Topological analysis of population activity in visual cortex

Gurjeet Singh Institute for Computational and Mathematical Engineering, Stanford University, Stanford, CA, USA

Facundo Memoli Department of Mathematics, Stanford University, Stanford, CA, USA

Tigran Ishkhanov Department of Mathematics, Stanford University, Stanford, CA, USA

Guillermo Sapiro Department of Electrical and Computer Engineering, University of Minnesota, Minnesota, MN, USA

Gunnar Carlsson Department of Mathematics, Stanford University, Stanford, CA, USA

Dario L. Ringach Departments of Neurobiology and Psychology, Jules Stein Eye Institute, David Geffen School of Medicine, University of California, Los Angeles, CA, USA

Information in the cortex is thought to be represented by the joint activity of neurons. Here we describe how 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 applied to the study of population activity in primary visual cortex (V1). We found that the topological structure of activity patterns when the cortex is spontaneously active is similar to those evoked by natural image stimulation and consistent with the topology of a two sphere. We discuss how this structure could emerge from the functional organization of orientation and spatial frequency maps and their mutual relationship. Our findings extend prior results on the relationship between spontaneous and evoked activity in V1 and illustrates how computational topology can help tackle elementary questions about the representation of information in the nervous system.

Keywords: computational topology, spontaneous activity, population coding, high-dimensional data, Betti numbers, persistent homology, natural images

Citation: Singh, G., Memoli, F., Ishkhanov, T., Sapiro, G., Carlsson, G., & Ringach, D. L. (2008). Topological analysis of population activity in visual cortex. *Journal of Vision, 8*(8):11, 1–18, http://journalofvision.org/8/8/11/, doi:10.1167/8.8.11.

Introduction

A key question about the function of the cortex is how sensory input and ongoing cortical activity combine to generate a response to a given stimulus. This issue has become increasingly important due to the recognition that the high variability of responses seen in cortical cells is due, in part, to the presence of ongoing (or spontaneous) activity (Arieli, Sterkin, Grinvald, & Aertsen, 1996; Kenet, Arieli, Grinvald, & Tsodyks, 1997).

Up until recently, it was thought that cortical spontaneous activity was random and unstructured, but we now know this is not the case. First, spontaneous activity in individual cells is tightly coupled to that of a specific population of neurons (or cell assembly) (Song, Sjöström, Reigl, Nelson, & Chklovskii, 2005; Tsodyks, Kenet, Grinvald, & Arieli, 1999; Yoshimura & Callaway, 2005). Second, these patterns of activity appear to be related to the intrinsic connectivity of the cortex (Blumenfeld, Bibitchkov, & Tsodyks, 2006; Goldberg, Rokni, & Sompolinsky, 2004; Kenet, Bibitchkov, Tsodyks, Grinvald, & Arieli, 2003). Theoretical models of V1 demonstrate how structured spontaneous activity could arise in such circuits (Ben-Yishai, Hansel, & Sompolinsky, 1997; Blumenfeld et al., 2006; Bressloff & Cowan, 2003; Hansel & Sompolinsky, 1996; Tsodyks et al., 1999). This is interesting because it suggests that, in principle, there is much to learn about neuronal connectivity in V1 just by observing its spontaneous patterns of activity.

Another intriguing experimental result is that spontaneous cortical activity resembles, in some respects, those observed during natural image stimulation (Fiser, Chiu, & Weliky, 2004; Ikegaya et al., 2004; Kenet et al., 2003;

doi: 10.1167/8.8.11 Received October 23, 2007; published June 30, 2008 ISSN 1534-7362 © ARVO

MacLean, Watson, Aaron, & Yuste, 2005). This might be a consequence of the low average contrast of natural scenes (Chirimuuta, Clatworthy, & Tolhurst, 2003; Clatworthy, Chirimuuta, Lauritzen, & Tolhurst, 2003). Altogether, these findings suggest a model of cortical computation where a relatively weak input effectively serves to push the patterns of activity around in a low-dimensional manifold. It is possible that the structure of this manifold is linked to the structure of the signals one naturally encounters in the environment (Simoncelli & Olshausen, 2001). Thus, natural signals may shape the architecture of V1 and this, in turn, may constrain the dynamics of the network. Our goal is to study some basic aspects of the patterns of activity in V1 evoked by natural images and during spontaneous activity.

Here we concentrate on a topological characterization of population activity in visual cortex. To introduce the topic, and to illustrate how such analysis can be useful, consider the result by Kenet et al. (2003), who found that spontaneous cortical states tend to reproduce the patterns evoked by oriented stimuli. If cortical activity is restricted to patterns evoked by an oriented stimulus, and considering that orientation is a circular variable, the resulting hypothesis is that activity patterns must have a topological structure equivalent to that of a circle (Ben-Yishai, Bar-Or, & Sompolinsky, 1995; Blumenfeld et al., 2006; Goldberg et al., 2004). In other words, the basic question about the structure of the data is topological in nature.

Unfortunately, testing such hypotheses using experimental data is difficult (Goldberg et al., 2004) and up until recently rigorous methods to study the topological structure of data were not available. Furthermore, one would like to develop methods that provide an estimate of the topological structure of a data set without any prior information about what those might be. Again, using the Kenet et al. (2003) study as an example, the authors found that only 20% of the time the cortex was in a state that correlated significantly to that evoked by an oriented stimulus. What is the cortical state during rest of the time? Perhaps, these investigators observed only part of the manifold of cortical activity, while the true underlying object is more complicated, but still low dimensional. One of the contributions of our study is to offer the first estimate of the underlying topological structure of V1 activity.

To anticipate the results, we found that the patterns of spontaneous and driven activity in V1 populations are similarly structured and consistent with the topology of the 2-sphere.

To familiarize the reader with the technique, we begin by introducing the basic ideas at a conceptual level, while a rigorous and detailed description is provided in Appendix A. We validate the method on simulated data by recovering the topological structure of data sets where the "ground truth" is known and study its robustness to changes in spike statistics and the homogeneity of the population. Then, we present our first experimental results from the application of the method to array recordings in macaque primary visual cortex. We conclude by providing a simple model of how such patterns may result from tiling of the cortex by cortical maps and their mutual relationships.

Methods

Conceptual description of topological analysis

Consider a world where objects are made of elastic rubber. Two objects are considered equivalent if they can be deformed into each other without tearing the material. If such a transformation between X and Y exists, we say they are topologically equivalent and write $X \sim Y$; otherwise, we write $X \nsim Y$. This notion of equivalence is illustrated in Figure 1, where the reader is invited to mentally visualize the possible transformations between the various objects to verify the stated equivalence relationships.

From this exercise, it is evident that a possible reason for two objects not to be equivalent is that they differ in the number of holes. Thus, simply counting holes can provide a signature for the object at hand. Holes can exist in different dimensions. A one-dimensional hole is exposed when a one-dimensional loop (a closed curve) on the object cannot be deformed into a single point without tearing the loop. If two such loops can be deformed into one another they define the same hole, which should be counted only once. Analogous definitions can be invoked in higher dimensions. For example, a two-dimensional hole is revealed when a closed two-dimensional oriented surface on the object cannot be deformed into a single point (Figure 2).

This notion of counting holes of different dimensions is formalized by the definition of Betti numbers. The Betti numbers of an object X can be arranged in a sequence, $b(X) = (b_0, b_1, b_2, \ldots)$, where b_0 represents the number of connected components, b_1 represents the number of one-dimensional holes, b_2 the number of two-dimensional holes, and so forth. An important property of Betti sequences is that if two objects are topologically equivalent (they can be deformed into each other) they share the same Betti sequence. One must note, as we will shortly illustrate, that the reverse is not always true: two objects can be different but have the same Betti sequence.

Before proceeding further, it helps to illustrate these notions by means of examples (Figure 2). A single point has a Betti sequence equal to (1, 0, 0, …) since it has only one component and no holes of any dimension (Figure 2a). A circle has a Betti sequence (1, 1, 0, …), as there is a single connected component and there is a single loop that cannot be deformed into a single point (Figure 2b). A

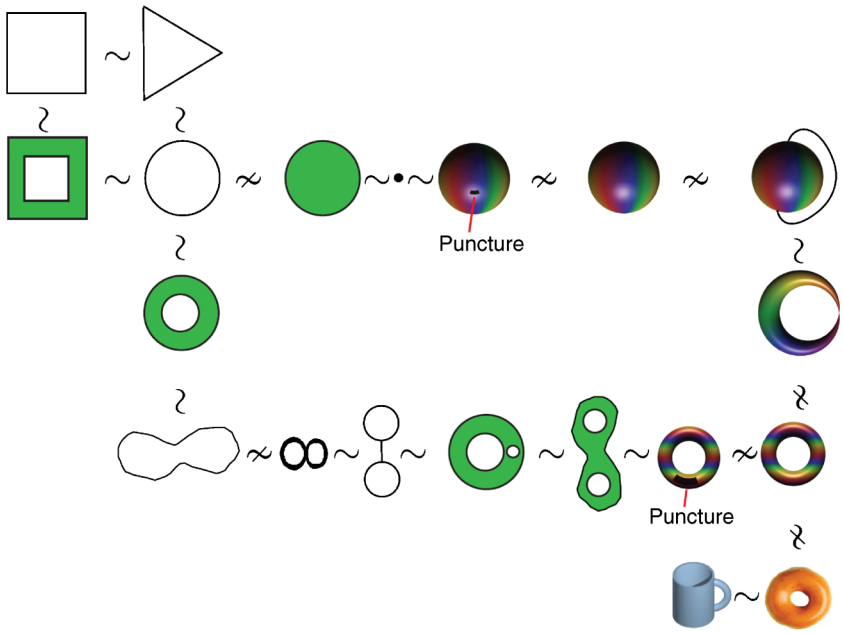


Figure 1. Topological equivalence in rubber-world. The figure illustrates the notion of equivalence by showing several objects (topological spaces) connected by the symbols $\sim$ when they are equivalent or by $\nsim$ when they are not. The reader should think that all the objects shown are made of an elastic material and visualize the equivalence of two spaces by imagining a deformation between to objects.

torus has a Betti sequence (1, 2, 1, 0, …), since it has a single connected component, two different loops that cannot be deformed into a point (shown in red in the bottom panel of Figure 2c), and there is a two-dimensional surface that cannot be deformed into a point (shown in orange in Figure 2c). The Klein bottle has the same sequence as the torus (1, 2, 1, 0, …). This shows that while two objects that are equivalent must have the same Betti sequences, two objects that are not equivalent do not necessarily have different sequences. Finally, a sphere has a sequence (1, 0, 1, 0, …), as any one-dimensional loop on its surface can be deformed into a point. The Betti sequence therefore provides a signature (albeit not unique) of the underlying topology of the object.

