MIT Open Access Articles

Neural Representation of Spatial Topology in the Rodent Hippocampus

Massachusetts Institute of Technology

Neural Representation of Spatial Topology in the Rodent Hippocampus

Zhe Chen
zhechen@mit.edu
Department of Brain and Cognitive Sciences and Picower Institute
for Learning and Memory, MIT, Cambridge, MA 02139, U.S.A.

Stephen N. Gomperts
sgomperts@partners.org
Picower Institute for Learning and Memory, MIT, Cambridge, MA 02139,
and Department of Neurology, Massachusetts General Hospital,
Harvard Medical School, Boston, MA 02116, U.S.A.

Jun Yamamoto
yamajun@mit.edu
Matthew A. Wilson
mwilson@mit.edu
Department of Brain and Cognitive Sciences and Picower Institute
for Learning and Memory, MIT, Cambridge, MA 02139, U.S.A.

Pyramidal cells in the rodent hippocampus often exhibit clear spatial tuning in navigation. Although it has been long suggested that pyramidal cell activity may underlie a topological code rather than a topographic code, it remains unclear whether an abstract spatial topology can be encoded in the ensemble spiking activity of hippocampal place cells. Using a statistical approach developed previously, we investigate this question and related issues in greater detail. We recorded ensembles of hippocampal neurons as rodents freely foraged in one- and two-dimensional spatial environments and used a "decode-to-uncover" strategy to examine the temporally structured patterns embedded in the ensemble spiking activity in the absence of observed spatial correlates during periods of rodent navigation or awake immobility. Specifically, the spatial environment was represented by a finite discrete state space. Trajectories across spatial locations ("states") were associated with consistent hippocampal ensemble spiking patterns, which were characterized by a state transition matrix. From this state transition matrix, we inferred a topology graph that defined the connectivity in the state space. In both one- and two-dimensional environments, the extracted behavior patterns from the rodent hippocampal population codes were compared against randomly shuffled spike data. In contrast to a topographic code, our results support

the efficiency of topological coding in the presence of sparse sample size and fuzzy space mapping. This computational approach allows us to quantify the variability of ensemble spiking activity, examine hippocampal population codes during off-line states, and quantify the topological complexity of the environment.

1 Introduction

Population codes derived from simultaneous recordings of ensembles of neurons have been studied in the representation of sensory or motor stimuli and in their relationship to behavior (Georgopoulos, Schwartz, & Kettner, 1986; Schwartz, 1994; Nirenberg & Latham, 1998; Sanger, 2003; Broome, Jayaraman, & Laurent, 2006). Uncovering the internal representation of such codes remains a fundamental task in systems neuroscience (Quian Quiroga & Panzeri, 2009). The rodent hippocampus plays a key role in episodic memory, spatial navigation, and memory consolidation (O'Keefe & Dostrovsky, 1971; O'Keefe & Nadel, 1978; Wilson & McNaughton, 1993, 1994; Buzsáki, 2006). Pyramidal cells in the CA1 area of the rodent hippocampus have localized receptive fields (RFs) that are tuned to the (measured) animal's spatial location during navigation in one-dimensional (1D) or two-dimensional (2D) environments. These cells are referred to as place cells, and their RFs are referred to as place fields (O'Keefe & Dostrovsky, 1971). However, the concept of place fields was invented by human observers for the purpose of understanding the tuning of place cells. It remains unclear how neurons downstream of the hippocampus can infer representations of space from hippocampal activity without place field information a priori. Two types of maps, topographical and topological, may be used for spatial representation. A topographical map contains metric information (such as distance and orientation) between two locations in the map, whereas a topological map contains only relative ordering or connectivity information between locations and is invariant to orientation (see the appendix for mathematical definitions). Since topological features are preserved despite deformation, twisting, and stretching of the space, a topological map would be an appealing candidate for abstract spatial representation. Uncovering the internal spatial representation of rodent hippocampal ensemble activity will not only help reveal the coding mechanism of the hippocampus, but will also provide a way to interpret ensemble spike data during off-line states.

