

A topological paradigm for hippocampal spatial map formation using persistent homology

Y. Dabaghian^{1,*}, F. Mémoli², L. Frank³, G. Carlsson²

¹Jan & Dan Duncan Neurological Research Institute, Baylor College of Medicine, Houston, TX 77030,
and Dept. of Computational & Applied Mathematics, Rice University, 6100 Main St. Houston, TX
77005, dabaghia@bcm.edu

²Department of Mathematics, Stanford University, Palo Alto, CA 94305
memoli@math.stanford.edu, carlsson@math.stanford.edu

³Department of Physiology, Keck Center for Integrative Neuroscience, University of California,
San Francisco, CA 94143-0444 loren@phy.ucsf.edu

*Corresponding author: dabaghia@bcm.edu

Running title: Topological model of hippocampal map formation (46 characters)

Author Summary

Our ability to navigate our environments relies on the ability of our brains to form an internal representation of the spaces we're in. The hippocampus plays a central role in forming this internal spatial map, and it is thought that the ensemble of active "place cells" (neurons that are sensitive to location) somehow encode metrical information about the environment, akin to a street map. Several considerations suggested to us, however, that the brain might be more interested in topological information—i.e., connectivity, containment, and adjacency, more akin to a subway map— so we employed new methods in computational topology to estimate how basic properties of neuronal firing affect the time required to form a hippocampal spatial map of three test environments. Our analysis suggests that, in order to encode topological information correctly and in a biologically reasonable amount of time, the hippocampal place cells must operate within certain parameters of neuronal activity that vary with both the geometric and

topological properties of the environment. The interplay of these parameters forms a “learning region” in which changes in one parameter can successfully compensate for changes in the others; values beyond the limits of this region, however, impair map formation.

Abstract

An animal’s ability to navigate through space rests on its ability to create a mental map of its environment. The hippocampus is the brain region centrally responsible for such maps, and it has been assumed to encode geometric information (distances, angles). Given, however, that hippocampal output consists of patterns of spiking across many neurons, and downstream regions must be able to translate those patterns into accurate information about an animal’s spatial environment, we hypothesized that 1) the temporal pattern of neuronal firing, particularly co-firing, is key to decoding spatial information, and 2) since co-firing implies spatial overlap of place fields, a map encoded by co-firing will be based on connectivity and adjacency, i.e., it will be a topological map. Here we test this topological hypothesis with a simple model of hippocampal activity, varying three parameters (firing rate, place field size, and number of neurons) in computer simulations of rat trajectories in three topologically and geometrically distinct test environments. Using a computational algorithm based on recently developed tools from Persistent Homology theory in the field of algebraic topology, we find that the patterns of neuronal co-firing can, in fact, convey topological information about the environment in a biologically realistic length of time. Furthermore, our simulations reveal a “learning region” that highlights the interplay between the parameters in combining to produce hippocampal states that are more or less adept at map formation. For example, within the learning region a lower number of neurons firing can be compensated by adjustments in firing rate or place field size, but beyond a certain point map formation begins to fail. We propose that this learning region provides a coherent theoretical lens through which to view conditions that impair spatial learning by altering place cell firing rates or spatial specificity.

INTRODUCTION

In order for an animal to be able to navigate a space, remember its route, find shortcuts, and so forth, it must have a fairly sophisticated internal representation of the spatial environment. This internal map is made possible by the activity of pyramidal neurons in the hippocampus known as place cells. Place cells are so named because of their striking spatial selectivity: as an animal (in experiments, typically a rat) explores a given environment, different place cells will fire a series of action potentials in different, discrete regions of the space. Each region, referred to as that cell's "place field," is defined by the pattern of neuronal firing (most intense at the center and attenuated toward the edges of the field) (**Figure 1**)—elsewhere, the cell remains silent [1].

The mechanism of this selectivity (why a place cell fires when the rat is here rather than there) is opaque, and how the ensemble of place cells forms a hippocampal map of the environment is only slightly less mysterious. It is believed that the ensemble of place cells activated in a given environment produces a sufficient number of place fields to cover the animal's vicinity [2]: indeed, a rat's path through a small space can later be re-traced with a high degree of accuracy by recording hippocampal spiking activity during its explorations and then analyzing the location, size, and firing rates of a mere 40-50 place fields [3-5]. Such experiments suggest that the information contained in place cell firing patterns encodes spatial navigation routes and somehow represents the spatial environment. The hippocampal map thus seems to form the basis of the animal's spatial memory and spatial cognition [6].

How does the brain convert the pattern of neuronal firing into an approximation of the surrounding space? And what information is most important to navigation and spatial memory? In theory, the mental map could represent metrical information (distances and angles), affine aspects (colinearity or parallels), or topological information (connectedness, adjacency, containment). The reigning paradigm is that the maps encode geometric information: in fact, most efforts to analyze cognitive maps derived from place fields are based explicitly on the geometry of both the place fields and the environment [2,4,6,7]. But this can not be how the hippocampus or neurons receiving hippocampal output decode place cell firing, because the brain has no direct access to the place fields mapped by neuroscientists. To understand what algorithms the brain might use to decode hippocampal place cell firing, then, we should rely solely on the information provided by place cell spiking activity [8,9].

If we restrict ourselves to cell spiking activity, the temporal features of the firing pattern become paramount: in particular, if spatial location is the primary determinant of each place cell's firing, then contemporaneous activity or co-firing of several place cells implies that the corresponding place fields overlap. It is, in fact, generally assumed that neurons downstream of the hippocampus interpret place cell spiking patterns based on co-firing. What is not often appreciated, however, is that if place cell co-firing

implies spatial overlap of place fields, then the map formed by co-firing is going to be based on connectivity, adjacency and containment—in other words, it will be a topological, rather than a geometric, map.

Indeed, the way place fields cover an environment calls to mind a basic theorem of algebraic topology: if one covers a space X with a sufficient number of discrete regions, then it is possible to reconstruct the topology of space X from the pattern of the overlaps between the regions [10]. We propose that the hippocampus actually does create a connectivity map derived from place cell co-firing patterns. Although we do not imply any specific interpreting mechanism, we propose nevertheless that it is possible to derive spatial information from place cell firing, with specific implications (and quantifiable predictions) for the qualities of the hippocampal spatial map. We hypothesize that the connectivity map is topological, i.e., any finite structure of overlaps between spatial regions, as represented by temporal overlap of spike trains, can be realized using regions of different shapes or sizes. One implication of this hypothesis is that the information contained in the spike trains is qualitative in nature and can be studied using topological techniques. This is not to deny that the hippocampal connectivity map could contain additional space encoding mechanisms for geometric information (scale, distances, angles)—this question would have to be answered experimentally. Nevertheless, a number of experiments [11-14], have demonstrated that smooth geometric variations of the environment produce continuous stretches of the place field layouts that preserve the relative timing between spikes, so that the temporal pattern of spiking remains largely invariant with respect to geometric transformation. This provides some experimental support for our mathematical intuition.

Here we investigate whether a topological connectivity map can be effectively and reliably derived from neuronal spiking patterns using computational tools recently developed in the field of algebraic topology. We show that there exist certain requirements for the firing activity to produce a stable topological map and that the experimentally observed characteristics of firing activity likely satisfy these requirements.

RESULTS

We will first outline the key concepts underlying our approach; more precise mathematical explanations are provided in the **Methods** section for interested readers.

Our topological framework: simplicial complexes, spatial and temporal

In algebraic topology, the topological features of a space X are defined by its topological invariants, i.e., those properties of the space that are invariant to applied transformations. Topological

invariants are described via indices, the simplest of which are the so-called Betti numbers that formalize the counting of loops and holes in various dimensions. The zeroth Betti number, $b_0(X)$, counts the connected components in the space X ; $b_1(X)$ gives the number of one-dimensional (1D) loops, $b_2(X)$ the number of two-dimensional (2D) loops, and so on (see **Methods** and [10]). The Betti numbers can be calculated by an algorithm that analyzes the “cover” of a space X by an ensemble of discrete regions [15]. This algorithm uses “nerve of the cover” or “nerve simplicial complex,” $\mathcal{N}(X)$, which has as many vertices as there are regions used to cover the space X . If two regions overlap, the corresponding vertices, say, v_i and v_j , are considered connected by a 1D bond v_{ij} . If three regions overlap, then three bonds, v_{ij} , v_{jk} , and v_{ki} , support a 2D triangular facet, and so on, as the number of overlaps and bonds increase. The complex $\mathcal{N}(X)$ obtained from a sufficiently dense cover of the space X will reproduce the correct topological indices of X (see **Methods** for a more precise definition of “sufficiently dense”). The structure of the simplicial complex thus approximates the structure of the environment (see **Methods**).

Drawing on this concept of the nerve simplicial complex $\mathcal{N}(X)$, whose structure can be used to deduce the structure of a space, we devised a temporal analogue, the temporal simplicial complex \mathcal{T} that should give a complete topological description of a space X in terms of place cell co-firing. The difference is that the structure of the temporal simplicial complex \mathcal{T} unfolds over time, $\mathcal{T} = \mathcal{T}(T)$: as the animal explores its environment and more place cells fire (and co-fire), the structure of the complex $\mathcal{T}(T)$ grows as T increases with the number of spikes. It should thus be possible to trace the emergence of topological information as more and more spikes are fired. When the animal is first introduced to the environment, there would be few data points from place cell spiking, but the data would accumulate as the animal explores, enabling the formation of an internally consistent topological map. Eventually, after a certain minimal time \bar{T}_{min} , the structure of $\mathcal{T}(T)$ should saturate, and its topological characteristics should stabilize and produce the correct topological indices, which would indicate the completeness of the topological information.