These definitions work for smooth continuous objects. But suppose now that instead of a continuous rubbery object we are faced with a finite set of (noisy) points sampled from it, which may represent actual experimental data. How can one estimate the Betti numbers of the original object from these samples? The proposed method

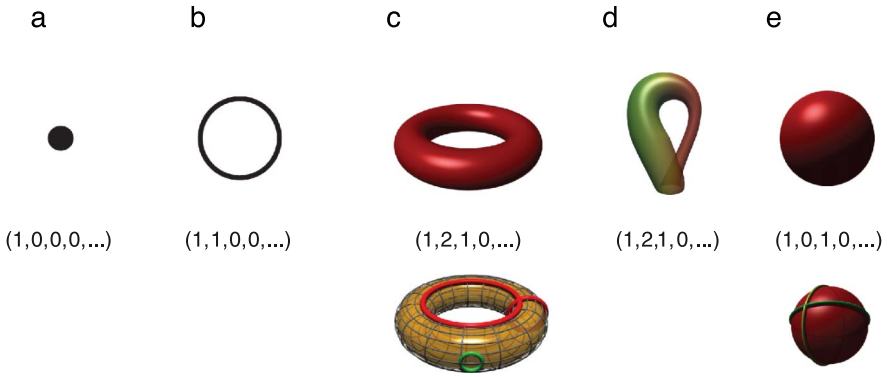


Figure 2. Betti numbers provide a signature of the underlying topology. Illustrated in the figure are five simple objects (topological spaces) together with their Betti number signatures: (a) a point, (b) a circle, (c) a hollow torus, (d) a Klein bottle, and (e) a hollow sphere. For the case of the torus (c), the figure shows three loops on its surface. The red loops are "essential" in that they cannot be shrunk to a point, nor can they be deformed one into the other without tearing the loop. The green loop, on the other hand, can be deformed to a point without any obstruction. For the torus, therefore, we have $b_1 = 2$. For the case of the sphere, the loops shown (and actually all loops on the sphere) can be contracted to points, which is reflected 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).

works by building an object, called the Rips complex, based on a given distance function $d(x, y)$ between any two points.

To build the Rips complex, we begin by selecting a fixed parameter ε, and we proceed to connect all points for which $d(x, y) < \varepsilon$ with edges, all triplets for which all pairwise distances are smaller than ε with triangles, all quadruplets for which all pairwise distances are smaller than ε with tetrahedra, and so on. The Betti numbers are then computed based on the Rips complexes at different values of ε. The parameter ε effectively controls the "spatial scale" of analysis.

The basic idea of the method is to track the different "holes" across different spatial scales of analysis. We visualize the results of the analysis by plotting "birth" and "death" intervals of individual holes of different dimensions as the spatial scale ε goes from zero to infinity. For each Betti number, we keep a separate graph. Connected components are drawn as horizontal lines in the b_0 graph, one-dimensional holes correspond to horizontal lines in the b_1 graph, two-dimensional holes in the b_2 graph, and so on. For each hole, the horizontal line has its endpoints at the values of ε at which the structure was first created and then destroyed. The set of all these lines together is called a barcode.

An example of the computation of the Rips complex at various levels of ε and the corresponding barcode is shown in Figure 3. The data are randomly sampled points from a torus. In each panel, the left three graphs show the barcode obtained from this calculation. The graph on the top corresponds to b_0, meaning that each horizontal line represents a different connected component; the middle graph corresponds to b_1, where each horizontal line corresponds to a one-dimensional loop; the bottom graph corresponds to b_2, where horizontal lines represent two-dimensional holes. The illustrations to the right of the barcode show the state of the Rips complex for the selected value of ε, which is indicated by the red vertical

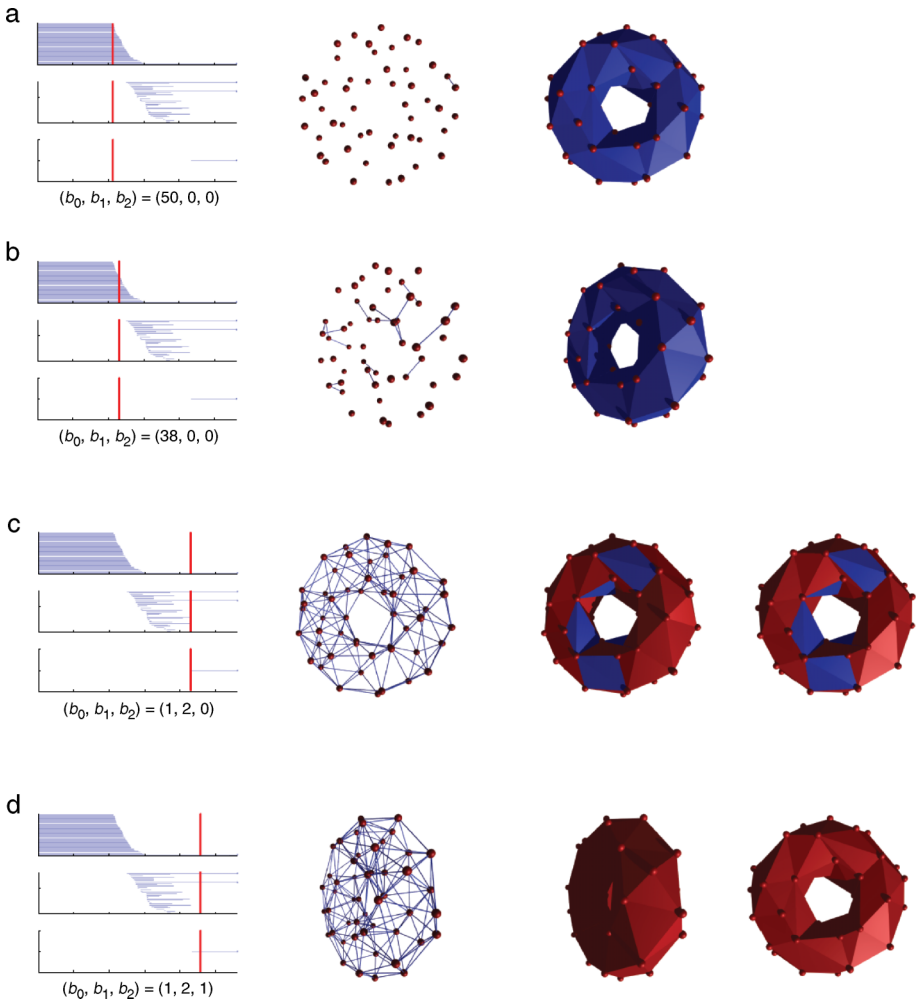


Figure 3. Barcodes and Rips complexes. The figure illustrates the construction of the Rips complex and the generation of barcodes (only the first three Betti numbers are displayed) for 50 points randomly sampled from the surface of a torus. Panels a to d show the barcode "sliced" at different values of ε (the horizontal axis) with the corresponding Rips complexes shown to the right. The corresponding Betti numbers for each level of ε can be obtained by counting the number of horizontal lines crossed by the vertical red line in each graph.

bars in the barcodes. The value of ε increases as we move from Figure 3a to Figure 3d. The blue edges in the Rips complex link the points for which $d(x, y) < \varepsilon$. The panels to the right show red triangles defined by triplets of points for which the pairwise distances satisfy $d(x, y) < \varepsilon$. The blue triangles show triangles that will later be added to the Rips complex at *higher* values of ε but should be considered "holes" for the level at which they are being shown. For simplicity, we do not show higher-order building blocks, such as tetrahedra.

