 computer played back the images during the experiment on a computer screen that measured 34.3 cm wide by 27.4 cm high. The refresh rate of the monitor was 90 Hz and each movie image was presented for six consecutive frames. The mean luminance of the display was 56 cd/m². Stimulation was monocular to the dominant eye (the other eye was occluded). The images subtended 6 deg 4.5 deg of visual angle and covered all the receptive fields under measurement.

6.4 Creating point clouds

The preparation of the data points for both the spontaneous and driven activity during natural image stimulation were identical. After spike-sorting signals from each electrode we sub-selected a group of 5 neurons that showed the highest firing rates. Subsequently, a point cloud was generated by binning spikes in 50ms windows. Spontaneous activity segments were collected in lengths of 10s each. Thus, each segment contained 200 points living in \mathbb{R}^5 . The software package PLEX was used with a weak witness complex construction. PLEX is a Matlab collection of functions for computational topology and is freely available from <http://math.stanford.edu/comptop/programs/>. We used a weak witness construction with 35 landmarks points which were selected using the 'maxmin' procedure. The 'maxmin' procedure was seeded with each one of the 200 points in the dataset, in order to eliminate dependence on our initial selection. We recorded the maximal length of persistence intervals for b_1 and b_2 for each of the 200 seeds.

6.5 Statistical Significance

To evaluate the probability that the barcodes could have resulted from independent firing of neurons across the electrodes we generated control datasets as follows. For each channel we computed first the total number of spikes. Then, we generated a new dataset by randomly positioning the same number of spikes in time. This keeps the total number of spikes for each neuron constant but destroys any potential relationship between them. The identical analysis done for the real dataset was performed for the control datasets (a total of 52700 times) generating a null distribution for bar lengths under the hypothesis of Poisson neurons firing independently. Finally, one can perform a one-tail rank sum test to verify that the median distribution of lengths in the data is significantly higher than that expected by chance. In all our experiments the results were highly significant (rank-sum test, $p < 10^{-10}$).

7 Acknowledgements

This work has been supported by DARPA 32905 (GS, FM, TI and GC), NSF DMS 0354543 (GC), ONR N000140310176 (GS), NSF 0309575 (GS), DARPA FA8650-06-1-7630 (GS), 14168480-32905-C (GS), NGIA HM1582-04-1-2023 (GS), NIIH-EY 12816 (DLR), FA8650-06-C-7633 (DLR).

A Simplicial Homology

In this section we present a technical overview of Homology as used in our procedures. For an intensive treatment we refer the reader to the excellent text of Hatcher [20].

Homology is an algebraic procedure for counting *holes* in topological spaces. There are numerous variants of homology: we use simplicial homology with Z_2 (binary, 0 or 1) coefficients. Given a set of points V , a k -simplex is an unordered subset $\{v_0, v_1, \dots, v_k\}$ where $v_i \in V$ and $v_i \neq v_j$ for all $i \neq j$. The faces of this k -simplex consist of all $(k-1)$ -simplices of the form $\{v_0, \dots, v_{i-1}, v_{i+1}, \dots, v_k\}$ for some $0 \leq i \leq k$. Geometrically, the k -simplex can be described as follows: given $k+1$ points in \mathbb{R}^m ($m \geq k$), the k -simplex is a convex body bounded by the union of $(k-1)$ linear subspaces of \mathbb{R}^m of defined by all possible collections of k points (chosen out of $k+1$ points). A simplicial complex is a collection of simplices which is closed with respect to inclusion of faces. Triangulated surfaces form a concrete example, where the vertices of the triangulation correspond to V . The orderings of the vertices correspond to an orientation. Any abstract simplicial complex on a (finite) set of points V has a geometric realization in some \mathbb{R}^n . Let X denote a simplicial complex. Roughly speaking, the homology of X , denoted $H_*(X)$, is a sequence of vector spaces $\{H_k(X) : k = 0, 1, 2, 3, \dots\}$, where $H_k(X)$ is called the k -dimensional homology of X . The dimension of $H_k(X)$, called the k th Betti number of X , is a coarse measurement of the number of different holes in the space X that can be sensed by using subcomplexes of dimension k .

For example, the dimension of $H_0(X)$ is equal to the number of connected components of X . These are the types of *features* (*holes*) in X that can be detected by using points and edges— with this construction one is answering the question: are two points connected by a sequence of edges or not? The simplest basis for $H_0(X)$ consists of a choice of vertices in X , one in each path-component of X . Likewise, the simplest basis for $H_1(X)$ consists of loops in X , each of which surrounds a *hole* in X . For example, if X is a graph, then the space $H_1(X)$ encodes the number and types of cycles in the graph, this space has the structure of a vector space. Let X denote a simplicial complex.

Define for each $k \geq 0$, the vector space $C_k(X)$ to be the vector space whose basis is the set of oriented k -simplices of X ; that is, a k -simplex $\{v_0, \dots, v_k\}$ together with an order type denoted $[v_0, \dots, v_k]$ where a change in orientation corresponds to a change in the sign of the coefficient: $[v_0, \dots, v_i, \dots, v_j, \dots, v_k] = -[v_0, \dots, v_j, \dots, v_i, \dots, v_k]$ if odd permutation is used.