Given a certain experimental, phenomenological or theoretical description of place cell firing, it should be possible to trace the accumulation of topological information with T and discover how much time \bar{T}_{min} is required to produce the correct topological signature of a particular environment. The temporal framework we propose does impose certain requirements, however: just as there must be a sufficient number of place fields covering a space X in order to produce a coherent map of that space, we predict that there are certain conditions that must be met by place cell activity if we are to be able to rely solely on the temporal overlap between neuronal spike trains. First, there should be sufficient co-firing of place cells. Second, the place cells should have sufficient spatial specificity (though there will be a certain amount of biological noise). Third, there should be a realistic learning period in which true signals can be distinguished by their persistence beyond biological noise. These criteria may or may not be met by the

hippocampus, given the high variability of biological systems [5], but they follow from our topological hypothesis. In this paper we use newly developed methods from Persistent Homology theory (see **Methods**) to test our hypothesis and to study the temporal dynamics of topological information emerging from the spikes.

Selecting parameters on which to base the model

The parameters that might be taken into account to define place cell activity are numerous and complex: there are biophysical variables (firing rates, spike amplitude, etc.), behavioral variables (the animal's running speed, etc.), out-of-field firing (not all place cell firing is for spatial encoding purposes), and so forth. For the sake of simplicity, at least for this first attempt to model place cell ensemble behavior, we zeroed in on just a few key parameters that will still enable us to ask key questions.

First we had to decide how to define temporal overlap between spike trains. There is some conjecture in the field that each theta cycle—the basic EEG cycle in the hippocampus, with a frequency of ~ 8 Hz—defines a temporal unit of processing [16,17]. This suggested to us that we might describe neuronal activity in terms of firing rates, defined over time bins comparable to the theta cycle. Absent any data that speaks directly to this, we defined co-firing as firing that occurred over two consecutive theta cycles.

For this initial analysis, we ignored the details of the spike train structure, such as spike bursting [18] and phase precession [19-21] focusing instead on the total number of cells (N), the firing rate (f), and a computationally derived place field size (s). (In other words, whereas place fields are typically created by mapping recorded neuronal firings onto an actual rat's trajectory through a particular environment, our algorithm ascribes size on the basis of the spread of simulated data points. See **Methods**.)

Thus, for an N -cell ensemble this approach produces $3N$ independent parameters, f_1, f_2, \dots, f_N and $s_{x,1}, s_{x,2}, \dots, s_{x,N}, s_{y,1}, s_{y,2}, \dots, s_{y,N}$. Since we are interested in the behavior of the cell ensemble and not just the firing rate of a single place cell or the size of an individual place field, however, we define the values f_i and s_i by the probability distributions $P(f)$ and $P(s)$. **Fig. 2** shows a typical shape for these distributions derived from experimental data collected on a linear track (unpublished data, Y.D. and L.F.); these data can be naturally fit by a log-normal distribution with a certain mean and a certain standard deviation, (\bar{f}, σ_f) and (\bar{s}, σ_s) , respectively. We further assumed that spiking dynamics are attributable solely to the rat's movement through the environment, i.e., that the probability distributions $P(f) = P(f|\bar{f}, \sigma_f)$, and $P(s) = P(s|\bar{s}, \sigma_s)$ do not depend on time. It is important to note, however, that place

fields can be highly plastic during the first few minutes in a new environment [22], so our estimates of the time required to build an accurate topological map are likely to be lower bounds.

Given these starting assumptions and simplifications, the individual firing rates f_i and place field sizes s_i are simply random variables drawn from the stationary distributions $P(f)$ and $P(s)$, so the total number N of cells can be treated as another independent parameter characterizing the ensemble. With N included, the activity regime of a place field ensemble is specified by just five parameters: \bar{f} , σ_f , \bar{s} , σ_s and N . To further reduce the number of parameters in our model, we capitalized on the fact that the mean and the standard deviation of the distributions $P(f)$ and $P(s)$ can be compared to actual experimental data and are thus not arbitrary (see **Fig. 2**). We can therefore avoid overly broad or overly narrow distributions $P(f)$ and $P(s)$ by imposing the additional conditions $\sigma_f = a\bar{f}$ and $\sigma_s = b\bar{s}$, (the higher the individual firing rates and place field sizes found in the ensemble, the wider the spread) and selecting the proportionality coefficients a and b so that the shape of the distributions $P(f)$ and $P(s)$ mimics those derived from the experiments. In our computations we used $a = 1.2$ and $b = 1.7$ (as in **Fig. 2**). This last simplification reduces the number of parameters to just three— \bar{f} , \bar{s} , and N —which gives us a 3D parameter space that can be readily simulated and visualized.

Given the temporal nature of our map formation model, we will adopt one more simplifying assumption, namely, that all the instances of co-activity that occur between $T = 0$ and $T = T_{min}$ are “remembered” and can be used to establish the structure of the temporal complex $\mathcal{T}(T)$. Clearly, any “forgetting” mechanism would cause the temporal complex to deteriorate; information provided by new spike trains could compensate for this loss, but transience of data would increase the map formation time T_{min} .

The model

The foregoing considerations led to the following (very simplified) working model of place cell activity:

1. Place cell firing activity is a stationary Poisson process described by the rate model [1] (see **Methods**). Theta oscillations, bursting and other effects are not considered.
2. Two cells are considered to be co-active if they fire within two consecutive periods of theta oscillations, i.e., within $\sim 1/4$ sec. We expect shorter time windows would require longer periods for map formation, so this value helps us establish a lower bound on the length of time required to extract connectivity information.
3. The firing rate amplitudes (f_i) and the computed place field sizes ($s_{x,i}, s_{y,i}$) of the cells in the ensemble are described by independent probability distributions $P(\mu, \sigma)$, where we used the mode

μ (to identify the peak of the distribution) and the standard deviation σ (e.g., the log-normal or the gamma distributions). In our simulations we used log-normal distributions with $\sigma_f = a\bar{f}$ and $\sigma_s = b\bar{s}$, $a = 1.2$ and $b = 1.7$.

4. Retained memory assumption: all firing events occurring up to time T can be used in the analysis. We ignore place cell firing that occurs because of the internal dynamics of the hippocampal network (such as reactivation of past experience [23]).

Our analysis is based on the dynamics of "cycles," objects that can be used to count the number of topological holes within the temporal complex $\mathcal{T}(T)$ (**Fig. 3, Methods**). The intuition informing our approach is that, at early stages of exploration, only a few co-firings will have occurred and so the complex will not adequately represent the topological structure of the environment. As the rat begins to explore an environment at time $T = 0$, the temporal complex will consist mostly of 0-cycles, marking the cells that have fired but not necessarily co-fired. As the rat continues to explore the environment, the co-firing cells will produce links between the vertices of $\mathcal{T}(T)$, and higher dimensional cycles will appear. **Fig. 3** shows the cycles in each dimension as a function of time: each horizontal bar represents the timeline of a particular cycle in the complex $\mathcal{T}(T)$.

At any time T , a vertical section will encompass the timelines of all the cycles that are present in $\mathcal{T}(T)$ at that moment. Once born, a cycle remains stable over a certain period of time, but as T increases, most cycles in each dimension will disappear as so much "topological noise," leaving only a few persisting cycles that express stable topological information. The beauty of this Persistent Homology method [24] (see **Methods**) is that it accommodates such noise so we can distinguish between cycles that persist across time (reflecting real topological characteristics) and transient cycles produced by the rat's behavior (e.g., circling in a particular spot during one trial).

The time required for the correct number of bars (cycles) to appear in every dimension is, by design, the time required to extract the correct topological signature of the environment, which can thus be interpreted as the minimal time \bar{T}_{min} required for the rat to learn the environment. Since this procedure can be applied to various place cell ensembles with different firing profiles, the Persistent Homology method allows us to determine how \bar{T}_{min} varies with the parameters of hippocampal firing activity. In effect, each set of parameters will produce a "barcode" that can be "scanned" in order to discern the topological structure of the environment.

Map formation depends on hippocampal state

We simulated map formation times using different place cell parameters and three separate planar

2 × 2 meter areas with 1 or 2 holes (**Fig. 4, top row**). (We chose this size because it is similar to that of experimental spaces used for neuroscientific studies of rat place cell properties.) The topological features to be detected are the number of holes in each environment. For the 2-hole environments we considered two cases with different hole sizes to vary the geometrical characteristics of the environment while keeping its topology fixed.

The trajectories were simulated to be: 1) sufficiently ergodic to represent non-preferential exploratory spatial behavior (i.e., there was no artificial circling or other ad hoc favoring of one segment of the environment over another). The spatial occupancy of the immediate vicinities of the holes and of the corners was therefore higher than the average, which is similar to patterns of spatial occupancy in the open field and linear track experiments. 2) The mean and the maximal speed were kept within the range of typical experimental values (based on our experience; the mean speed was chosen to be slightly higher than a typical experimental mean value in order to get a lower estimate for the learning time \bar{T}_{min}). Lastly, 3) the distribution of the moment-to-moment changes in the direction of the simulated rat's movement, $\Delta\varphi_{sim}$, matches the experimental histogram of $\Delta\varphi_{exp}$.

We asked whether, and for which ensembles, the place cell spiking signals would be able to produce a temporal simplicial complex with the correct number of topological loops (Betti numbers; see **Methods**) in every dimension—the connectivity of the space (0D loops) and the correct number of 1D loops. In each of these environments we simulated a set of 1000 place cell ensembles by independently varying three parameters of neuronal activity. We probed ten distributions of firing rates, with the mean maximal rate, \bar{f} , ranging from 2 to 40 Hz, and ten distributions for the place field sizes with \bar{s} ranging between ~10 and ~90 cm. The size of the population varied independently from $N = 50$ to $N = 400$ cells. In each case, the centers of the place fields were scattered randomly and uniformly over the environment. For each combination of the parameters \bar{f} , \bar{s} , and N —which we can say defines a hippocampal ‘state’—the computation was repeated 10 times (total $10 \times 10 \times 10 \times 10 = 10,000$ ensemble simulations), which allowed us to compute the average time \bar{T}_{min} required for the emergence of the correct topological signature for each specific choice of the ensemble parameters, \bar{f} , \bar{s} , and N . Although the simulated trajectory was fixed, we chose a new set of place field centers for each set of \bar{f} , \bar{s} , and N for each repetition.