For a small value of ε, only one edge exists and the resulting structure has many different connected components and no holes of any dimension (Figure 3a, top panels). The Betti sequence for this value of ε can be recovered by counting how many horizontal lines (corresponding to different holes) the red vertical line crosses in each of the graphs. The resulting sequence is $(b_0, b_1, b_2) = (50, 0, 0)$. At a higher value of ε, we see more edges being added (thereby reducing the number of connected components) but still no holes of any dimensions (Figure 3b). The corresponding sequence is $(b_0, b_1, b_2) = (38, 0, 0)$. At the next higher value of ε, we finally obtain a single connected component (Figure 3c). The vertical bar in the graph corresponding to b_1 crosses two horizontal lines, meaning that there are two one-dimensional holes. However, the red line does not intersect any horizontal lines in the graph for b_2 meaning that there are no two-dimensional holes at this scale (this is because some of the triangles still need to be filled-in). The Betti sequence for this value of ε is then $(b_0, b_1, b_2) = (1, 2, 0)$. Finally, at a slightly higher value of ε, the correct signature of the torus emerges $(b_0, b_1, b_2) = (1, 2, 1)$ (Figure 3d). This Betti sequence then persists for a long interval of ε.

As is often the case in scale–space methods, a "rule of thumb" is it that a signature $(b_0, b_1, b_2, \ldots)$ that is stable over a "long" interval of spatial scales $(\varepsilon_{low}, \varepsilon_{high})$ is a good candidate for our estimate of the topological structure of the data set. In this case, one would propose $(b_0, b_1, b_2) = (1, 2, 1)$ as our guess for the underlying topology of the space. We will see how a statistical method can be developed to estimate the probability that such signature could have arisen by chance.

Useful animations of the construction of barcodes and Rips complexes for simulated data sets obtained by noisy samples from a circle, sphere and torus are shown in Figure 4. The format of the animation parallels the one in

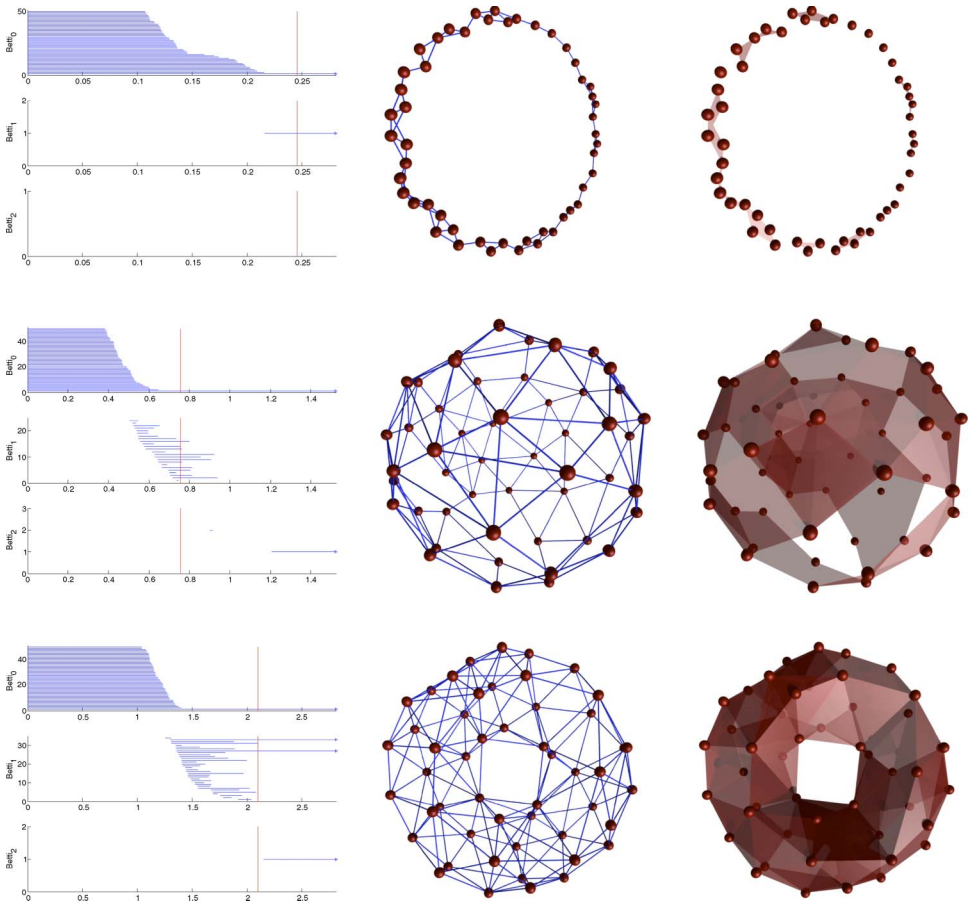


Figure 4. Animations of barcodes and Rips complexes for data points consistent with a circle (top), sphere (center), and torus (bottom). The animation proceeds from low values of ε to high values (that is, from fine to coarse spatial scales). Clicking on each of the figures will link to a movie.

Figure 3. The red bar on the barcode indicates value of ε, which moves from left to right, as the spatial scale goes from fine to coarse over the animation sequence. The middle structure illustrates the data in 3D space (rotating over time for visualization purposes) and the edges (or 1-simplices) being formed as ε increases. The graph on the right shows the triangulation (or 2-simplices) being created as ε increases. In all three examples, it can be seen how at an appropriate scale of analysis, the constructions yield the correct Betti numbers for these objects.

The basic strategy is reminiscent of prior scale–space approaches, such as the analysis of signals by tracking the location of zero-crossings across scale (Florack, Romeny, Koenderink, & Viergever, 1992; Hummel & Moniot, 1989; Yuille & Poggio, 1985, 1986). Specifically, one could attempt an analogy between the barcodes, which track topological structures as a function of scale, and zero-crossing "fingerprints," which track the location of zero-crossings (and their elimination) across scale (Babaud, Witkin, Baudin, & Duda, 1986; Witkin, Terzopoulos, & Kass, 1987; Yuille & Poggio, 1985).

Validation of the method

Before applying the method to multi-unit recordings in primary visual cortex, it seemed prudent to evaluate how the technique performs with simulated data where the underlying topology was known to us. Another objective of these computational studies was to estimate 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. Such information could subsequently be used to design our experiments.

Validation for a circle

As mentioned above, a circle is expected if cortical activity is dominated by neuronal responses to stimulus orientation. We evaluated the method's ability to recover the structure of a circle from the simulation of a population of cells with homogeneous tuning curves and preferred orientations equally spaced around the circle. The tuning curve for the ith cell is given by:

$$\lambda_i(\theta) = r_{\max}\ \exp(\kappa\ \cos(\theta - 2\pi\, i/N))/\exp(\kappa), \qquad (1)$$

for $i = 0, \ldots, N - 1$. Here, N represents the total number of cells in the population, κ determines the bandwidth of tuning, and $r_{\max}$ the maximum spike rate. These rate functions were used to generate spike counts that were Poisson distributed. We selected a value $\kappa = 2$ to match the average tuning as observed experimentally (Ringach, Shapley, & Hawken, 2002) (Figure 5a).

To assess the amount and quality of data required to recover the structure of the data, we ran simulations for various combinations of ($r_{\max}$, N). In all cases, we simulated 100 presentations of 18 orientations equally spaced around the circle. Thus, in all conditions, our data set (or point cloud) consisted of 1800 points (Figure 5b). Given this simulated data, we applied our algorithm to calculate the maximum interval of the parameter ε for which we observe the signature of a circle: (b_0, b_1, b_2) = (1, 1, 0) (Figure 5c). To estimate the likelihood that the result could have resulted by chance, we shuffled the elements of the data matrix in Figure 5b and re-computed the maximal length a total of 100 times. The probability that the measured length was obtained by chance was assessed from this empirical distribution (Figure 5d).

The average probability across 10 realizations of such simulations for each pair yielded $p(r_{\max}, N)$, which is plotted in the pseudo-color image of Figure 6a. 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 fixed significance level. 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 an average of ~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. This dependence indicates that the total number of spikes collected is a key variable controlling the statistical power of the technique.

Validation for a torus

A toroidal representation may arise from a neuronal population responding to two circular variables, such as orientation and color hue. 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

$$\lambda(\theta,\phi) = r_{\max}\ \exp\big(\kappa_\theta\ \cos\theta + \kappa_\phi\ \cos\phi\big)/\exp\big(\kappa_\theta + \kappa_\phi\big). \qquad (2)$$

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 ith cell was given

$$\lambda_i(\theta,\phi) = r_{\max}\frac{\exp\big(\kappa_\theta\ \cos(\theta-\theta_i) + \kappa_\phi\ \cos(\phi-\phi_i)\big)}{\exp\big(\kappa_\theta + \kappa_\phi\big)}, \qquad (3)$$

for $i = 1, \ldots, N$. For this simulation, we used values of $\kappa_\theta = 2$ and $\kappa_\phi = 1.5$. Here, $r_{\max}$ represents the mean