It has been suggested that an inner physiological space could emerge from the population spiking activity of hippocampal neurons (O'Keefe & Nadel, 1978). A few groups have recently examined the issue of topological coding of space in the rat hippocampus (Curto & Itskov, 2008; Dabaghian, Cohn, & Frank, 2011; Dabaghian, Memoli, Frank, & Carlsson, 2012), but all reported studies were limited to computer simulations or theoretic

conceptualization. The essential question of interest is how the output of rodent hippocampal place cells (without the explicit construction of place fields) might be used by downstream structures in order to reconstruct both the animal's position and the topology of the environment. From a data analysis point of view, the question is this: How do we transform the temporal patterns of spiking activity in the form of multiple time series into a spatial pattern of place fields? To investigate these questions in greater detail, here we use a computational and data-driven approach (Chen, Kloosterman, Brown, & Wilson, 2012) to investigate the neural representation of spatial topology embedded in rodent hippocampal population codes. The new contributions beyond our previous work are threefold. First, to our knowledge, this is the first systematic investigation on this topic. We explore this issue across different species (rat versus mouse), environments (1D versus 2D), and behaviors (locomotion versus quiet wakefulness). Second, we extract quantitative statistics (such as the topological complexity) from topological codes and topological graphs. Third, we systematically investigate the representation efficiency between topological and topographic codes in a 2D environment. We hope that the investigations and findings presented here, while still empirical, will add insight into well-studied rodent hippocampal population codes.

2 Inferring Rodent Hippocampal Population Codes

2.1 Probabilistic Modeling. We used a hidden Markov model (HMM) to model the population spiking activity from simultaneously recorded C rodent hippocampal neurons. We assumed that the animal's spatial location during locomotion, being modeled as a latent state process, followed a first-order discrete-state Markov chain $\{S(t)\} \in \{1, \ldots, m\}$ (where m denotes the size of the discrete state space). We also assumed that conditional on the hidden state $S(t)$ at time t, the spike counts of individual hippocampal neurons followed a Poisson probability with their own tuning curves. Stated mathematically, we used the following probabilistic mapping (Chen et al., 2012),

$$S(t-1) \to S(t) \sim P_{S_{t-1}S_t}, \tag{2.1}$$

$$y_c(t)|S(t) = j \sim Poisson(y_c(t); \lambda_c(j)), \tag{2.1}$$

where $\boldsymbol{P} = \{P_{ij}\}$ denotes an m-by-m state transition probability matrix, with the element P_{ij} representing the transition probability from state i to state j; $y_c(t)$ denotes the number of spike counts within the tth temporal bin from the cth neuron, which has the tuning curve λ_c with respect to the state space. Given the multiple time series of spike counts $\boldsymbol{y} = \{\boldsymbol{y}(1), \ldots, \boldsymbol{y}(T)\}$ (where $\boldsymbol{y}(t) = [y_1(t), \ldots, y_C(t)]$ is a C-dimensional population vector), our goal is to infer the mostly likely hidden state sequence $\boldsymbol{S} = \{S(1), \ldots, S(T)\}$ and

the unknown parameters $\theta = \{\pi, P, \Lambda\}$, where π denotes the initial state probability vector and $\Lambda = \{\lambda_c(j)\}$ denotes a C-by-m tuning curve matrix that can be interpreted as the virtual place fields or state fields of all hippocampal neurons. Given the animal's locomotion behavior as well as the spatial topology of the environment, the ground truth transition probability matrix P captures important information related to the spatial environment. The computational task is to infer the transition probability matrix P from the ensemble spike data alone (without assuming any knowledge of the animal's behavior).

In this probabilistic modeling framework, we represented a continuous topographic space $\mathcal{S}$ by a finite discrete alphabet $\mathcal{A}$ using a code book: $\mathcal{S} = f(\mathcal{A})$. The consistency of this representation requires a one-to-one mapping between $\mathcal{S}$ and $\mathcal{A}$. First, any element in $\mathcal{S}$ is not simultaneously represented by A_i and A_j ($i \neq j, A_i \in \mathcal{A}, A_j \in \mathcal{A}$). Second, the same A_i does not represent two or more distinct regions in $\mathcal{S}$ (except for neighboring regions that can be merged). Of note, A_i and A_j may encode two regions, each with different spatial coverage.