For k larger than the dimension of X , we set $C_k(X) = 0$. The boundary map is defined to be the linear transformation $\partial : C_k \rightarrow C_{k-1}$ which acts on basis elements $[v_0, \dots, v_k]$ via

$$\partial[v_0, \dots, v_k] := \sum_{i=0}^k (-1)^i [v_0, \dots, v_{i-1}, v_{i+1}, \dots, v_k]. \quad (1)$$

This gives rise to a chain complex: a sequence of vector spaces and linear transformations

$$\dots \xrightarrow{\partial} C_{k+1} \xrightarrow{\partial} C_k \xrightarrow{\partial} C_{k-1} \dots \xrightarrow{\partial} C_2 \xrightarrow{\partial} C_1 \xrightarrow{\partial} C_0$$

Consider the following two subspaces of C_k : the **cycles** (those subcomplexes without boundary) and the **boundaries** (those subcomplexes which are themselves boundaries) formally defined as:

- k – cycles: $Z_k(X) = \ker(\partial : C_k \rightarrow C_{k-1})$
- k – boundaries: $B_k(X) = \text{im}(\partial : C_{k+1} \rightarrow C_k)$

A simple lemma demonstrates that $\partial \circ \partial = 0$; that is, the boundary of a chain has empty boundary. It follows that B_k is a subspace of Z_k . This has great implications. The k -cycles in X are the basic objects which count the presence of a “hole of dimension k ” in X . But, certainly, many of the k -cycles in X are measuring the same hole; still other cycles do not really detect a hole at all – they bound a subcomplex of dimension $k+1$ in X . We say that two cycles ζ and η in $Z_k(X)$ are **homologous** if their difference is a boundary:

$$[\zeta] = [\eta] \iff \zeta - \eta \in B_k(X).$$

The k -dimensional *homology* of X , denoted $H_k(X)$ is the quotient vector space

$$H_k(X) := \frac{Z_k(X)}{B_k(X)}. \quad (2)$$

Specifically, an element of $H_k(X)$ is an equivalence class of homologous k -cycles. This inherits the structure of a vector space in the natural way $[\zeta] + [\eta] = [\zeta + \eta]$ and $c[\zeta] = [c\zeta]$ for $c \in \mathbb{Z}_2$.

A map $f : X \rightarrow Y$ is a **homotopy equivalence** if there is a map $g : Y \rightarrow X$ so that $f \circ g$ is homotopic to the identity map on Y and $g \circ f$ is homotopic to the identity map on X . This notion is a weakening of the notion of *homeomorphism*, which requires the existence of a continuous map g so that $f \circ g$ and $g \circ f$ are equal to the corresponding identity maps. The less restrictive notion of homotopy equivalence is useful in understanding relationships between complicated spaces and spaces with simple descriptions. We say two spaces X and Y are **homotopy equivalent**, or have the same **homotopy type** if there is a homotopy equivalence from X to Y . This is denoted by $X \sim Y$.

By arguments utilizing barycentric subdivision, one may show that the homology $H_*(X)$ is a **topological invariant** of X : it is indeed an invariant of homotopy type. Readers familiar with the Euler characteristic of a triangulated surface will not find it odd that intelligent counting of

simplicies yields an invariant. For a simple example, the reader is encouraged to contemplate the “physical” meaning of $H_1(X)$. Elements of $H_1(X)$ are equivalence classes of (finite collections of) oriented cycles in the 1-skeleton of X , the equivalence relation being determined by the 2-skeleton of X .

A.1 Induced homomorphisms

Is it often remarked that homology is functorial, by which it is meant that things behave the way they ought. A simple example of this which is crucial to our applications arises as follows. Consider two simplicial complexes X and X' . Let $f : X \rightarrow X'$ be a continuous simplicial map: f takes each k -simplex of X to a k' -simplex of X' , where $k' \leq k$. Then, the map f induces a linear transformation $f_\# : C_k(X) \rightarrow C_k(X')$. It is a simple lemma to show that $f_\#$ takes cycles to cycles and boundaries to boundaries; hence there is a well-defined linear transformation on the quotient spaces

$$f_* : H_k(X) \rightarrow H_k(X'), \quad f_*([\zeta]) = [f_\#(\zeta)].$$

This is called the **induced homomorphism** of f on H_* . Functoriality means that (1) if $f : X \rightarrow Y$ is continuous then $f_* : H_k(X) \rightarrow H_k(Y)$ is a group homomorphism; and (2) the composition of two maps $g \circ f$ induces the composition of the linear transformation: $(g \circ f)_* = g_* \circ f_*$.

A.2 Building simplicial complexes from the data

How is a simplicial complex built from the data? The basic idea is to take a finite set of points \mathbb{X} with distance function d , together with a parameter ϵ , and construct from it some simplicial complex, for example the **Rips complex**, denoted $R_\epsilon(\mathbb{X})$. The complex will have \mathbb{X} as its vertex set, and a collection $\{x_0, \dots, x_k\} \subset \mathbb{X}$ will span a k -simplex in $R_\epsilon(\mathbb{X})$ if and only if $d(x_i, x_j) \leq \epsilon$ for all $0 \leq i, j \leq k$, where d denotes the metric (distance) which is chosen depending on the problem at hand.