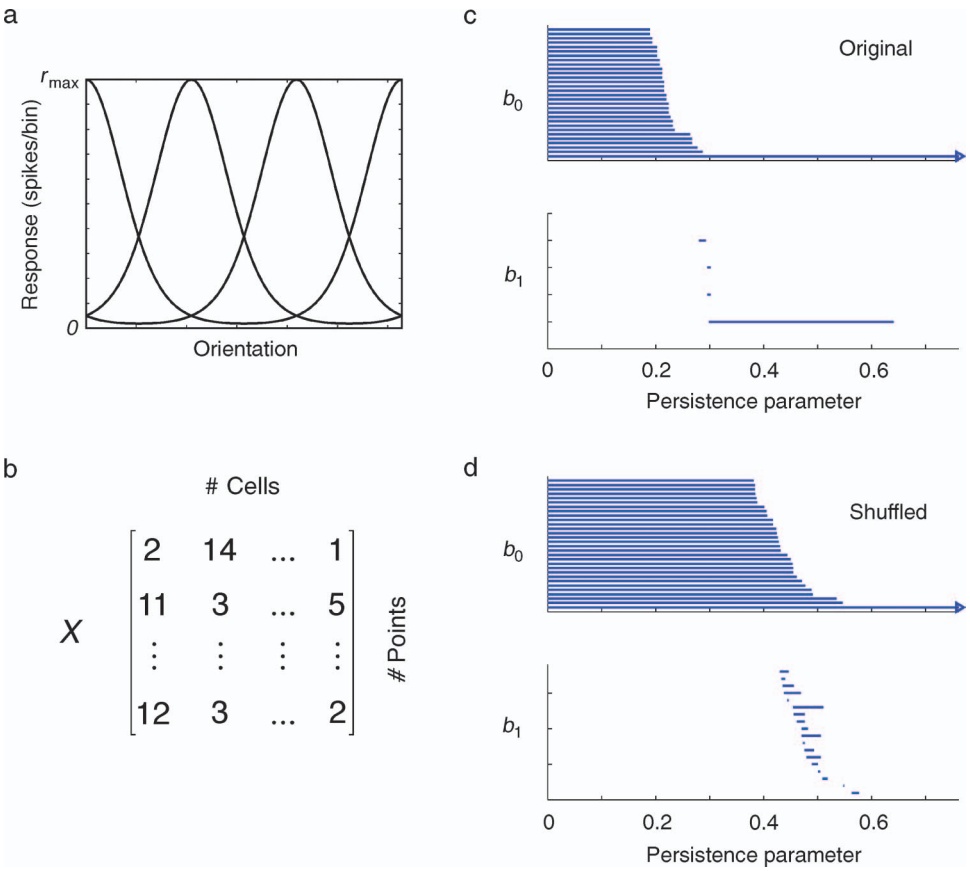


Figure 5. Testing the statistical significance of barcodes. (a) We assume an initial population of Poisson-spiking neurons tuned for orientation. (b) The simulated response of this population to the presentation of different orientations is collected into a data matrix (or point cloud). (c) Analysis of the simulated data shows a long interval with a signature (b_0, b_1) = (1, 1), which correctly identifies the underlying object as a circle. (d) We also compute the barcodes 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 that there the coordinates of the points are independent.

number of spike counts in a time bin. In all cases, we simulated 100 presentations of all 400 stimuli ($\pi k/20$, $\pi l/20$), where $k, l = 0, \ldots, 19$. Thus, in all situations there are a total of 40,000 points in the data set. We then calculated the maximal length of the signature (b_0, b_1, b_2) = (1, 2, 1). 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 computed $p(r_{max}, N)$, which is shown in the pseudo-color image of Figure 6b. The dashed line shows the approximate boundary for detecting a torus at a significance level of 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, the method would require an average of ~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.

Robustness of the method

The above results were obtained for a set of homogeneous cells with the same maximal rates and with Poisson firing rates. To explore the robustness of the technique to different spike count statistics and to heterogeneity in the maximal responses of across the population, we performed three other simulations where these were varied for the case of the circle.

We first simulated a population with Poisson firing but with maximum rates distributed uniformly in the range ($0.5 \times r_{max}$, $1.5 \times r_{max}$) (Figure 7, top). We then used a negative binomial distribution to generate counts with a variance/mean ratio of 1.9 (Kang, Shapley, & Sompolinsky, 2004) but with a uniform rate across the population (Figure 7, middle). Finally, we ran a simulation with non-Poisson firing and non-uniform firing rates (Figure 7, bottom). All the other parameters of the simulation were exactly the same as those in the simulations of Figure 6a. We can see that these manipulations have

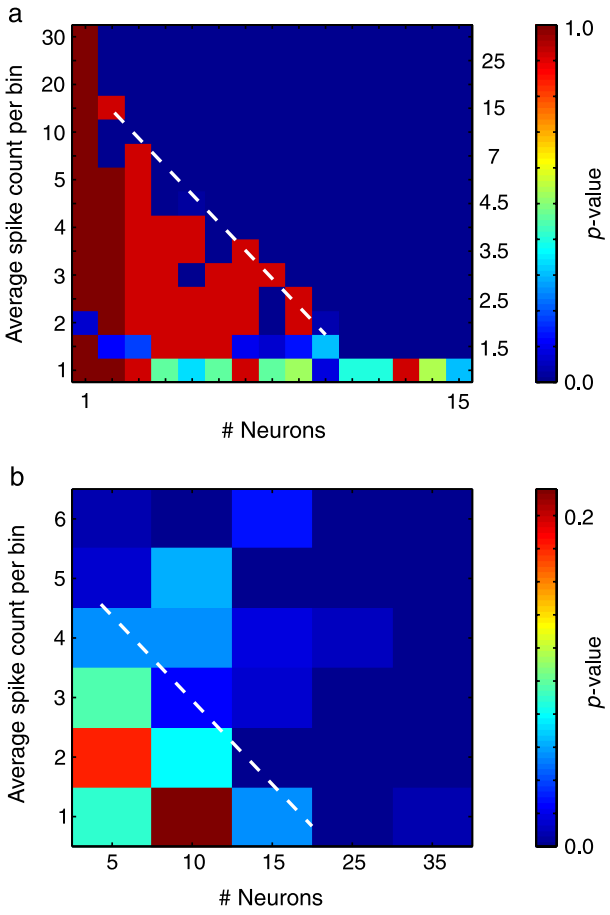


Figure 6. The ability of the method in recovering the underlying structure depends on the mean firing rate and number of cells. Small *p*-values are regions where the algorithm correctly identified the circle (a) and the torus (b), as the likelihood of obtaining such barcodes by chance is very low. There is a trade-off 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 statistical significance.

little effect on our ability to detect the structure of the object, demonstrating the robustness of the technique.

Experimental methods

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*. Experiments were performed on three old-world monkeys (*Macaca fascicularis,* 3.2–4.5 kg). Initially, animals were sedated with acepromazine (30–60 μg/kg) and anesthetized with ketamine (5–20 mg/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 Tobradex (Alcon Laboratories, Texas) and custom-made gas permeable contact lenses.

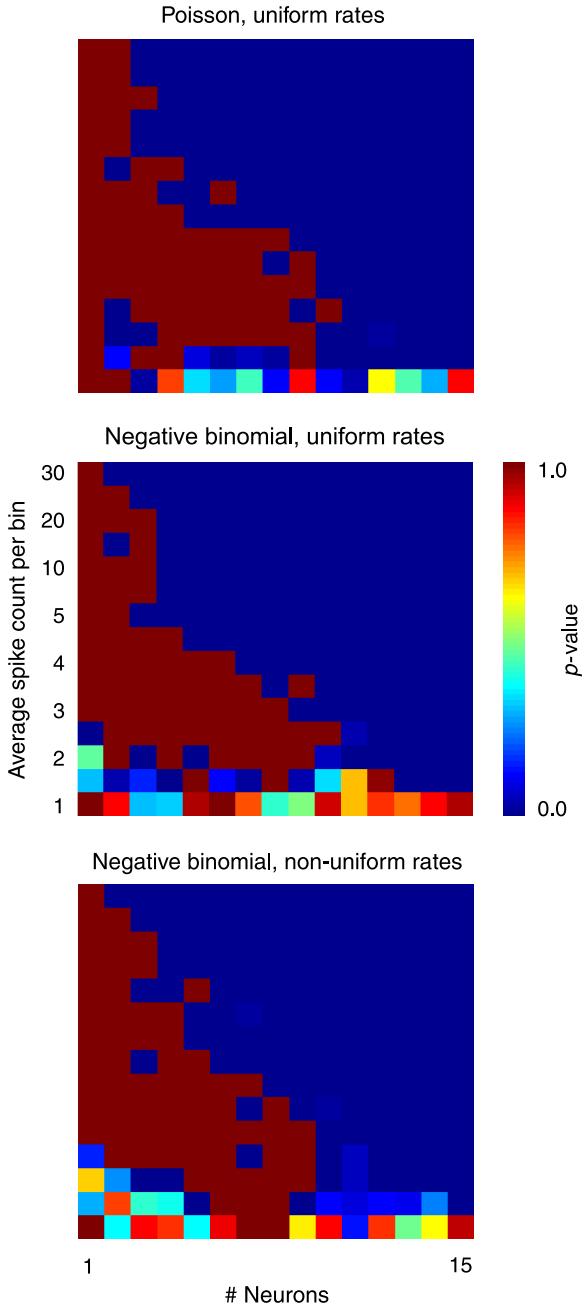


Figure 7. Robustness of the method to changes in count statistics and in the homogeneity of firing rates in the population. Top, simulation of Poisson firing and non-uniform rates. Middle, non-Poisson statistics and uniform rates. Bottom, non-Poisson statistics and non-uniform rates. It can be seen that there is little impact of these in the ability of the method to recover the structure of the embedded circle. Compare with Figure 6a.

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 μg/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 CO_2, SpO_2, 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 37.5°C. Expired CO_2 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.

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 10 × 10 electrodes 1.5 mm long. The distance between neighboring electrodes was 400 μm. 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 O_2 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/m^2. 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.

Visual stimulus

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 O_2 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/m^2. 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.