The results are shown in **Fig. 4**. The three panels across the second row show the mean map formation times \bar{T}_{min} (in minutes) for each of the three environments in the top row. Each point on this diagram represents a particular place cell ensemble with a certain mean ensemble firing rate \bar{f} , mean place field size \bar{s} , and number of cells N . The sizes of the dots represent the percentage of repetitions in which a given set of parameters (\bar{f} , \bar{s} , and N) produced the correct outcome: the largest dots correspond to

the most successful ensembles, and the sizes of the smaller dots represent the percentage of trials producing the correct outcome for that set of parameters. The color of the dots represents the value of the mean map formation time \bar{T}_{min} (see Figure legend). Some ensembles consistently produced the correct topological signature for all 10 repetitions, even in very short time frames (large blue dots), whereas other ensembles either produced the correct signature in only a fraction of repetitions (smaller dots) or repeatedly failed to produce the correct result over long periods of time (smallest red dots; $\bar{T}_{min} \gtrsim 30$ mins, i.e., $\sim 20,000$ theta cycles).

These data illustrate, first, that the firing activity of smaller place cell ensembles ($N \leq 150$), characterized by low mean firing rates ($\bar{f} \leq 10$ Hz) and by small mean place field sizes ($\bar{s} \leq 20$ cm), consistently failed to produce the correct topological characteristics of the environment. Similarly, ensembles with very large place field sizes (low spatial selectivity of the place cell's firing, high \bar{s} values) also failed to produce the correct topological signature. Both types of cases are represented by the dots on the periphery of each cloud (**Fig 4, second row**). In contrast, larger place cell ensembles with higher firing rates and well-tuned place fields reliably captured the topological structure of the environment within 2-5 mins. As a result, each point cloud can be conceived as containing within its fuzzy boundaries a learning region L : a submanifold in the space encompassing the hippocampal states that produce the correct topological map within a biologically plausible time-frame. This pattern is clarified by the 2D sections of the 3D diagram (**Fig. 5** and Suppl. Figs. 1 and 2).

It is noteworthy that the points with intermediate sizes, representing the partially failing ensembles, tend to diffuse out from the center of L to the sparser boundary region of the cloud. This neatly illustrates the transition that occurs between the hippocampal states that consistently produce stable, topologically accurate maps (interior points of L), and those that do not (dots outside of L). Thus, for all hippocampal states within L , the \bar{T}_{min} values show an orderly, regular dependence on all three variables \bar{f} , \bar{s} , and N . Despite the stochastic nature of the model, then, the minimal map formation time \bar{T}_{min} can be approximated by a well-defined, continuous function of the parameters, $\bar{T}_{min} = \bar{T}_{min}(\bar{f}, \bar{s}, N)$. If the firing activity regime moves out of L , then the time \bar{T}_{min} abruptly increases at the boundary of this region.

The map formation region is stable and robust

It is noteworthy that at the core of L , the characteristic minimal map formation time is $\bar{T}_{min} \approx 2 - 5$ mins, which is comparable to the biological learning time in rats and mice in simple environments [25,26]. Indeed, the characteristic time \bar{T}_{min} is shorter than the time it takes the trajectories themselves to cover environments A, B, and C (see **Fig. 6**); in other words, the topological model forms maps more

rapidly than simply computationally covering the simulated space. As noted above, one of the key hypotheses of our model is that map formation time \bar{T}_{min} is included in the biological learning time. We reasoned that before the “topological noise” stabilizes, it is not possible to tell how many correct loops there will be or which ones are going to persist (see **Fig. 3**), so that prior to \bar{T}_{min} the spatial information encoded by place cell firing is unstable and probably incomplete. Therefore, if the spatial map produced by hippocampal activity is based on interpreting the co-firing patterns, one of the main qualitative predictions of this approach is that the biological learning time can be estimated by \bar{T}_{min} . If, for example, the map formation time for a place cell ensemble in Rat A is \bar{T}_{min}^A , and for a place cell ensemble in Rat B is \bar{T}_{min}^B , and $\bar{T}_{min}^A > \bar{T}_{min}^B$, then Rat A will take longer to learn an environment than Rat B. This difference should be observed in the Morris water maze task and other behavioral experiments.

Figure 4 also shows that the size of the learning region L depends on the complexity of the environment. L is largest for environment C, in which topological connectivity is defined by the quasi-linear order of place cell firing. The region L for environment B is the most compact, reflecting the fact that this environment is topologically the most complex because the navigational paths are indexed by two topological indices (defined by the fundamental group π_1), so that two persistent loops have to be extracted from a set of non-persisting loops. At the same time, it is also more geometrically complex than the quasi-linear environment (C), because it allows $2D$ motion.

It is also important to note that the mean map formation time \bar{T}_{min} produces a stable, robust core of the learning region L . To characterize the stability of \bar{T}_{min} values, we considered the standard deviation of the minimal map formation times, $\Delta\bar{T}_{min}$, computed for each fixed value of \bar{f} , \bar{s} , and N . The third row of panels in **Fig. 4** demonstrates the relative standard deviation of the map formation times, $\Delta T_{min}/\bar{T}_{min}$ as a function of \bar{f} , \bar{s} , and N . The variations in map formation times increase towards the boundary of the learning region, i.e., the place cell ensembles defined by parameters at the boundary of L are successful only a fraction of the time—these ensembles are “unstable” in their map formation ability, or we might say the maps themselves are unstable. (One can imagine a rat with some impairment being unable to learn a space because its mean place cell firing rate is a little too low to produce consistent information about the test environment.) Inside L , the values of \bar{T}_{min} vary less. To emphasize this, the subregion $R_* \subset R$ of points $x = (\bar{f}, \bar{s}, \text{ and } N)$ in which the relative standard deviation of the T_{min} values is less than 15% ($\Delta T_{min}/\bar{T}_{min} \leq 0.15$) (**Fig 4, third row**).

Finally, in order to single out the hippocampal states for which the mean map formation time \bar{T}_{min} is not only biologically plausible but also robust, we combined the requirements $\bar{T}_{min} < 30$ mins (used to build **Fig. 4ii**) and $\Delta T_{min}/\bar{T}_{min} \leq 0.15$. The resulting robust learning region at the core of L (**Fig. 4iv**) shows that the parameters of the place cell ensemble that guarantee reliable map formation form a

well-defined submanifold in the parameter space. This submanifold can be computed using the methods outlined above for each specific environment and for each specific model of neuronal firing, allowing us to relate the geometric and the topological features of the space with the biological parameters of hippocampal place cell activity.

In summary, it is usually assumed that an ensemble of cells with spatially selective firing will naturally encode a spatial map. Our results demonstrate that the spatial selectivity of firing does not, by itself, guarantee a reliable mapping of the actual environment. The geometric shape of the learning region \mathbf{L} and the distribution of the \bar{T}_{min} values within \mathbf{L} depend on the global geometry and topology of the environment. This means that place cells cannot be ‘agnostic’ about the scope and nature of the spatial encoding task: the geometry of the environment sets limits on the parameters of neuronal activity that are able to lead to a coherent topological map. Despite the stochastic nature of the system, well-defined mean map formation times \bar{T}_{min} not only exist inside of the stable learning region \mathbf{L} , but their values can be approximated by a continuous function of the place cell ensemble statistics. The latter implies that a continuous variation of the hippocampal state within \mathbf{L} will result in a continuous change of the mean map formation time value \bar{T}_{min} . The hippocampus can thus change its operating state inside \mathbf{L} without compromising the integrity of the topological map, such that the size and the shape of \mathbf{L} reflect the scope of the biological variability that the hippocampus can afford in a given environment. The larger the region \mathbf{L} , the more stable the map.

DISCUSSION

We have examined the dynamics of hippocampal spatial map formation beginning with arbitrary place cell activity regimes, both those that resemble biological cells and those that do not. We created a computational program to simulate map formation with three independent variables: the firing rate of the place cells, the size of the place field, and the number of cells. We then tested the model on three different scenarios (which included two topological configurations and two different geometries), and repeated the simulation in each scenario 10 times prior to statistical analysis. Our simulations show that in order to form a reliable topological map of the environment, the place cell ensemble must operate within certain parameters—outside these parameters, place cells can be spatially specific but will not be able to produce a reliable map. It is noteworthy that the parameters for place cell firing and place field size that produced a robust map formation region \mathbf{L} correspond well with experimentally observed place cell firing rates and place field sizes: when the simulated place cells fired at rates either below or above a certain range, or when the simulated place field sizes fell above or below a certain range, what we call the learning region

failed to form (**Figure 5, Suppl. Figures 1 and 2**). Mathematically, the model could have required any set of values to work: firing rates of 500 Hz and place field sizes of 2 cm, for example. There was no *a priori* reason that the parameters should fall so neatly into biological range. The fact that they do lends support to our topological paradigm, despite the simplicity of this first model.

Other parameters and models of place cell behavior

Our current model relies on a simple spike train structure based on a Gaussian firing rate (**Methods**). We ignored many biological parameters of place cell activity, such as synaptic connections, theta phase precession and spike bursting, and out-of-field firing. The ensembles we used were also rather small, ~400 neurons, which is less than 1% of the number of cells that are believed to be active in a rat's hippocampus during its exploration of a new environment [27] and about 2% of that number for mice [28]. Larger numbers of cells can be incorporated into future versions of the model, which will lead to a more realistic description of the hippocampal spatial map. Although we expect that the quantitative predictions of the model will change as more subtle neurophysiological phenomena are included, we do not anticipate that the overall structure outlined in the current, basic model will change qualitatively. For example, preliminary analyses suggest that the phenomenon of theta precession and multiply connected place fields affect map formation time, i.e., the size and the shape of region L_{model} but do not change the fact of L 's existence or the existence of the function $\bar{T}_{min} = \bar{T}_{min}(\bar{f}, \bar{s}, \text{and } N)$.

One could conceivably choose any valid set of parameters to define hippocampal states that produce a model-defined learning region L_{model} . The result will correspond to the actual, biological place cell map only to the extent that the starting model accurately captures relevant aspects of place cell activity. For example, the Continuous Attractor Neural Network Models [29], which includes (among other things) synaptic efficacies, could be tested for the topological completeness and robustness of the map that it produces. In the absence of exact knowledge about place cell activity in a specific animal, the structure of L_{bio} can be studied using statistically defined (experimental or model-generated) characteristics of neuronal activity. The approach we have outlined here is thus one means of testing the efficiency of other place cell activity models in forming spatial maps.