2.2 Bayesian Inference. We applied a variational Bayes (VB) algorithm to estimate the unknown hidden state S and unknown parameters $\theta = \{\pi, P, \Lambda\}$ (Chen et al., 2012). For Bayesian inference, we imposed informative conjugate priors for θ, specified by $p(\pi, P, \Lambda)$. Specifically, we used a Dirichlet prior for the multinomial likelihood (vector π and row vectors of P) and a gamma prior for $\{\lambda_c(j)\}$. The goal of VB inference was to maximize the lower bound of the marginal log-likelihood $\log p(y)$, also known as the free energy:

$$
\begin{aligned}
\log &\, p(y) \\
&= \log \int d\pi \int dP \int d\Lambda \sum_S q(\pi, P, \Lambda, S) \frac{p(\pi, P, \Lambda)p(y, S|\pi, P, \Lambda)}{q(\pi, P, \Lambda, S)} \\
&\geq \int d\pi \int dP \int d\Lambda \sum_S q(\pi, P, \Lambda, S) \log \frac{p(\pi, P, \Lambda)p(y, S|\pi, P, \Lambda)}{q(\pi, P, \Lambda, S)} \\
&= \langle \log p(y, S, \pi, P, \Lambda) \rangle_q + \mathcal{H}_q(\pi, P, \Lambda, S) \equiv \mathcal{F}(q), \quad (2.3)
\end{aligned}
$$

where $p(y, S|\pi, P, \Lambda)$ defines the complete data likelihood and $q(\pi, P, \Lambda, S) \approx q(\pi)q(P)q(\Lambda)q(S)$ represents the factorial variational posterior that approximates the joint posterior of the hidden state and parameter $p(\pi, P, \Lambda, S|y)$. The term $\mathcal{H}_q$ represents the Shannon entropy of the distribution q. The best approximation in q to the joint posterior yields the tightest lower bound on $\log p(y)$.

We applied an iterative expectation-maximization (EM) type algorithm to optimize the free energy $\mathcal{F}$ until it reached a local maximum. In the

VB-E step, we estimated the sufficient statistics using a standard forward-backward algorithm. In the VB-M step, we estimated the variational posteriors $q_{\theta}(\theta)$ and the posterior mean $\tilde{\theta}$ statistics. (Details of the method are in Chen et al., 2012.) During the testing mode, given the posterior mean statistics of the estimated parameters and the ensemble spiking activity from the same hippocampal neurons, we ran a modified version of the EM algorithm to obtain the maximum a posteriori (MAP) estimate of the state sequence as well as the free energy score of the tested ensemble spike data.

2.3 Model Selection and Assessment. We used the Bayesian deviance information criterion (DIC) as a guiding principle for selecting the state dimensionality m. The DIC is defined as the sum of the expected deviance and the model complexity measure p_D (McGrory & Titterington, 2009):

$$DIC = \mathbb{E}_{p(\theta|y)}[-2\log p(y|\theta)] + p_D$$

$$\approx -2\log p(y|\tilde{\theta}) - 2\int q_{\theta}(\theta)\log\frac{q_{\theta}(\theta)}{p(\theta)}d\theta + 2\log\frac{q_{\theta}(\tilde{\theta})}{p(\tilde{\theta})}, \qquad (2.4)$$

where $\tilde{\theta}$ denotes the posterior mean computed with respect to $q_{\theta}(\theta)$ and $p(y|\tilde{\theta})$ can be computed from the forward-backward algorithm (Chen et al., 2012). Unless stated otherwise, the model with the smallest DIC value would be selected.

In our analysis of spike data during periods of locomotion, we used a state-space map to visualize the accuracy of the estimation. The map was used to check the consistency of the one-to-one mapping between the discrete state ID and the animal's spatial position. Visualized as a matrix, the element of the state-space map represents the counting number of mapping between the state and the specific spatial position. Ideally, the state-space map would consist of segmented nonoverlapping narrow "stripes," where the width of the stripes would reflect the uncertainty of the reconstructed position.

2.4 Visualization of Spatial Topology Graph. On the completion of Bayesian inference, we obtained the estimated state trajectory S as well as the posterior mean of the parameters $\tilde{\theta}$. To visualize the spatial topology by an undirected graph, we represented the states with m distinct nodes. The presence of the edges between the nodes indicated that two nodes were connected in space, and the strength of the edge between two nodes (shown in terms of the darkness of the color) was proportional to the transition probability value between two states. The estimated transition probability matrix was fed to a custom graph-drawing force-based algorithm to produce a 2D topology graph (Battista, Eades, Tamassia, & Tollis, 1998; Chen et al., 2012). To reduce the impact of noise (i.e., estimation bias), we often thresholded

the estimated transition probability matrix with a typical threshold value varying between 0.1 and 0.2. From the same transition matrix, the exact shape of the derived topology graphs might appear differently due to randomly initialized conditions of the force-based algorithm, but the underlying topology spaces were homeomorphic.