Creating point clouds

The preparation of the data points for both the spontaneous and driven activity during natural image stimulation was identical. After spike-sorting signals from each electrode we sub-selected a group of 5 neurons that showed the highest firing rates. We worked with such subsets because, in general, the spontaneous rate of spiking in cortical cells is very low. To allow a reasonable comparison between spontaneous and driven activity, we thus concentrated in cells with high spontaneous rate. Otherwise, the selection of the cells was completely unbiased. Subsequent to the channel selection, a point cloud was generated by binning spikes in 50-ms windows. This was selected based on the fact that in voltage sensitive measurements the transitions between states take about 80 ms.

Spontaneous and evoked activity segments were collected in lengths of 10 s each. According to the published data (Kenet et al., 2003), these records should be long enough to allow the detection of structured activity. Each of these segments contained 200 points living in 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 available in http://math.stanford.edu/comptop/programs/. We used a weak witness construction with 35 landmarks points which were selected using the "max–min" procedure (see Appendix A). The "max–min" procedure was seeded with each one of the 200 points in the data set 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.

Statistical significance

To evaluate the probability that the barcodes could have resulted from independent firing of neurons across the electrodes, we generated control data sets as follows. For each channel, we computed first the total number of

spikes. Then, we generated a new data set 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 data set was performed for the control data sets (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, $p < 10^{-10}$).

Fourier analysis

As our data originates from a time series, one must consider that neural activity may be influenced by periodic (or quasi-periodic) physiological processes, such as respiration, heart rate, or vasomotor activity. If neuronal firing rate were to be modulated by these processes they could potentially create spurious one-dimensional loops in the resulting topology. To control for this we search for the presence of peaks in the amplitude spectrum of the same segments of data that formed the basis of our experimental results. We found no statistically significant peaks that would suggest a periodic component in our data. Thus, we rule out the possibility that our findings just reflect some underlying physiological mechanism unrelated to the stimulus.

Results

Topological structure of spontaneous and driven activity in primary visual cortex

Prior studies have revealed similarities between activity patterns in spontaneous and evoked conditions. Here we set out to investigate what topological structure these patterns correspond to.

Our data were obtained from multi-electrode recordings in primary visual cortex of macaque V1 (Figure 8a). Using these electrode arrays, we recorded the population activity in two experimental conditions. In the spontaneous activity condition, population activity was measured when the two eyes were occluded. In the natural image stimulation condition, one eye was opened and we presented a video sequence obtained by sampling different movie clips and the evoked spike trains from the neuronal population were recorded (Figure 8b). A total

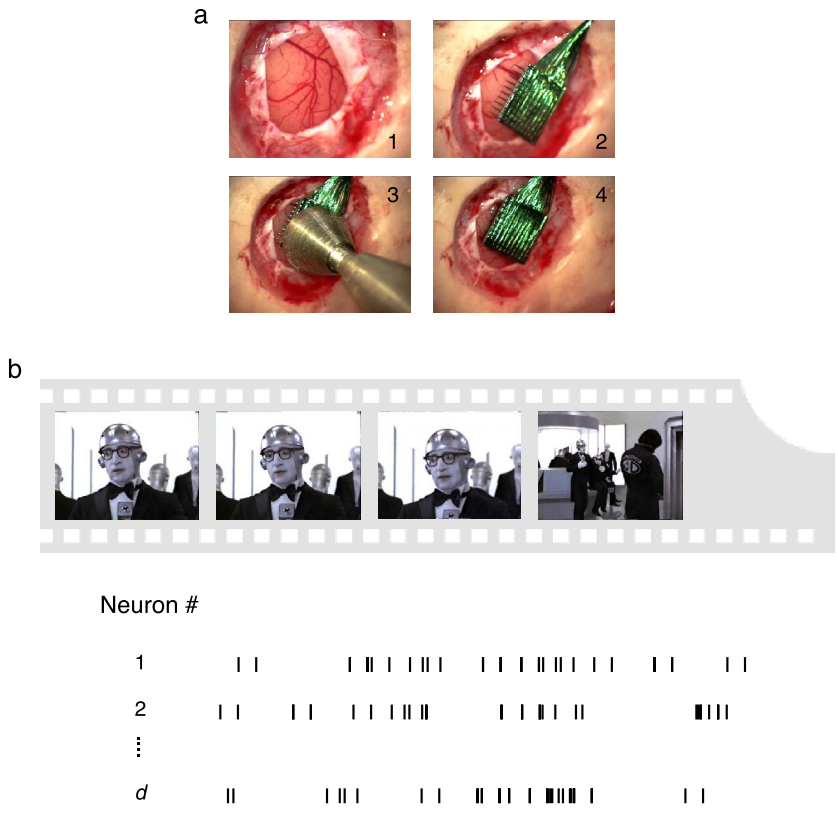


Figure 8. Experimental recordings in primary visual cortex. (a) Insertion sequence of a multi-electrode array into primary visual cortex (Nauhaus & Ringach, 2007). (b) Natural image sequences, sampled from commercial movies, were used to stimulate all receptive fields of neurons isolated by the array. In the spontaneous condition, activity was recorded while both eyes were occluded.

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 Methods. We asked what would be the distribution of the "topological signatures" of the data segments in both conditions.

Figure 9a illustrates the different topological signatures observed in 10 s segments of our data 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 Figure 9b, where the x-axis represents the same ordering of signatures as depicted in Figure 9a. 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 interval 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. Nevertheless, all the topological features shown are statistically different from noise, as Monte Carlo simulations using shuffled data show that the probability of obtaining segments of b_1

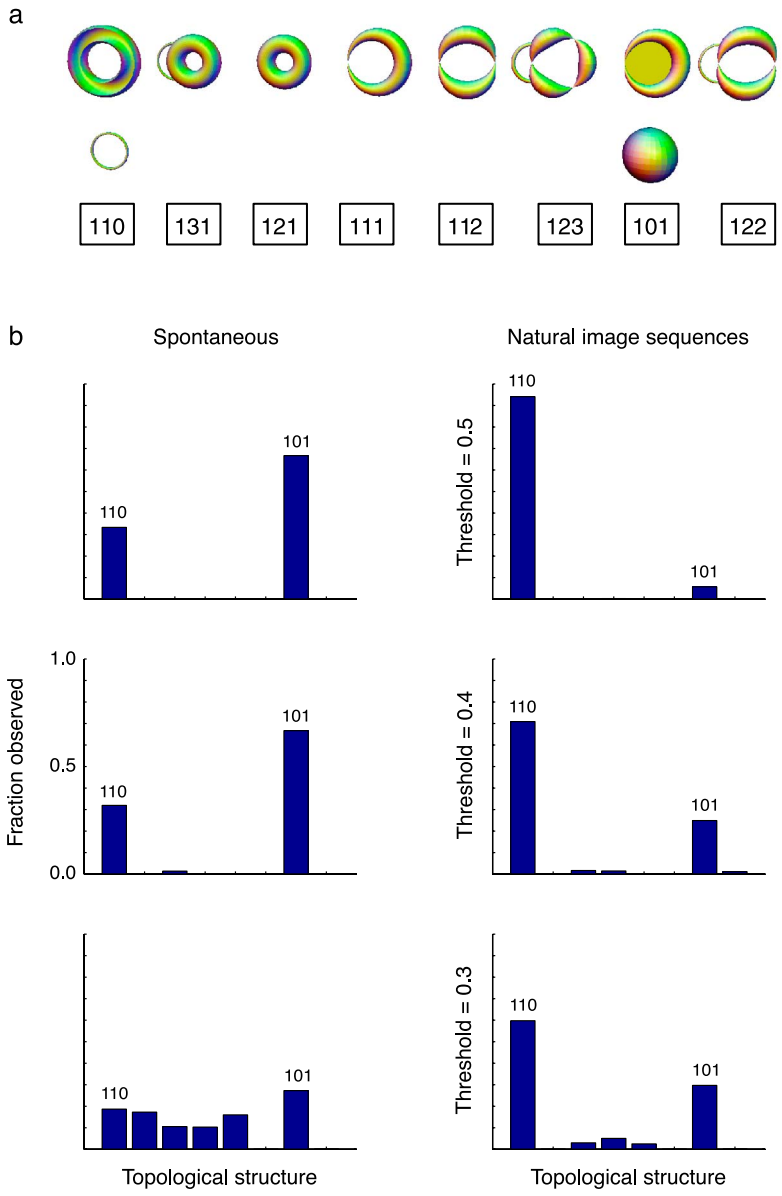


Figure 9. Estimation of topological structure in driven and spontaneous conditions. (a) Ordering of topological signatures observed in our experiments. Each triplet (b_0, b_1, b_2) is shown along an illustration of objects consistent with these signature. (b) Distribution of topological signatures in the spontaneous and natural image stimulation conditions pooled across the three experiments performed. Each row correspond to signatures with a minimum interval length (denoted as the threshold) expressed as a fraction of the covering radius of the data cloud (see Appendix A for the definition of the covering radius).

or b_2 longer than 0.3 by chance (which is the smallest threshold used) was less than 0.005.

We find that 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, population activity patterns in both spontaneous and natural image stimulation have similar topological distributions. However, the relative frequencies of observation of these signatures are 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 (Figure 9b, bottom panels).