The topological model predicts that: (1) the parameters describing the hippocampal place cell map in healthy animals should fall inside of the stable learning region L computed for the given environment, and (2) the hippocampal state might drift towards the boundary of stability or even leave the stability region as a result of a deterioration of neuronal activity. As long as the parameters used in the model are phenomenological characteristics of neuronal firing, the structure of the learning region L_{model} will define the effect that a particular parameter has on the hippocampal place cell map as a whole. With

this approach, there is no need for *a priori* assumptions about place cell firing rates or the parameters that define place field sizes. Instead, the correct values can be estimated based on the structure of the computed stability region. The fact that the values typically observed in electrophysiological recording experiments fall within the region of stability shown in **Fig. 4** is a testament to the validity of the grounding assumptions of our model as outlined above.

Implications of the topological model for spatial learning.

Despite its simplifications, the current model allows us to examine whether a particular set of place cell parameters can be used to map a given environment and vice versa, and to reason about the effect of the geometry and topology of an environment on place cell behavior. For example, **Figure 4** demonstrates that greater topological complexity reduces the size of the stable learning region L by constraining the range of hippocampal states capable of forming accurate maps in a reasonable amount of time. This is, in fact, what is observed in experiments that include a large number of objects (enriched environments) and are thus more geometrically and topologically complex than the standard environments: the firing rates and the number of active cells tend to increase [30,31], and the place fields become more sharply tuned [6].

Although the current model does not describe the formation of place fields themselves, it provides some insight into the process of learning in novel environments. Place fields show considerable plasticity over the course of learning new environments, expanding in adaptation to large environments [32] or over the course of several days of learning (with a concomitant decrease in the number of place cells firing at high rates) [33]. (As presciently noted by Shen et al. [34] in a study of aging rats, the expansion of place fields increases the amount of place field overlap, which can encode more information, at least up to a point.) Furthermore, Karlsson et al. [33] reported that a stable high rate cell population ($\bar{f} > 25 \text{ Hz}$) emerges over the course of learning a new environment. More specifically, while the overall population firing rate diminishes with learning, the spatial specificity of a small proportion of active cells increases, while neurons that are weakly spatially tuned are suppressed. This is precisely the sort of compensation within L predicted by our model: the hippocampus is free to adopt the most efficient parameters within the learning region once a space is learned, and map formation remains stable.

Indeed, perhaps the most striking aspect of the current study is not that it supports the hypothesis that the hippocampus encodes topological information about the environment, but that the learning region L , which reflects the scope of biological variability that the hippocampus can afford in a given environment, is rather large. Given the importance of spatial navigation, and thus spatial map formation, to the lives of most animals, it is not surprising that there should be such a wide range of possible firing

rates, place field sizes, or cell numbers capable of forming a map of a simple space. (Lose the ability to navigate reliably, and one's lifespan shortens dramatically.) Our model would predict that a degenerating brain that is losing place cells might initially compensate by upregulating firing rate and that such compensation might take place for quite some time before function is noticeably impaired.

Numerous studies have documented spatial learning deficits and changes in place field characteristics in mice bearing specific genetic mutations, but the connection between behavioral changes and the changes in place field properties has been unclear. We suggest that significant alterations of place cell behavior result in hippocampal states hovering at or beyond the boundaries of L that cannot consistently support spatial learning. In mouse models of Alzheimer disease (AD), for example, the place fields are larger (less spatially specific), the firing rates lower, and the number of active cells smaller [35,36]. We speculate that the hippocampal map in AD does not do its job because the parameters of place field activity fall outside the core of the learning region L_{bio} and therefore cannot reliably encode spatial information. Similarly, acute ethanol intoxication causes place fields to lose their specificity temporarily suppresses place cell firing rate in a dose-dependent manner [37], and the place fields concomitantly lose their spatial specificity [38]; according to our proposed model, the lowest doses of ethanol do not compromise the rat's navigational ability because they allow the place cells still to operate within the learning region. Our model could thus help shed light not only on the process of learning in novel environments, but also on how such abilities can be lost.

METHODS

We open this section by outlining the assumptions we made about place fields and place cells in this first attempt at a model of hippocampal spatial map formation. We then define key theoretical concepts from algebraic topology that motivated our particular computational approach, particularly relating to the relatively new tools of Persistent Homology theory.

The three environments and simulations of rat trajectories. Each experimental environment depicted in the top row of Figure 4 is 2 meters square. The hole in scenario A is $50\text{ cm} \times 50\text{ cm}$; the two holes in scenario B are $50\text{ cm} \times 50\text{ cm}$ and $50\text{ cm} \times 1\text{ m}$, respectively; both holes in scenario C are $180\text{ cm} \times 80\text{ cm}$. We simulated rat movement to have a mean speed of 25 cm/sec (with a range from 0 cm/sec to 50 cm/sec) and designed the trajectories to mimic how a rat moves in actual open field experiments: moment by moment, the animal's head changes position by some amount $\Delta\varphi_{sim}$, and we reconstructed the histogram of the $\Delta\varphi_{exp}$ distributions from recorded trajectories so they match the bimodal distribution as found in [39].

Simulating place cell firing. For this initial analysis, we ignored the details of the spike train

structure, such as spike bursting [18] and phase precession [19], and used the simplest cell firing model based on the time rescaling theorem (modeling spiking as ‘‘Poisson noise’’) [40]. In this approach, place cell firing is represented by an inhomogeneous Poisson process with a time-independent rate function $\lambda = \lambda(\vec{r})$, which is a function of the animal's position, \vec{r} and which produces stochastic firing around place field centers. These Gaussian place fields are not characterized by sharply defined boundaries. Our model thus allows for noise from ‘‘erroneous spikes’’ that may connect PFs in one case and not connect them in another. In the simplest case that is commonly used for place cell activity modeling (cf. [3]), the firing rate λ_i of an individual cell c_i is modulated by a single peak 2D Gaussian function,

$$\lambda_i(\vec{r}) = f_i e^{-\frac{(\vec{r}-\vec{r}_i^c)^2}{2s_i^2}},$$

centered around the point \vec{r}_i^c (the center of the i th place field) with the variance $s_i^2 = (s_{i,x}^2, s_{i,y}^2)$.

Assumptions regarding place fields. In this model, we assume that: 1) Place fields are ellipsoid and omni-directional, as typically recorded in open field environments [41] and are derived from a 2D Gaussian rate function (see below.) 2) Place fields within each given simulation are stable, i.e., $P(s)$, the probability distribution of place fields does not change over time. 3) The place field sizes vary according to a log-normal distribution, the center being s -mean (cf. Fig. 2). Because our simulations varied both the number of place cells and the size and shape of place fields, we expect that some combinations of parameters tested—e.g., a combination of low cell number N and small mean place field size s —will fail to cover the experimental space, while a large N should cover the area uniformly. Indeed, we found that low N and s -mean produced high rates of map formation failure (see Figures 4 and 5).

Simplicial complexes. Simplicial complexes are used to approximate the structure of topological spaces [42]. For example, a tetrahedron, as a simplicial complex, consists of four triangular facets, six linear edges, and four points in Euclidean space. Each one of these elements is by itself a smaller simplex; this hierarchy is captured in the notion of an abstract simplicial complex, in which the tetrahedron is thought of simply as a set of 4 elements, and any of its 3-element subsets corresponds to a facet, any subset of two corresponds to a segment or edge, and any subset of one corresponds to a vertex. Therefore, given a set of vertices 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 facets 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 Euclidean space 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). Any abstract simplicial complex on a (finite) set of points V has a geometric realization in some

R^n .

It can be shown that topological features, e.g., holes in the environment, correspond to loops in the simplicial complex, which can be detected through combinatorics of the simplices. It is possible to determine, for example, whether two points in the complex are connected by a sequence of edges or not. The simplicial complex produced by the overlaps between the place fields covering the environment is known in algebraic topology as the "nerve of the cover" or the "nerve simplicial complex" $\mathcal{N}(x)$ [10,43,44]. We use the abstract simplicial complex to interpret the pattern of temporal overlaps between the place cell spike trains.

Homology theory. The hypothesis that drives this project is that the hippocampus encodes a topological map. To begin our investigation we ask whether the topological map produced by the place cells captures the most basic topological features of the environment, namely, the number of holes in it. This question can be addressed using homology theory, which aims to detect homologous loops and to categorize holes in a space. Since the structure of the nerve simplicial complex approximates the structure of the environment, we can use homology theory to count the loops in the simplicial complex and therefore the number of holes in the environment.

There are numerous variants of homology: we use simplicial homology with \mathbb{Z}_2 coefficients (the algebraic system consisting of the Boolean values 0 and 1, equipped with "and" as the multiplication and "exclusive or" as addition).

Betti numbers and Homology groups. Let Σ denote a simplicial complex. Roughly speaking, the homology of Σ , denoted $H_*(\Sigma)$, is a sequence of vector spaces $H_k(\Sigma)$, $k = 0, 1, 2, \dots$, where $H_k(\Sigma)$ is called the k -dimensional homology of Σ . The dimension of $H_k(\Sigma)$, called the k th Betti number of Σ , is a coarse measurement of the number of different k -dimensional structures, e.g., "loops" in Σ , that cannot be collapsed or deformed into one another (see **Fig. 7**). For example, the simplest basis for $H_0(\Sigma)$ consists of a choice of vertices, one in each path-component of Σ . Hence the dimension of $H_0(\Sigma)$ is equal to the number of connected components of Σ . Likewise, the simplest basis for $H_1(\Sigma)$ consists of looping sequences of 1D edges in Σ , which surround *holes* in Σ . For example, if Σ is a 1D graph, then the space $H_1(\Sigma)$ encodes the number and types of loops in the graph.

Cycles, boundaries and homotopies. For each $k \geq 0$, let $C_k(\Sigma)$ be the vector space whose basis is the set of *oriented* k -simplices of Σ ; that is, k -simplices $\{v_0, \dots, v_k\}$ 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 an odd permutation is used. For k larger than the dimension of Σ , we set $C_k(\Sigma) = 0$.