Another possible construction is the **Witness complex**. Given a finite set of points \mathbb{X} equipped with a distance function d , a set of points $L \subset \mathbb{X}$, the **landmark set**, and $\epsilon \geq 0$, we say that a point $x \in \mathbb{X}$ is an ϵ -witness for a $k+1$ -tuple $\{l_0, l_1, \dots, l_k\}$ of points in L if $\max_i d(x, l_i) \leq \epsilon + m_x$, where m_x denotes the $k+1$ smallest value of $d(x, l)$ as l varies over all of L . We next associate a simplicial complex $W_\epsilon(\mathbb{X}, L)$ to \mathbb{X} , L and ϵ , by letting the vertex set of $W_\epsilon(\mathbb{X}, L)$ be L and declaring that a collection $\{l_0, l_1, \dots, l_k\}$ spans a k -simplex in $W_\epsilon(\mathbb{X}, L)$ if and only if there is an ϵ -witness for the collection $\{l_0, l_1, \dots, l_k\}$ and for all its faces. We note that if $\epsilon \leq \epsilon'$, there is an evident inclusion $W_\epsilon(\mathbb{X}, L) \subset W_{\epsilon'}(\mathbb{X}, L)$. Consequently, we have an increasing family of simplicial complexes, parameterized by the real line, just as we did for the Rips complexes. In practice the landmark set is built either by uniform random sampling over \mathbb{X} or by the *max-min* procedure: one first randomly picks a point l_1 from \mathbb{X} . Then, the second point l_2 is chosen so as to maximize $d(l_1, l_2)$. Subsequently, points are chosen to maximize the distance to the set of points already chosen.

Earlier work has shown that this much smaller complex accurately represents topology in simple cases, and we regard it as a computationally tractable proxy for the Rips complex (Carlsson and DeSilva, 2004).

A.3 Persistent Homology: barcodes

Given these definitions, it is clear that ϵ establishes the “spatial scale” of analysis. Assume that \mathbb{X} were sampled from an underlying space X . When ϵ is very small, the result will be a discrete set of points; when ϵ is large, the result will be a single simplex of dimension $\#\mathbb{X} - 1$. However, there is typically a middle range of values of ϵ where $R_\epsilon(\mathbb{X})$ has homology isomorphic to that of the original space, and therefore has Betti numbers equal to those of X . Thus, one of the key concepts below is that the analysis will have to be done for a range of values, from low to high, and investigate those scales where the topological structure remains invariant.

When the space X is a Riemannian manifold, for example, one can explicitly estimate a range of values of ϵ for which this is the case (Niyogi et al., 2006). In our situation we only have the finite sample and no a priori information about the underlying space, therefore, obtaining such estimates is not practical.

Edelsbrunner and colleagues, however, made the following observation (Edelsbrunner et al., 2000). Given $\epsilon \leq \epsilon'$ there is a natural inclusion of simplicial complexes $R_\epsilon(\mathbb{X}) \subset R_{\epsilon'}(\mathbb{X})$, and because of the functoriality property described above, one obtains a linear transformation $H_k(R_\epsilon(\mathbb{X})) \rightarrow H_k(R_{\epsilon'}(\mathbb{X}))$ for any k . What Edelsbrunner et al observed was that in order to study the homology of a given space using a point cloud sampled from it, one should keep track of the entire system of vector spaces $H_k(R_\epsilon(\mathbb{X}))$ along with all the linear transformations described above. Such a system is called a *persistence vector space*. Importantly, it was shown that persistence vector spaces admit a classification analogous to the classification result for finite dimensional vector spaces (Zomorodian and Carlsson, 2004), which asserts that two vector spaces of the same dimension are isomorphic. In the case of persistence vector spaces, it turns out that attached to each persistence vector space, there is an invariant called a **barcode** which is just a finite collection of intervals (perhaps infinite to the right), and that any two persistence vector spaces with the same barcodes are isomorphic.

With computational efficiency considerations in mind one could opt to compute barcodes using the Witness complex construction.

We point out that even the witness complex can become intractable if ϵ is permitted to go to infinity. This is because for sufficiently large ϵ we will construct the full complex with the given number of landmark points. If the set of landmarks is large, this may become intractable as well. For this reason, we introduce a number R_0 associated with a choice of landmark points L , which is the *covering radius* of the set L , defined by $R_0 := \max_{x \in \mathbb{X}} \min_{l \in L} d(x, l)$. In practice we use this as an upper bound for the persistence parameter, and express lengths of persistence intervals as fractions of R_0 . When we have data which is the result from independent repeats of the same experiment, we explore the resulting topological objects obtained by plotting the relative frequency of observation for different topological signatures (sequences of Betti numbers) for different lengths of the persistence interval.

In analyzing both simulated and experimental data we used PLEX, a collection of Matlab functions for computational topology that implements the concepts described above. It is freely available from <http://math.stanford.edu/comptop/programs/>.