Conclusions

We have discussed how 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 patterns of activity, which provides qualitative information about its structure, such as the number of clusters (connected components) and holes of various dimensions. A key concept we used, originating in the work of Edelsbrunner and colleagues (Edelsbrunner, 1998; Edelsbrunner, Letscher, & Zomorodian, 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. These ideas have previously been applied with success in other domains, such in the analysis of natural images (Carlsson, Ishkhanov, DaSilva, & Zomordian, 2008), range images (Adams & Carlsson, 2007), sensor networks (DaSilva & Ghrist, 2007), and the analysis of chaotic signals (Gameiro, Mischaikow, & Kalies, 2004). Here we applied computational topology, for the first time, to the analysis of neuronal population activity in macaque V1.

Through the use of computer simulations we demonstrated that the method works as expected, recovering the underlying structure of the data in artificial data sets 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 that the high-dimensional structures of the data were similar in both cases, consistent with prior results suggesting that natural stimulation modulates ongoing activity only weakly (Arieli, Shoham, Hildesheim, & Grinvald, 1995; Fiser et al., 2004; Grinvald, Arieli, Tsodyks, & Kenet, 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 have similar topological structures, with the signatures of the circle and the sphere dominating the results.

A natural question is how a spherical structure could emerge from what is known about the organization of primary visual cortex. One possibility is that these results from two known facts: (a) the tiling of primary visual cortex by maps of preferred orientation and spatial frequency (Blasdel & Salama, 1986; Grinvald, Lieke, Frostig, Gilbert, & Wiesel, 1986; Issa, Trepel, & Stryker, 2000) and (b) the tendency for extreme spatial frequencies to align with orientation pinwheels (Issa et al., 2000; but see Sirovich & Uglesich, 2004) (Figure 10a).

In addition to these experimental observations, we further assume that when we restrict our view to a single hypercolumn (Figure 10a), neural activity is restricted to a small region of the cortex (Figure 10a, white transparent disk). If the center of this activity profile moves around on the cortical surface, either spontaneously or driven by visual input, the resulting population activity would be equivalent to that of sphere (Figure 10b). Here, extreme

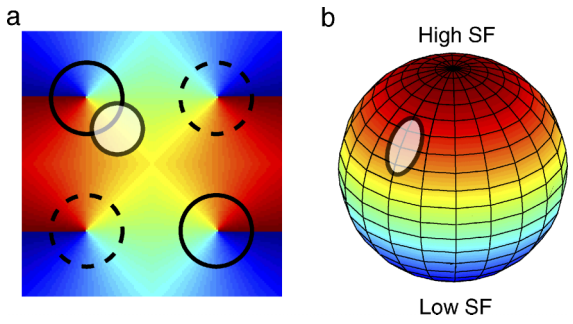


Figure 10. A possible origin of the spherical topological structure in visual cortex. (a) Cartoon organization of orientation and spatial frequency maps in V1. An orientation map tiling V1 is shown along the locations of extremes spatial frequency tuning. Extreme spatial frequency selectivity (solid circles indicates high spatial frequency preference; dashed circles represents low spatial frequency preference) tends to overlap with orientation pinwheels. The white transparent disk indicates a localized region of activity in this hypercolumn. (b) A the activity profile shifts to different locations on the cortex the resulting topological structure of the population response is equivalent to that of a sphere, where extreme spatial frequencies are mapped to the poles and orientation is coded by the azimuth (Bressloff & Cowan, 2003).

spatial frequencies are mapped to the poles and orientation is coded by the azimuth as a result of the spatial relationship between the two maps (Bressloff & Cowan, 2003).

The fact that under natural image stimulation there is a higher tendency for a circular signature may be due to the fact that in short movie segments (10 s long in our analyses) there could be one dominant spatial frequency that effectively restricts the cortical state to travel along one of the circle at a fixed elevation. A possible test of this idea is to analyze *both* the structure of the actual image sequences (Carlsson et al., 2008) and the structure of the cortical responses and check if there is any correlation between the topological signatures across independent trials.

Our findings are consistent with previous voltage-sensitive dye imaging data (Kenet et al., 2003) and, in fact, may offer an explanation for why 80% of the time the cortical state is *uncorrelated* with the orientation maps. If the proposed spherical model is correct, and the cortical 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. An experimental prediction from this model is that if one were to measure orientation *and* spatial frequency maps, when activity is uncorrelated with the orientation maps it must be correlated with the spatial frequency maps (and vice versa). This is readily testable using the methods of Kenet et al. (2003).

At first sight there may appear to be an inconsistency between our results and prior studies in that spike-triggered correlation coefficients in the data of Tsodyks et al. (1999) are Gaussian shaped, while a spherical encoding predicts a uniform distribution (Goldberg et al., 2004). However, we note that we only studied the topological structure of the activity; our results should not be interpreted as implying that the manifold is a sphere in the metric sense.

One limitation of our study is that the data originates mostly from complex cells in the superficial layers. Some of the structure that may relate to spatial-phase selectivity may have gone undetected. It is possible that larger population sizes that included both simple and complex cells may reveal structures more complex than a sphere (a set of gratings with fixed orientation and varying spatial phase generate a Klein bottle; Swindale, 1996; Tanaka, 1995). We are planning a more detailed exploration of these issues using polytrode measurements (Blanche & Swindale, 2006) and two-photon imaging of calcium signals in V1 (Ohki, Chung, Ch'ng, Kara, & Reid, 2005; Ohki et al., 2006).

This study described our first attempt at using topological analysis in the study of brain function. We think the method shows promise and that there may be other areas where the topological analysis of neural activity can help guide further research. First, the technique can be used to test specific hypotheses, such as "is the activity consistent with a single loop?," which is, for example, the key question about the data of Kenet et al. (2003). Second, it provides a rigorous tool to study the phenomenon of cortical "songs," where repeated patterns of activity have been interpreted as activity attractors (Cossart, Aronov, & Yuste, 2003; Deneve, Latham, & Pouget, 2001; Ikegaya et al., 2004; Latham & Nirenberg, 2004; Tsodyks, 1999). The statistical analysis 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, Wiener, Lestienne, & Richmond, 1999; Richmond, Oram, & Wiener, 1999; Wiener & Richmond, 2003). The presence of distinct stable fixed points, or line attractors, is something that could be tested with our methods as well, as they would show up as different connected components in the analysis. Third, topological analysis may be appropriate to explore 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 object shapes is a good example (DiCarlo & Cox, 2007; Edelman, 1998; Feldman & Richards, 1998; Kayaert, Biederman, Op de Beeck, & Vogels, 2005; Kourtzi & DiCarlo, 2006; Tanaka, Saito, Fukada, & Moriya, 1991). Fourth, understanding the topological structure of population activity may help in the design of better decoding methods for use in brain-machine interfaces (Andersen, Musallam, & Pesaran, 2004; Donoghue, 2002; Jazayeri & Movshon, 2006; Nicolelis, 2003; Nicolelis & Chapin, 2002; Ohnishi, Weir, & Kuiken, 2007; Santhanam, Ryu, Yu, Afshar, & Shenoy, 2006; Serruya, Hatsopoulos, Fellows, Paninski, & Donoghue, 2003; Shoham et al., 2005). For example, if one were to know that the activity of a population in a high-dimensional space is equivalent to that of a circle, one can collapse the entire activity to single number (such as the intrinsic distance from a reference data point).

In summary, we have shown that it is possible to apply topological methods in the study of neural activity and learn some new aspects about the structure of the data. More experience with the method is needed, but we expect topological techniques to have a broad impact in the analysis of neural activity across different systems.

Appendix A

Technical description of the method

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 (Hatcher, 2002). Homology is an algebraic procedure for counting holes in topological spaces. There are different variants of homology: we use simplicial homology with Z_2 (binary, 0 or 1) coefficients. This is the reason the Betti sequences for the torus and the Klein

bottle are the same in our calculations (the two objects can be differentiated if the homology is computed over a different field). Given a set of points V, a k-simplex is an unordered subset $\{v_0, v_1, \ldots, v_k\}$ 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 $\{v_0, \ldots, v_{i-1}, v_{i+1}, \ldots, v_k\}$ for some $0 \leq i \leq k$.

Geometrically, the k-simplex can be described as follows: given $k+1$ points in R^m ($m \geq k$), the k-simplex is a convex body bounded by the union of $(k+1)$ linear subspaces of R^m 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 R^n. Let X denote a simplicial complex. Roughly speaking, the homology of X, denoted $H_*(X)$, is a sequence of vector spaces $\{H_k(X)\colon k = 0, 1, 2, 3\ldots\}$, where $H_k(X)$ is called the k-dimensional homology of X. The dimension of $H_k(X)$ is the kth Betti number of X, $b_k(X)$, which is a measurement of the number of different holes in the space X that can be sensed by using sub-complexes 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 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 $\{v_0, v_1, \ldots, v_k\}$ together with an order type denoted $[v_0, v_1, \ldots, v_k]$ where a change in orientation corresponds to a change in the sign of the coefficient: $[v_0, \ldots, v_i, \ldots, v_j, \ldots, v_k] = -[v_0, \ldots, v_j, \ldots, v_i, \ldots, v_k]$ if an 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 $\partial\colon C_k(X) \rightarrow C_{k-1}(X)$, $k \geq 1$, producing the sequence

$$\ldots \xrightarrow{\partial} C_{k+1} \xrightarrow{\partial} C_k \xrightarrow{\partial} C_{k-1} \ldots \xrightarrow{\partial} C_1 \xrightarrow{\partial} C_0. \tag{A1}$$

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

$$\begin{aligned} &k\text{-cycles} : Z_k(X) = \ker(\partial : C_k \rightarrow C_{k-1}) \\ &k\text{-boundaries} : B_k(X) = \mathrm{im}(\partial : C_{k+1} \rightarrow C_k). \end{aligned} \tag{A2}$$

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 sub-complex of dimension $k+1$ in X. We say that two cycles ξ and η in Z_k are *homologous* if their difference is a boundary:

$$[\xi] = [\eta] \quad \leftrightarrow \quad \xi - \eta \in B_k(X). \tag{A3}$$

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

$$H_k(X) := \frac{Z_k(X)}{B_k(X)}. \tag{A4}$$

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 $[\xi] + [\eta] = [\xi + \eta]$ and $c[\xi] = [c\xi]$ for $c \in \mathrm{Z}_2$.