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

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

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

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

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

A simple lemma demonstrates that $\partial \circ \partial = 0$; that is, the boundary of a chain has an empty boundary. It follows that B_k is a subspace of Z_k . This has significant implications. Just as 1D loops on graphs, the k -cycles in Σ are the basic objects which count the presence of "holes of dimension k " in Σ . Certainly, many of the k -cycles in Σ 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 Σ . We say that two cycles ζ and η in $Z_k(\Sigma)$ are **homologous** if their difference is a boundary:

$$[\zeta] = [\eta] \leftrightarrow \zeta - \eta \in B_k(\Sigma).$$

The k -dimensional **homology** of Σ , denoted $H_k(\Sigma)$ is then the quotient vector space $H_k(\Sigma) = Z_k(\Sigma)/B_k(\Sigma)$.

Specifically, an element of $H_k(\Sigma)$ 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 $c \in \mathbb{Z}_2$. The k -th **Betti number** of Σ is then formally defined as the dimension of the k -dimensional homology group:

$$b_k(\Sigma) = \dim H_k(\Sigma)$$

A map $f: \Sigma \rightarrow Y$ is a **homotopy equivalence** if there is a map $g: Y \rightarrow \Sigma$ so that $f \circ g$ is homotopic to the identity map on Y and $g \circ f$ is homotopic to the identity map on Σ . 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.

By arguments utilizing barycentric subdivision, one may show that the homology $H_*(\Sigma)$ is a **topological invariant** of Σ : 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 simplices yields an invariant. For a simple example, the reader is encouraged to contemplate the "physical" meaning of $H_1(\Sigma)$. Elements of $H_1(\Sigma)$ are equivalence classes of (finite collections of) oriented cycles in the 1-

skeleton of Σ , the equivalence relation being determined by the 2-skeleton of Σ .

Building simplicial complexes from the spike data: Moving from spatial overlap of place fields to temporal overlap of spike trains. Given a set of place fields $\{PF_1, PF_2, \dots, PF_N\}$, with specified shapes and locations, one can use a simple algorithm to construct the simplicial complex \mathcal{N} with vertex set $\{v_1, \dots, v_N\}$ (one vertex per cell/place field). Two vertices v_i and v_j are connected by an oriented 1D bond $\sigma_{ij} = [ij]$, if the corresponding regions PF_i and PF_j overlap. Three vertices support an oriented 2D facet $\sigma_{ijk} = [ijk]$, if there exists an overlap of three regions PF_i, PF_j and PF_k , and so on. In general, a simplex $\sigma_{i_1 i_2, \dots, i_k} = [i_1 i_2, \dots, i_k]$ is in \mathcal{N} if and only if

$$PF_{i_1} \cap PF_{i_2} \cap \dots \cap PF_{i_k} \neq \emptyset.$$

This is the so-called "Czech simplicial complex" or the "nerve" complex \mathcal{N} [10,44]. It can be shown that if the set of PFs, $\{PF_1, PF_2, \dots, PF_N\}$, covers the space X ,

$$X = \bigcup_{k=1}^N PF_k$$

sufficiently densely, then, under fairly general conditions, the nerve complex \mathcal{N} has the same homotopy type as the underlying space X , and so the topological invariants computed from \mathcal{N} will agree with those corresponding to X [10,15,44]. To be precise, "sufficiently dense" here means that each point of space X is contained in at least one place field, and each finite intersection of the fields is contractible.

In the context of studying a hippocampal map formation, in which the analysis is based on temporal characteristics of place cell activity, the simplicial complex can be constructed using the notion of *temporal* overlap between the spike trains rather than *spatial* overlap between place fields. The intuition is the following: If the rat happens to visit the location in space spanned by $PF_{i_1} \cap PF_{i_2} \cap \dots \cap PF_{i_k} \neq \emptyset$, then there is a non-zero probability that the cells $c_{i_1}, c_{i_2}, \dots, c_{i_k}$ will produce spikes at roughly the same time. Then the coactivity of the place cells can be interpreted as spatial connectivity: if at any a moment of time t during the observation period, two neurons c_i and c_j cofire, then there is a link between the corresponding vertices; if three neurons c_i, c_j and c_k cofire, then there is a 2D facet between the vertices and so on. We are therefore led to considering the following rule: fix $\varepsilon > 0$ and $m \in \mathbb{N}$. Then simplex

$$[i_1 i_2, \dots, i_k] \in \widehat{\mathcal{N}}_N \leftrightarrow \exists t \in [0, T] \text{ such that } \min_{j \in \{1, \dots, k\}} |s_{i_j} \cap [t - \varepsilon, t + \varepsilon]| \geq m.$$

This defines a "temporal simplicial complex" \mathcal{T} , which is a direct analogue of the "spatial" simplicial complex, which summarizes the information contained in the pattern of temporal overlaps between the spike trains and gives a complete topological description of the space X .

This construction achieves the goal of providing us with a topological method that can tell us

whether cells are indeed receiving all the information necessary for reconstructing the topology of the environment. The main question discussed in the paper is whether and to what extent different hippocampal states (as defined by variations in the mean firing rate \bar{f} , mean place field size \bar{s} , and the number of cells N) affect the network's ability to encode topological information.

In theory, there are two ways in which one can build simplicial complexes in order to describe the topological information contained in place cell firing activity: use place field geometry or place cell spike trains. How are the corresponding simplicial complexes \mathcal{N} and \mathcal{T} related?

It is often remarked that homology is functorial, by which it is meant that it faithfully represents topological information. To clarify this point, consider two simplicial complexes Σ and Σ' . Let $f: \Sigma \rightarrow \Sigma'$ be a continuous simplicial map: f takes each k -simplex of Σ to a k' -simplex of Σ' , where $k' \leq k$. Then, the map f induces a linear transformation $f_{\#}: C_k(\Sigma) \rightarrow C_k(\Sigma')$. 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

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

This is called the **induced homomorphism** of f on H_* . Functoriality means that (1) if $f: \Sigma \rightarrow Y$ is continuous then $f_*: H_k(\Sigma) \rightarrow H_k(Y)$ is a group homomorphism; and (2) the composition of two maps $g \circ f$ induces the composition of the linear transformation: $(g \circ f)_* = g_* \circ f_*$. This correspondence allows us to not only relate the spatial and the temporal complexes, but to consider the *dynamics* of simplicial complexes used in this paper to study the formation of different of hippocampal maps, using the idea of Homological Persistence.

Persistent homologies and barcodes. Consider the collection s_1, s_2, \dots, s_N of spike trains corresponding to the N cells. Consider the full simplicial complex K_N over N vertices. Given $\varepsilon > 0$ and $m \in \mathbb{N}$ we define a function $\mathbf{f}: K_N \rightarrow R^+$ as follows:

$$\mathbf{f}([i_1 i_2, \dots, i_k]) = \min \left\{ t \in [0, T], \min_{j \in \{1, \dots, k\}} |s_{i_j} \cap [t - \varepsilon, t + \varepsilon]| \geq m \right\}.$$

This function is a **filtration** on the simplicial complex K_N , and the pair (K_N, \mathbf{f}) is called a **filtered simplicial complex**. The concept of filtration is best understood by imagining that the simplicial complex is built across time. One starts with an empty simplicial complex and as time goes by, that is, as the rat explores the environment, the firing of cells "witnesses" [45] the formation of links between the vertices of the simplicial complex. For example, for a cell i , by definition $\mathbf{f}(i)$ equals the first time that a significant firing is observed in the spike train s_i . More precisely, $\mathbf{f}(i)$ equals the first time t that the spike count for s_i is above m in a window centered at t of size $\pm \varepsilon$.

Note that by definition, for any simplex σ containing i , $\mathbf{f}(\sigma) \geq \mathbf{f}(i)$. This implies, in particular, that a vertex is added to the simplicial complex earlier than any edge containing the vertex. More

generally, one also sees that $\mathbf{f}(\sigma) \geq \mathbf{f}(\tau)$ for any $\tau \subset \sigma$. Thus we have an increasing family of simplicial complexes, parameterized by the real line. Indeed, for each $t \geq 0$ let

$$K_N(t) = \{\sigma \in K_N, \mathbf{f}(\sigma) \leq t\}.$$

Then, if $t_1 \leq t_2 \leq \dots \leq t_n$ are all the different values taken by $\mathbf{f}(\sigma)$ as σ ranges in K_N , we have the increasing sequence of simplicial complexes

$$\emptyset \subset K_N(t_1) \subset K_N(t_2) \subset \dots \subset K_N(t_n) \subset K_N.$$

Notice that the last inclusion may be strict since it can very well happen that not all the simplices in K_N are witnessed [45]. The simplicial complex $K_N(t_n)$ above is the one that could be regarded as a proxy for \mathcal{N} : it contains all the connectivity information produced by all the co-firings that occurred before t .

Edelsbrunner and colleagues, however, made the following observation [46]. Since given $t \leq t'$ there is a natural inclusion of simplicial complexes $K_N(t) \subset K_N(t')$, because of the functoriality property described above, one obtains a linear transformation $H_k(K_N(t)) \rightarrow H_k(K_N(t'))$ for any k . What Edelsbrunner *et al.* observed was that in order to study the homology of a given space one should keep track of the entire system of vector spaces $H_k(K_N(t))$ 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 [47], 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 is a barcode (see above and Fig. 3). The barcode corresponds to the persistent cycles in the simplicial complex, and any two persistence vector spaces with the same barcodes are isomorphic. In the case of the temporal simplicial complex \mathcal{T} , these barcodes can be interpreted as the "time lines" traced by the topological loops, which characterize the stability of the topological structure defined by place cell co-activity patterns.

To analyze 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/>.

ACKNOWLEDGEMENTS

We thank the anonymous reviewers and V. Brandt for their critical reading of the manuscript. This work was supported by the Sloan and Swartz Foundations and NIH 5F32NS054425 (Y.D.), MH080283 (Y.D. and L.F.), DARPA grants HR0011-05-1-0007 and FA8650-06-1-7629, ONR grant N00014-09-1-0783 and AFOSR Grant FA9550-09-1-0643, (F.M. and G.C) Princeton Subaward 0001716-2 (G.C.)

REFERENCES

1. O'Keefe J, Dostrovsky J (1971) The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res* 34: 171-175.
2. O'Keefe J, Nadel L (1978) *The hippocampus as a cognitive map*. New York: Clarendon Press; Oxford University Press. xiv, 570 pp.
3. Zhang K, Ginzburg I, McNaughton BL, Sejnowski TJ (1998) Interpreting neuronal population activity by reconstruction: unified framework with application to hippocampal place cells. *J Neurophysiol* 79: 1017-1044.
4. McNaughton BL, Barnes CA, O'Keefe J (1983) The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. *Exp Brain Res* 52: 41-49.
5. Brown EN, Frank LM, Tang D, Quirk MC, Wilson MA (1998) A statistical paradigm for neural spike train decoding applied to position prediction from ensemble firing patterns of rat hippocampal place cells. *J Neurosci* 18: 7411-7425.
6. Best PJ, White AM, Minai A (2001) Spatial processing in the brain: the activity of hippocampal place cells. *Annu Rev Neurosci* 24: 459-486.
7. O'Keefe J, Burgess N (1996) Geometric determinants of the place fields of hippocampal neurons. *Nature* 381: 425-428.
8. Curto C, Itskov V (2008) Cell groups reveal structure of stimulus space. *PLoS Comput Biol* 4: e1000205.
9. Igel'nik B (2011) *Computational modeling and simulation of intellect : current state and future perspectives*. Hershey, PA: Information Science Reference. xxix, 655 p. p.
10. Hatcher A (2002) *Algebraic topology*. Cambridge ; New York: Cambridge University Press. xii, 544 p. p.
11. Muller RU, Kubie JL (1987) The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *J Neurosci* 7: 1951-1968.
12. Gothard KM, Skaggs WE, McNaughton BL (1996) Dynamics of mismatch correction in the hippocampal ensemble code for space: interaction between path integration and environmental cues. *J Neurosci* 16: 8027-8040.
13. Gothard KM, Skaggs WE, Moore KM, McNaughton BL (1996) Binding of hippocampal CA1 neural activity to multiple reference frames in a landmark-based navigation task. *J Neurosci* 16: 823-835.
14. Diba K, Buzsaki G (2008) Hippocampal network dynamics constrain the time lag between pyramidal cells across modified environments. *J Neurosci* 28: 13448-13456.

15. Dubrovin BA, Fomenko AT, Novikov SP (1992) *Modern geometry--methods and applications*. New York: Springer-Verlag. 432 pp.
16. Buzsaki G (2005) Theta rhythm of navigation: link between path integration and landmark navigation, episodic and semantic memory. *Hippocampus* 15: 827-840.
17. Buzsaki G (2002) Theta oscillations in the hippocampus. *Neuron* 33: 325-340.
18. Lisman JE (1997) Bursts as a unit of neural information: making unreliable synapses reliable. *Trends Neurosci* 20: 38-43.
19. Huxter JR, Senior TJ, Allen K, Csicsvari J (2008) Theta phase-specific codes for two-dimensional position, trajectory and heading in the hippocampus. *Nat Neurosci* 11: 587-594.
20. O'Keefe J, Recce ML (1993) Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* 3: 317-330.
21. Skaggs WE, McNaughton BL, Wilson MA, Barnes CA (1996) Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus* 6: 149-172.
22. Frank LM, Stanley GB, Brown EN (2004) Hippocampal plasticity across multiple days of exposure to novel environments. *J Neurosci* 24: 7681-7689.
23. Davidson TJ, Kloosterman F, Wilson MA (2009) Hippocampal replay of extended experience. *Neuron* 63: 497-507.
24. Zomorodian AJ (2005) *Topology for computing*. Cambridge, UK ; New York: Cambridge University Press. xiii, 243 p. p.
25. D'Hooge R, De Deyn PP (2001) Applications of the Morris water maze in the study of learning and memory. *Brain Res Brain Res Rev* 36: 60-90.
26. Morris RG, Garrud P, Rawlins JN, O'Keefe J (1982) Place navigation impaired in rats with hippocampal lesions. *Nature* 297: 681-683.
27. McNaughton BL, Barnes CA, Gerrard JL, Gothard K, Jung MW, et al. (1996) Deciphering the hippocampal polyglot: the hippocampus as a path integration system. *J Exp Biol* 199: 173-185.
28. Abusaad I, MacKay D, Zhao J, Stanford P, Collier DA, et al. (1999) Stereological estimation of the total number of neurons in the murine hippocampus using the optical disector. *J Comp Neurol* 408: 560-566.
29. Samsonovich A, McNaughton BL (1997) Path integration and cognitive mapping in a continuous attractor neural network model. *J Neurosci* 17: 5900-5920.
30. Eckert MJ, Bilkey DK, Abraham WC (2010) Altered plasticity in hippocampal CA1, but not dentate gyrus, following long-term environmental enrichment. *J Neurophysiol* 103: 3320-3329.
31. Eckert MJ, Abraham WC (2010) Physiological effects of enriched environment exposure and LTP

- induction in the hippocampus in vivo do not transfer faithfully to in vitro slices. *Learn Mem* 17: 480-484.
32. Fenton AA, Kao HY, Neymotin SA, Olypher A, Vayntrub Y, et al. (2008) Unmasking the CA1 ensemble place code by exposures to small and large environments: more place cells and multiple, irregularly arranged, and expanded place fields in the larger space. *J Neurosci* 28: 11250-11262.
 33. Karlsson MP, Frank LM (2008) Network dynamics underlying the formation of sparse, informative representations in the hippocampus. *J Neurosci* 28: 14271-14281.
 34. Shen J, Barnes CA, McNaughton BL, Skaggs WE, Weaver KL (1997) The effect of aging on experience-dependent plasticity of hippocampal place cells. *J Neurosci* 17: 6769-6782.
 35. Cacucci F, Yi M, Wills TJ, Chapman P, O'Keefe J (2008) Place cell firing correlates with memory deficits and amyloid plaque burden in Tg2576 Alzheimer mouse model. *Proc Natl Acad Sci U S A* 105: 7863-7868.
 36. Nithianantharajah J, Hannan AJ (2006) Enriched environments, experience-dependent plasticity and disorders of the nervous system. *Nat Rev Neurosci* 7: 697-709.
 37. White AM, Best PJ (2000) Effects of ethanol on hippocampal place-cell and interneuron activity. *Brain Res* 876: 154-165.
 38. Matthews DB, Simson PE, Best PJ (1996) Ethanol alters spatial processing of hippocampal place cells: a mechanism for impaired navigation when intoxicated. *Alcohol Clin Exp Res* 20: 404-407.
 39. Moser EI Open field trajectory data is freely available from. <http://www.ntnu.no/cbm/moser/gridcell>.
 40. Brown EN, Barbieri R, Ventura V, Kass RE, Frank LM (2002) The time-rescaling theorem and its application to neural spike train data analysis. *Neural Comput* 14: 325-346.
 41. Mizumori SJ (2008) *Hippocampal place fields : relevance to learning and memory*. Oxford ; New York: Oxford University Press. xx, 409 p. p.
 42. Aleksandrov PS (1965) *Elementary concepts of topology*. New York: F. Ungar Pub. Co. 63 p. p.
 43. Munkres JR (2000) *Topology*. Upper Saddle River, NJ: Prentice Hall. xvi, 537 p. p.
 44. Dubrovin BA, Fomenko AT, Novikov SP (1984) *Modern geometry--methods and applications*. New York: Springer-Verlag.
 45. de Silva V, and Carlsson, G. Topological estimation using witness complexes. In: Alexa M, and Rusinkiewicz, S. , editor; 2004; Zürich. pp. 157-166.
 46. Edelsbrunner H, Letscher, D., and Zomorodian, A. (2002) Topological Persistence and Simplification. *Discrete & Computational Geometry* 28: 511--533.
 47. Zomorodian A, and Carlsson, G. (2005) Computing persistent homology. *Discrete & Computational Geometry* 33: 249--274.

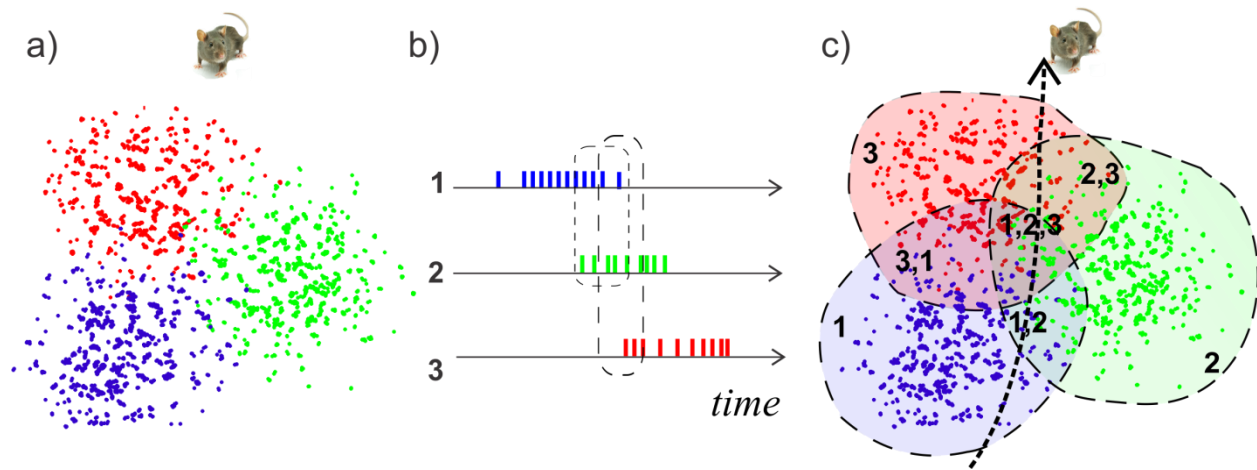


Figure 1. Place fields can be derived from place cell spike trains. (a) As a rat explores a given environment, various place cells will fire in spatially discrete locations. Here, for the sake of simplicity, we depict three place fields as they might arise from spike trains from three place cells, as in the next panel. (b) Schematic representation of spike trains fired from three different place cells as a rat explores an environment. Note that there is contemporaneous spiking activity, or co-firing. (c) The place fields derived from the three place cells in (b): the co-firing patterns indicates areas of overlap of the place fields. When the rat makes a straightforward trajectory through an explored environment, different place cells will be activated and their place fields can overlap.