A map $f\colon X \rightarrow Y$ is a *homotopy equivalence* if there is a map $g\colon Y \rightarrow X$ 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 that two spaces X and Y are *homotopy equivalent* or that have the same *homotopy type* if there is a homotopy equivalence from X to Y. This is denoted by $X \sim Y$.

Using arguments based on 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.

Induced homomorphisms

Is it often remarked that homology is functorial, by which it is meant that things behave the way they ought to. 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

$$f_* : H_k(X) \rightarrow H_k(X'), \quad f_*([\xi]) = [f(\xi)]. \tag{A5}$$

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 transformations: $(g \circ f)* = (g* \circ f*)$.

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 X with distance function d, together with a parameter ε, and construct from it some simplicial complex, for example the Rips complex, denoted $R_\varepsilon(\mathrm{X})$. This complex will have X as its vertex set, and a collection $\{x_0, x_1, \ldots, x_k\} \subset \mathrm{X}$ will span a k-simplex in $R_\varepsilon(\mathrm{X})$ if and only if $d(x_i, x_j) \leq \varepsilon$ for all $0 \leq i, j \leq k$, where d denotes the metric (distance) which is chosen depending on the problem at hand.

The construction of $R_\varepsilon(\mathrm{X})$ may be computationally intractable for very large data sets, as it requires the calculation of all pairwise distances between points in the set. Another possible construction is the *Witness complex*. Given a finite set of points X equipped with a distance function d, a set of points $L \subset \mathrm{X}$, the *landmark set*, and $\varepsilon \geq 0$, we say that a point $x \in \mathrm{X}$ is an ε-witness for a $k+1$-tuple $\{l_0, l_1, \ldots, l_k\}$ of points in L if $\max_i d(x, l_i) \leq \varepsilon + m_x$, 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 $W_\varepsilon(\mathrm{X}, L)$ to X, L and ε, by letting the vertex set of $W_\varepsilon(\mathrm{X}, L)$ be L and declaring that a collection $\{l_0, l_1, \ldots, l_k\}$ spans a k-simplex in $W_\varepsilon(\mathrm{X}, L)$ if and only if there is an ε-witness for the collection $\{l_0, l_1, \ldots, l_k\}$ and for all its faces.

We note that if $\varepsilon \leq \varepsilon'$, there is an evident inclusion $W_\varepsilon(\mathrm{X}, L) \subset W_{\varepsilon'}(\mathrm{X}, L)$. 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 X or by the max–min procedure: one first randomly picks a point l_1 from 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 & DeSilva, 2004).

Persistent homology: Barcodes

It is clear from our discussion that ε establishes the "spatial scale" of analysis. Assume that X was 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 #X − 1. However, there is typically a middle range of values of ε where $R_\varepsilon(\mathrm{X})$ 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 for which this is the case (Niyogi, Smale, & Weinberger, 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 (2000), however, made the following observation. Given $\varepsilon \leq \varepsilon'$ there is a natural inclusion of simplicial complexes $R_\varepsilon(\mathrm{X}) \subset R_{\varepsilon'}(\mathrm{X})$, and because of the functoriality property discussed above, one obtains a linear transformation $H_k(R_\varepsilon(\mathrm{X})) \rightarrow H_k(R_{\varepsilon'}(\mathrm{X}))$ for any k. What Edelsbrunner et al. discovered 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 $H_k(R_\varepsilon(\mathrm{X}))$ 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 & 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 $R_0 := \max_{x \in X} \min_{l \in L} d(x, l)$. 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 are the result of 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 (which we referred to as the "threshold" in the body of the article).

Finally, 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/.

Acknowledgments

This work has been supported by DARPA HR0011-05-1-0007 (GC, GS, FM, TI and GS), 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), ARO W911NF-07-1-0473 (GS), NIH-EY 12816 (DLR), and FA8650-06-C-7633 (DLR).

Commercial relationships: none.
Corresponding author: Dario L. Ringach.
Email: dario@ucla.edu.
Address: UCLA Department of Neurobiology, 405 Hilgard Ave., Los Angeles, CA 90095, USA.

References

Adams, H., & Carlsson, G. (2007). On the non-linear statistics of range image patches. Manuscript submitted for publication.

Andersen, R. A., Musallam, S., & Pesaran, B. (2004). Selecting the signals for a brain-machine interface. *Current Opinion in Neurobiology, 14,* 720–726. [PubMed]

Arieli, A., Shoham, D., Hildesheim, R., & Grinvald, A. (1995). Coherent spatiotemporal patterns of ongoing activity revealed by real-time optical imaging coupled with single-unit recording in the cat visual-cortex. *Journal of Neurophysiology, 73,* 2072–2093. [PubMed]

Arieli, A., Sterkin, A., Grinvald, A., & Aertsen, A. (1996). Dynamics of ongoing activity: Explanation of the large variability in evoked cortical responses. *Science, 273,* 1868–1871. [PubMed]

Babaud, J., Witkin, A. P., Baudin, M., & Duda, R. O. (1986). Uniqueness of the Gaussian kernel for scale–space filtering. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 8,* 26–33.

Ben-Yishai, R., Bar-or, R. L., & Sompolinsky, H. (1995). Theory of orientation tuning in visual cortex. *Proceedings of the National Academy of Sciences of the United States of America, 92,* 3844–3848. [PubMed] [Article]

Ben-Yishai, R., Hansel, D., & Sompolinsky, H. (1997). Traveling waves and the processing of weakly tuned inputs in a cortical network module. *Journal of Computational Neuroscience, 4,* 57–77. [PubMed]

Blanche, T. J., & Swindale, N. V. (2006). Nyquist interpolation improves neuron yield in multiunit recordings. *Journal of Neuroscience Methods, 155,* 81–91. [PubMed]

Blasdel, G. G., & Salama, G. (1986). Voltage-sensitive dyes reveal a modular organization in monkey striate cortex. *Nature, 321,* 579–585. [PubMed]

Blumenfeld, B., Bibitchkov, D., & Tsodyks, M. (2006). Neural network model of the primary visual cortex: From functional architecture to lateral connectivity and back. *Journal of Computational Neuroscience, 20,* 219–241. [PubMed]

Bressloff, P. C., & Cowan, J. D. (2003). A spherical model for orientation and spatial-frequency tuning in a cortical hypercolumn. *Philosophical Transactions of the Royal Society of London B: Biological Sciences, 358,* 1643–1647. [PubMed] [Article]

Carlsson, G., & DeSilva, V. (2004). Topological estimation using witness complexes. Paper presented at the Symposium on Computational Geometry.

Carlsson, G., Ishkhanov, T., DaSilva, V., & Zomordian, A. (2008). On the local behavior of spaces of natural images. *International Journal of Computer Vision, 76,* 1–12.

Chirimuuta, M., Clatworthy, P. L., & Tolhurst, D. J. (2003). Coding of the contrasts in natural images by visual cortex (V1) neurons: A Bayesian approach. *Journal of the Optical Society of America A, Optics, Image Science, and Vision, 20,* 1253–1260. [PubMed]

Clatworthy, P. L., Chirimuuta, M., Lauritzen, J. S., & Tolhurst, D. J. (2003). Coding of the contrasts in natural images by populations of neurons in primary visual cortex (V1). *Vision Research, 43,* 1983–2001. [PubMed]

Cossart, R., Aronov, D., & Yuste, R. (2003). Attractor dynamics of network UP states in the neocortex. *Nature, 423,* 283–288. [PubMed]

DaSilva, V., & Ghrist, R. (2007). Coverage in sensor networks via persistent homology. *Algebraic and Geometric Topology, 7,* 339–358.

Deneve, S., Latham, P. E., & Pouget, A. (2001). Efficient computation and cue integration with noisy population codes. *Nature Neuroscience, 4,* 826–831. [PubMed]

DiCarlo, J. J., & Cox, D. D. (2007). Untangling invariant object recognition. *Trends in Cognitive Sciences, 11,* 333–341. [PubMed]

Donoghue, J. P. (2002). Connecting cortex to machines: Recent advances in brain interfaces. *Nature Neuroscience, 5,* 1085–1088. [PubMed]

Edelman, S. (1998). Representation is representation of similarities. *Behavioral and Brain Sciences, 21,* 449–467. [PubMed]

Edelsbrunner, H. (1998). Shape reconstructoin with Delauney complex. Paper presented at the Latin '98: Theoretical informatics.

Edelsbrunner, H., Letscher, D., & Zomorodian, A. (2000). Topological persistence and simplification. Paper presented at the Proceedings of the 41st Annual Symposium on Foundations of Computer Science.