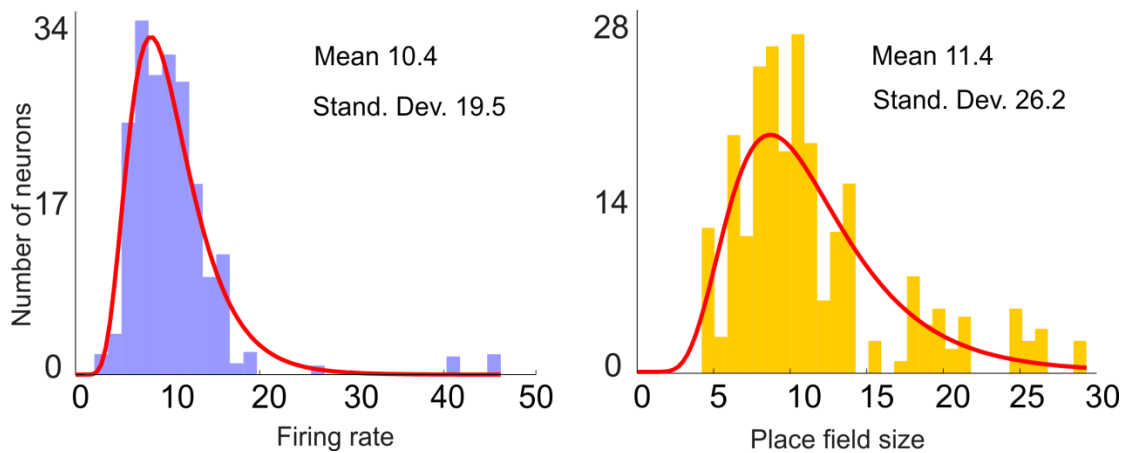


Figure 2. Distributions of firing rate and place field size collected by recording place cell firing as a rat explores a linear track. A typical place cell fires at a rate of ~10-20 Hz and place fields typically range from 10 to 30 cm across. These experimentally derived distributions serve as realistic constraints on our simulated data by providing proportionality coefficients a and b so that the shape of the distributions $P(f)$ and $P(s)$ mimics those derived from the experiments. From the data depicted here, $a = 1.2$ and $b = 1.7$.

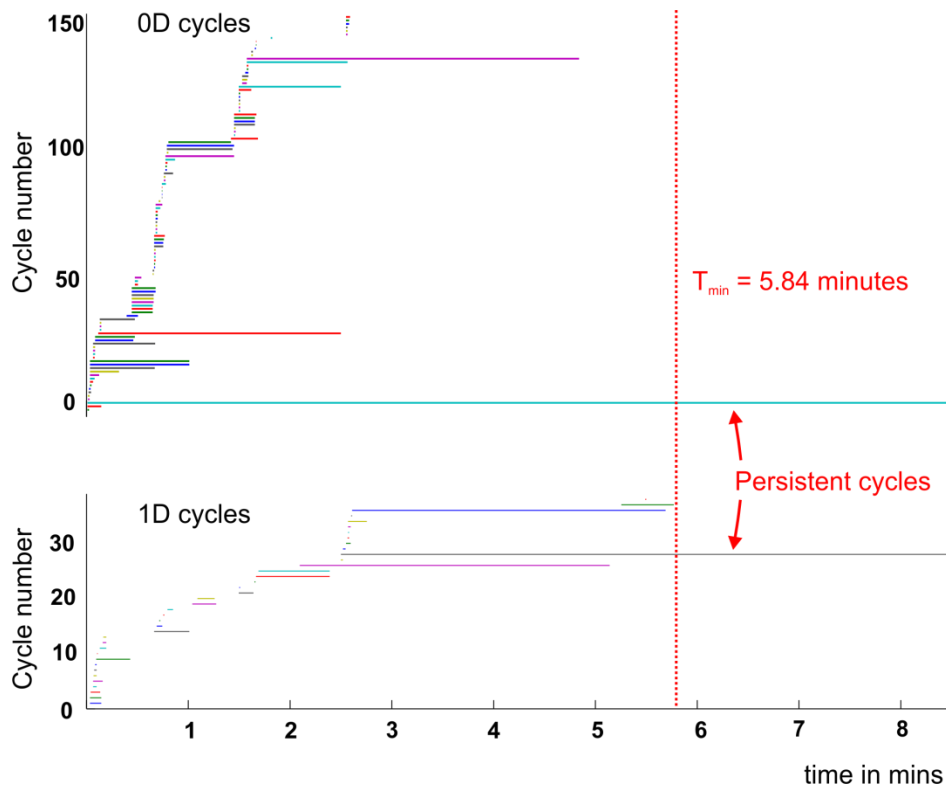




Figure 3. Persistent cycles form a topological barcode. Top and bottom graphs show which $0D$ and $1D$ cycles, respectively, persist in this cell ensemble. Each colored horizontal line represents one $0D$ cycle (top panel) or one $1D$ cycle (bottom panel). Initially, until cells begin co-firing, each $0D$ cycle corresponds to one cell. At later times, both $0D$ and $1D$ cycles are emergent phenomena, produced by co-firing of groups of cells. The dotted red vertical line at 5.84 minutes marks the moment when the correct number of loops appears in both $0D$ and in $1D$, which is the minimal map formation time \bar{T}_{min} . The series of short horizontal bars in both panels (some quite miniscule) and the longer lines that disappear before \bar{T}_{min} represent topological noise, i.e., cycles that fail to persist. The one persistent $1D$ cycle indicates that the environment in question has one physical (topological) loop, and the single $0D$ cycle indicates that the space is connected, of one piece. Together, this pattern of stable bars forms a barcode that can be ‘scanned’ to discern the topological structure of the environment (see Methods).

i. experimental scenarios (2x2m)


A. One-hole






B. Two holes






C. Quasi Linear



ii. mean map formation times



iii. relative variance of map formation times



iv. stable learning submanifold






Figure 4. Variations in topology place different demands on hippocampal state. The *top row* depicts three experimental configurations, each two meters square, for our computational simulations; note that the second and third scenarios (B and C) are topologically equivalent but geometrically different, and that scenario C will force our simulated rat to adopt a quasi-linear trajectory. The dense network of gray lines represents the simulated trajectories. *Second row*: Point cloud approximations that reveal mean map formation times for each space configuration. Each dot represents a hippocampal state as defined by the three parameters (\bar{f} , \bar{s} , and N); the size of the dot reflects the proportion of trials in which a given set of parameters produced the correct outcome; the color of the dot is the mean time over ten simulations. If, for example, one set of parameters produced the correct topological information in 6 out of 10 trials, the dot will be 60% of the size of the largest dot, and the color will reflect the mean map formation time for the correct trials. (Blue represents success within the first 25% of the total time; green within the first 50%, yellow-orange within the first 75%, and red means success took nearly the whole time period. The maximal observed time was 4.3 minutes for configuration A, 11.7 minutes for B, and 9.3 minutes for C.) Note how the third scenario (C) contains a preponderance of blue dots, indicating that the majority of hippocampal states easily mapped this environment. This is because the two holes are so large that a rat is virtually forced into a straight-line trajectory. *Third row*: Each dot represents the relative standard deviation of map formation times T_{min} for successful trials where ΔT_{min} is very small (<0.15). *Fourth row*: Combining the mean map formation times (second row) with the robustness requirement $\Delta T_{min}/\bar{T}_{min} < 0.15$ (third row) reveals a domain of stable, robust map formation times that we call the core of the region L in the text.

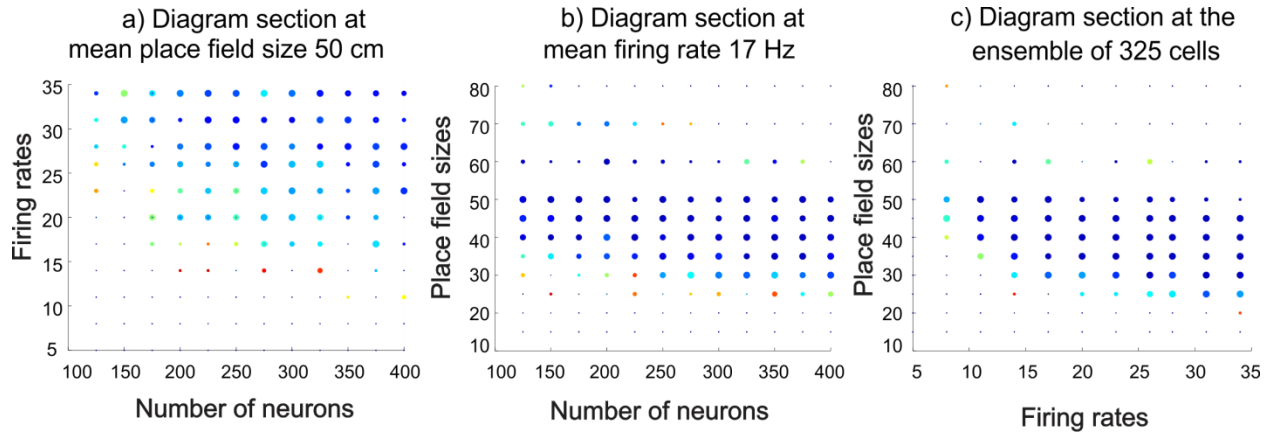


Figure 5. 2D sections highlight dependence of map formation times on hippocampal state. These 2D sections are based on the point cloud data in Figure 4C, second row (far right). Dot sizes and colors represent the same characteristics as described in Figure 4 (i.e., the larger and bluer the dot, the more successful and more rapid the map formation). Graph A fixes the mean place field size at 50 cm, and shows that robust map formation in this case requires a larger number of cells firing at a higher rate. Graph B shows that, at a mean firing rate of 17 Hz, any ensemble size between 100 and 400 neurons can fairly rapidly form a correct topological map as long as the place fields are between 50 and 80 cm. Graph C shows that an ensemble of 325 cells can have mean firing rates from 10 to 35 Hz and form maps quickly and accurately with place field sizes from 40-80 cm. In short, smaller place cell ensembles, with low mean firing rates ($\bar{f} < 10$ Hz) and too small ($\bar{s} < 20$ cm) or too large ($\bar{s} > 100$) mean place field sizes, fail to produce the correct topological signature. In contrast, sufficiently large place cell ensembles with higher firing rate neurons and well-tuned place fields reliably capture the topological structure of the environment in a time frame comparable to the experimentally observed map formation period.

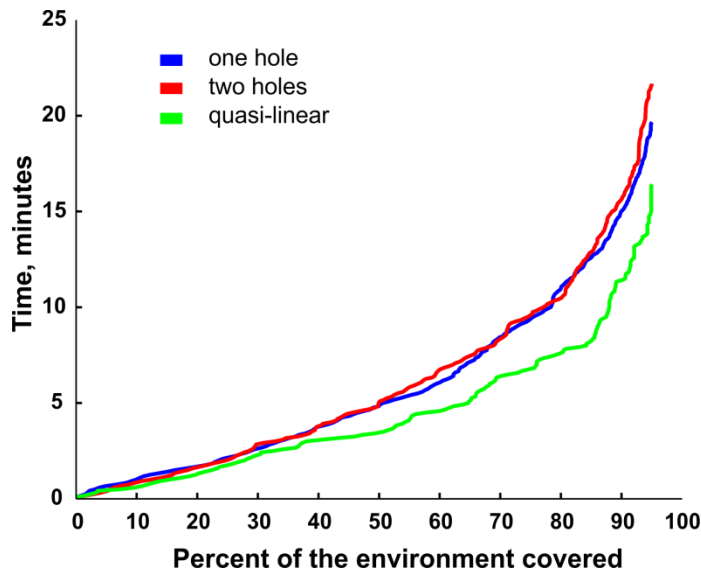


Figure 6. Ergodicity times for the three environments shown in Figure 4. For each environment, the graph shows how much time is required to cover a certain percentage of the 3×3 cm spatial bins. This ergodic time scale shows that it takes approximately ten minutes for a rat to cover 80% of the environment; by comparison, the topological map formation time for stable regimes is much lower.

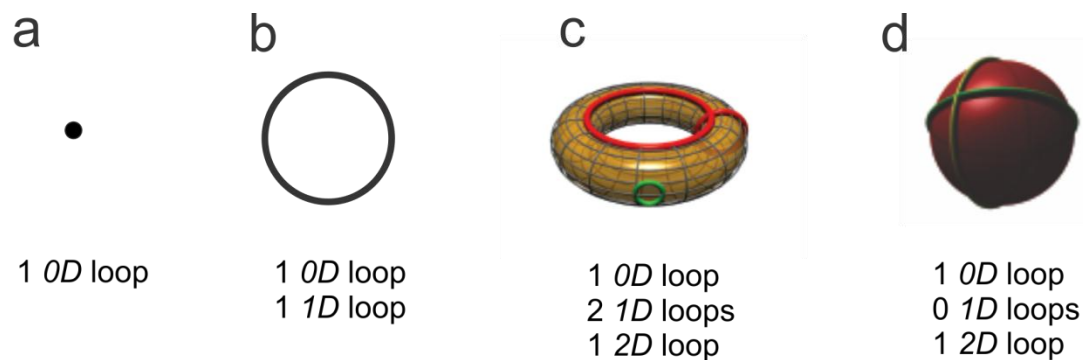

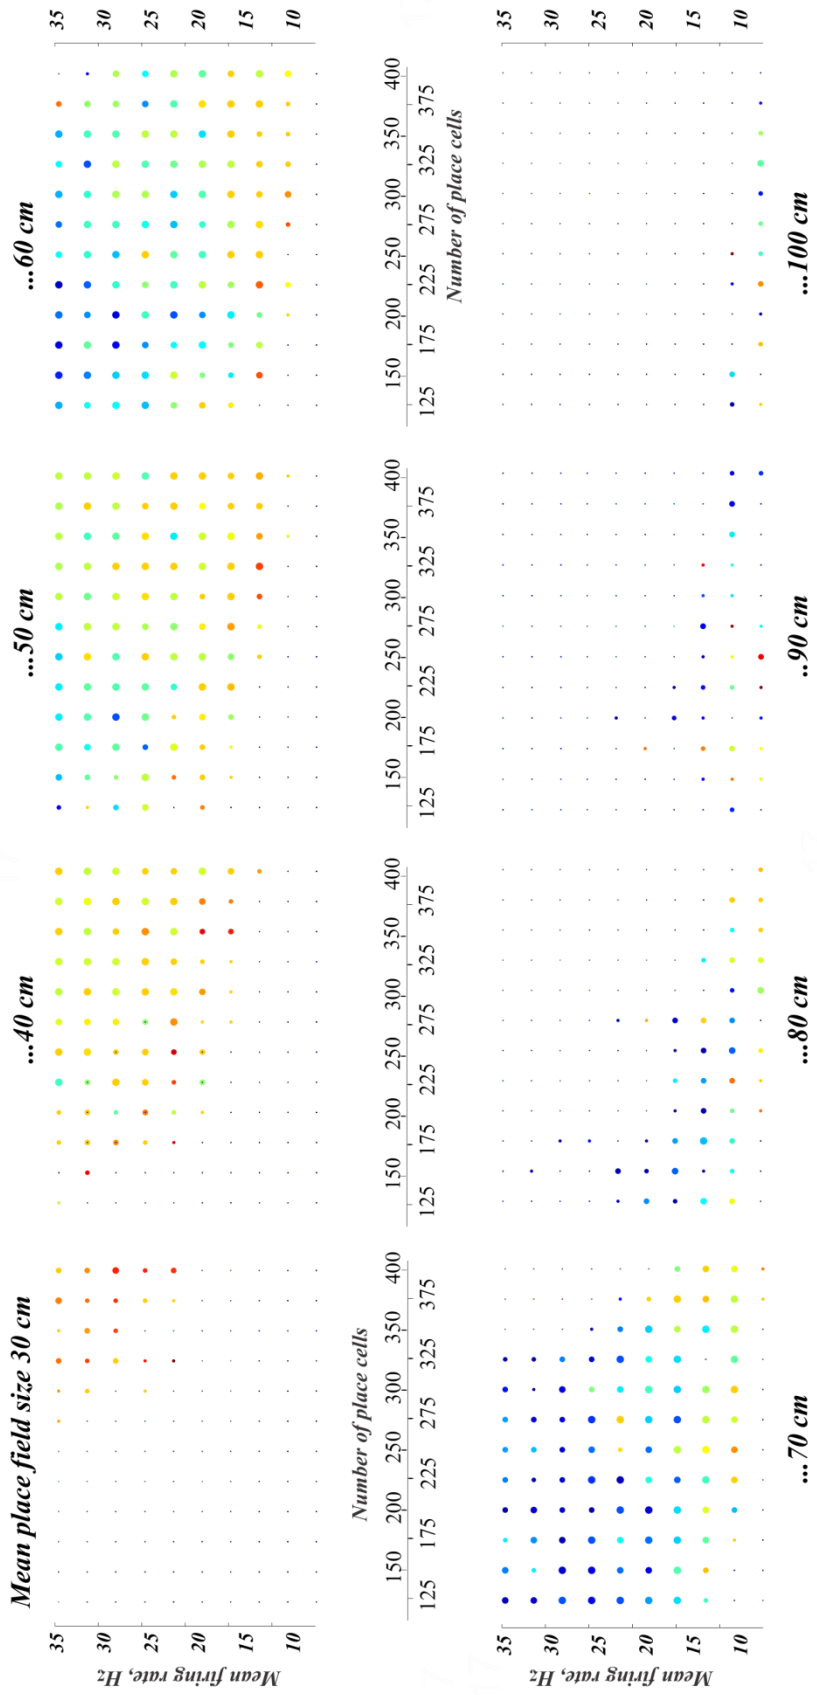


Figure 7. Examples of low-dimensional manifolds and their Betti numbers with some of the corresponding loops. (a) A point is a $0D$ loop; no higher dimensional loops are present. Thus, each manifold containing at least one point has a $0D$ loop, so every list of Betti numbers starts with a "1". (b) A circle is a $1D$ loop, with no other loops in higher dimensions. (c) A $2D$ torus with two examples of non-contractible (red) $1D$ loops, and an example of a $1D$ loop contractible into a point (green). The $2D$ surface of the torus is the $2D$ loop listed. (d) A $2D$ sphere, with two exemplary contractible $1D$ loops. The $2D$ surface of the sphere "loops" onto itself.



Supplementary Figure 1. 2D slices of point cloud data in Figure 4 (environment A, fourth row), with steady variation in mean firing rate. Colors and sizes of dots code for the same meanings as described in the legend of Figure 4: the large blue dots represent the most successful hippocampal states for map formation, with the most rapid map formation times. Here the graphs show a gradual increase in mean firing rate (from 5-31 Hz) and how this affects the overall shape of the learning region. At low firing rates (upper left panels) there is no successful map formation; at 8 Hz, we begin to see some map formation occurring at the largest place field sizes (80-90 cm), especially as the number of neurons increases to 300-350. By 17 and 20 Hz (lower left panels), there is fairly good and rapid map formation with place field sizes around 60 cm. By the time the firing becomes very rapid (31 Hz), smaller place field sizes of 20 cm are able to produce topologically accurate maps, sometimes, but map formation time is long (red dots) so the process is not very efficient.



Supplementary Figure 2. 2D slices of point cloud data in Figure 4, (environment A, fourth row) with steady variation in mean place field size. From these data it appears that the hippocampal state is less sensitive to the chosen range of mean place field sizes (especially between 50 and 80 cm) than it is to firing rate. The graphs show that at this mid-range of place field size, map formation is rapid and accurate (lots of blue dots) for a fairly wide range of firing rates and number of cells in the ensemble.