Feldman, J., & Richards, W. (1998). Mapping the mental space of rectangles. *Perception, 27,* 1191–1202. [PubMed]

Fiser, J., Chiu, C., & Weliky, M. (2004). Small modulation of ongoing cortical dynamics by sensory input during natural vision. *Nature, 431,* 573–578. [PubMed]

Florack, L. M. J., Romeny, B. M. T., Koenderink, J. J., & Viergever, M. A. (1992). Scale and the differential structure of images. *Image and Vision Computing, 10,* 376–388.

Gameiro, M., Mischaikow, K., & Kalies, W. (2004). Topological characterization of spatial-temporal chaos. *Physical Review E: Statistical, Nonlinear, and Soft Matter Physics, 70,* 035203. [PubMed]

Goldberg, J. A., Rokni, U., & Sompolinsky, H. (2004). Patterns of ongoing activity and the functional architecture of the primary visual cortex. *Neuron, 42,* 489–500. [PubMed] [Article]

Grinvald, A., Arieli, A., Tsodyks, M., & Kenet, T. (2003). Neuronal assemblies: Single cortical neurons are obedient members of a huge orchestra. *Biopolymers, 68,* 422–436. [PubMed]

Grinvald, A., Lieke, E., Frostig, R. D., Gilbert, C. D., & Wiesel, T. N. (1986). Functional architecture of cortex revealed by optical imaging of intrinsic signals. *Nature, 324,* 361–364. [PubMed]

Hansel, D., & Sompolinsky, H. (1996). Chaos and synchrony in a model of a hypercolumn in visual cortex. *Journal of Computational Neuroscience, 3,* 7–34. [PubMed]

Hatcher, A. (2002). *Algebraic topology*. Cambridge: Cambridge University Press.

Hummel, R., & Moniot, R. (1989). Reconstructions from zero crossings in scale space. *IEEE Transactions on Acoustics Speech and Signal Processing, 37,* 2111–2130.

Ikegaya, Y., Aaron, G., Cossart, R., Aronov, D., Lampl, I., Ferster, D., et al. (2004). Synfire chains and cortical songs: Temporal modules of cortical activity. *Science, 304,* 559–564. [PubMed]

Issa, N. P., Trepel, C., & Stryker, M. P. (2000). Spatial frequency maps in cat visual cortex. *Journal of Neuroscience, 20,* 8504–8514. [PubMed] [Article]

Jazayeri, M., & Movshon, J. A. (2006). Optimal representation of sensory information by neural populations. *Nature Neuroscience, 9,* 690–696. [PubMed]

Kang, K., Shapley, R. M., & Sompolinsky, H. (2004). Information tuning of populations of neurons in primary visual cortex. *Journal of Neuroscience, 24,* 3726–3735. [PubMed] [Article]

Kayaert, G., Biederman, I., Op de Beeck, H. P., & Vogels, R. (2005). Tuning for shape dimensions in macaque inferior temporal cortex. *European Journal of Neuroscience, 22,* 212–224. [PubMed]

Kenet, T., Arieli, A., Grinvald, A., & Tsodyks, M. (1997). Cortical population activity predicts both spontaneous and evoked single neuron firing rates. *Neuroscience Letters,* S27–S27.

Kenet, T., Bibitchkov, D., Tsodyks, M., Grinvald, A., & Arieli, A. (2003). Spontaneously emerging cortical representations of visual attributes. *Nature, 425,* 954–956. [PubMed]

Kourtzi, Z., & DiCarlo, J. J. (2006). Learning and neural plasticity in visual object recognition. *Current Opinion in Neurobiology, 16,* 152–158. [PubMed]

Latham, P. E., & Nirenberg, S. (2004). Computing and stability in cortical networks. *Neural Computation, 16,* 1385–1412. [PubMed]

MacLean, J. N., Watson, B. O., Aaron, G. B., & Yuste, R. (2005). Internal dynamics determine the cortical response to thalamic stimulation. *Neuron, 48,* 811–823. [PubMed] [Article]

Mokeichev, A., Okun, M., Barak, O., Katz, Y., Ben-Shahar, O., & Lampl, I. (2007). Stochastic emergence of repeating cortical motifs in spontaneous membrane potential fluctuations in vivo. *Neuron, 53,* 413–425. [PubMed]

Nauhaus, I., & Ringach, D. L. (2007). Precise alignment of micromachined electrode arrays with v1 functional maps. *Journal of Neurophysiol, 97,* 3781–3789. [PubMed] [Article]

Nicolelis, M. A. (2003). Brain-machine interfaces to restore motor function and probe neural circuits. *Nature Reviews, Neuroscience, 4,* 417–422. [PubMed]

Nicolelis, M. A., & Chapin, J. K. (2002). Controlling robots with the mind. *Scientific American, 287,* 46–53. [PubMed]

Niyogi, P., Smale, M., & Weinberger, S. (2006). Finding the homology of submanifolds with high confidence from random samples. Discrete and Computational Geometry, To appear.

Ohki, K., Chung, S., Ch'ng, Y. H., Kara, P., & Reid, R. C. (2005). Functional imaging with cellular resolution reveals precise micro-architecture in visual cortex. *Nature, 433,* 597–603. [PubMed]

Ohki, K., Chung, S., Kara, P., Hübener, M., Bonhoeffer, T., & Reid, R. C. (2006). Highly ordered arrangement of single neurons in orientation pinwheels. *Nature, 442,* 925–928. [PubMed]

Ohnishi, K., Weir, R. F., & Kuiken, T. A. (2007). Neural machine interfaces for controlling multifunctional powered upper-limb prostheses. *Expert Review of Medical Devices, 4,* 43–53. [PubMed]

Oram, M. W., Wiener, M. C., Lestienne, R., & Richmond, B. J. (1999). Stochastic nature of precisely timed spike patterns in visual system neuronal responses. *Journal of Neurophysiology, 81,* 3021–3033. [PubMed] [Article]

Richmond, B. J., Oram, M. W., & Wiener, M. C. (1999). Response features determining spike times. *Neural Plasticity, 6,* 133–145. [PubMed]

Ringach, D. L., Shapley, R. M., & Hawken, M. J. (2002). Orientation selectivity in macaque V1: Diversity and laminar dependence. *Journal of Neuroscience, 22,* 5639–5651. [PubMed] [Article]

Santhanam, G., Ryu, S. I., Yu, B. M., Afshar, A., & Shenoy, K. V. (2006). A high-performance brain-computer interface. *Nature, 442,* 195–198. [PubMed]

Serruya, M., Hatsopoulos, N., Fellows, M., Paninski, L., & Donoghue, J. (2003). Robustness of neuroprosthetic decoding algorithms. *Biological Cybernetics, 88,* 219–228. [PubMed]

Shoham, S., Paninski, L. M., Fellows, M. R., Hatsopoulos, N. G., Donoghue, J. P., & Normann, R. A. (2005). Statistical encoding model for a primary motor cortical brain-machine interface. *IEEE Transactions on Biomedical Engineering, 52,* 1312–1322. [PubMed]

Simoncelli, E. P., & Olshausen, B. A. (2001). Natural image statistics and neural representation. *Annual Review of Neuroscience, 24,* 1193–1216. [PubMed]

Sirovich, L., & Uglesich, R. (2004). The organization of orientation and spatial frequency in primary visual cortex. *Proceedings of the National Academy of Sciences of the United States of America, 101,* 16941–16946. [PubMed] [Article]

Song, S., Sjöström, P. J., Reigl, M., Nelson, S., & Chklovskii, D. B. (2005). Highly nonrandom features of synaptic connectivity in local cortical circuits. *Plos Biology, 3,* e68. [PubMed] [Article]

Swindale, N. V. (1996). Visual cortex: Looking into a Klein bottle. *Current Biology, 6,* 776–779. [PubMed] [Article]

Tanaka, K., Saito, H., Fukada, Y., & Moriya, M. (1991). Coding visual images of objects in the inferotemporal cortex of the macaque monkey. *Journal of Neurophysiology, 66,* 170–189. [PubMed]

Tanaka, S. (1995). Topological analysis of point singularities in stimulus preference maps of the primary visual-cortex. *Proceedings of the Royal Society B: Biological Sciences, 261,* 81–88.

Tsodyks, M. (1999). Attractor neural network models of spatial maps in hippocampus. *Hippocampus, 9,* 481–489. [PubMed]

Tsodyks, M., Kenet, T., Grinvald, A., & Arieli, A. (1999). Linking spontaneous activity of single cortical neurons and the underlying functional architecture. *Science, 286,* 1943–1946. [PubMed]

Wiener, M. C., & Richmond, B. J. (2003). Decoding spike trains instant by instant using order statistics and the mixture-of-Poissons model. *Journal of Neuroscience, 23,* 2394–2406. [PubMed] [Article]

Witkin, A., Terzopoulos, D., & Kass, M. (1987). Signal matching through scale space. *International Journal of Computer Vision, 1,* 133–144.

Yoshimura, Y., & Callaway, E. M. (2005). Fine-scale specificity of cortical networks depends on inhibitory cell type and connectivity. *Nature Neuroscience, 8,* 1552–1559. [PubMed]

Yuille, A. L., & Poggio, T. (1985). Fingerprints theorems for zero crossings. *Journal of the Optical Society of America A, Optics, Image Science, and Vision, 2,* 683–692.

Yuille, A. L., & Poggio, T. A. (1986). Scaling theorems for zero crossings. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 8,* 15–25.

Zomorodian, A., & Carlsson, G. (2004). *Computing persistent homology*. Paper presented at the ACM Symposium on Computational Geometry, Brooklyn, New York.