2.5 Hypothesis Testing and Statistical Significance Test. The estimated state transition matrix was inferred from the observed population spike trains based on our statistical model (see equations 2.1 and 2.2). Due to limited sampling, the estimate would be subject to statistical bias and variance. To test the statistical significance of the extracted structure of the population spike data, we created randomly shuffled data and compared their statistics. The null hypothesis was that the population codes were observed by chance and therefore could not be captured by our statistical model due to the lack of structure. To test against the null hypothesis, we randomly shuffled the population spike count matrix in both row (unit identity shuffle) and column (time shuffle). The time shuffle was done by randomly jittering the temporal bins uniformly between 1 and 5 seconds. In each Monte Carlo shuffle trial, the row and column shuffles were conducted 1000 times. We performed 100 Monte Carlo shuffle trials. We then compared the distributions of the converged free energy value derived from the raw and shuffled data based on the same Poisson population firing model. The Monte Carlo P-value was reported using a two-sample Kolmogorov-Smirnov (KS) test.

2.6 Neural Encoding and Decoding with Topographic Maps. For experimental comparison (see section 4), we also implemented a topographic coding method. In the encoding phase, we estimated the place fields of hippocampal pyramidal neurons using a simple nonparametric model based on the Poisson assumption. Specifically, we discretized the spatial environment and estimated the individual neuronal firing rate based on the spiking activity during the running epochs. In the decoding phase, we used a maximum likelihood estimator (or empirical Bayesian estimator with uniform priors) (Zhang, Ginzburg, McNaughton, & Sejnowski, 1998). Note that our strategy of using such a simple encoding-decoding method was for illustrating the topographic code instead of searching for optimal decoding performance, although improved decoding accuracy would be expected by employing more sophisticated neural encoding-decoding methods (Brown, Frank, Tang, Quirk, & Wilson, 1998; Barbieri et al., 2004).

3 Results on Neural Topological Representations

3.1 Experimental Data. The experimental data were collected under several different protocols, in which animals freely foraged in different
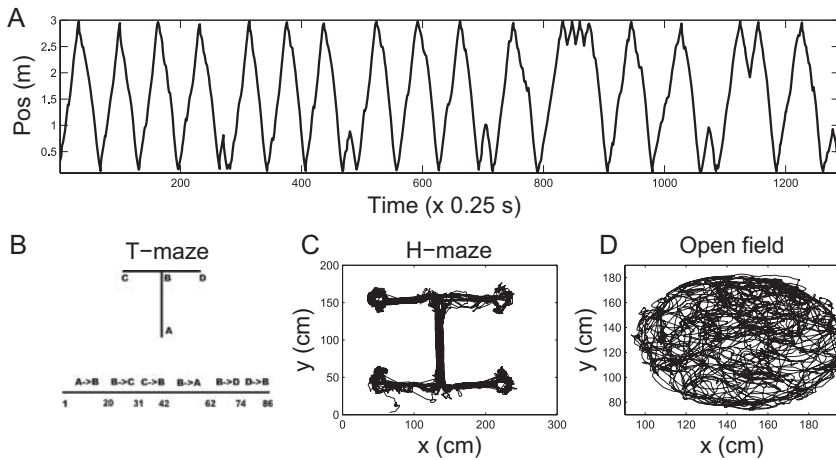
Figure 1: Examples of spatial environments used in the analysis. (A) linear track. (B) T-maze. (C) H-maze. (D) Open field. In panel A, the animal's linearized run trajectory is shown, and in panels C and D, the animal's actual spatial trajectories are shown. Note that among three one-dimensional environments (A–C), the linear track has two end points and no junction, the T-maze has three end points and one junction, and the H-maze has four end points and two junctions.

spatial environments (see Figure 1) (Barbieri et al., 2004; Yamamoto & Wilson, 2008; Davidson, Kloosterman, & Wilson, 2009). All experiments were conducted under the supervision of the MIT Committee on Animal Care and followed the guidelines of the U.S. National Institutes of Health. Details of experimental protocols and recordings are shown in the appendix. Data from 10 animals (8 rats and 2 mice) were collected and analyzed. Due to limitations of space, only eight data sets (see Table 1) were presented in our experimental data analysis. For each data set, we binned the spiking activity of rodent hippocampal putative pyramidal neurons with a 250 ms bin size to obtain the raw spike count statistics for the locomotion period.

We analyzed the hippocampal population spike data from both rats and mice during periods of locomotion. In the 1D environment, hippocampal place cell activity depends on both location and the direction of navigation (Markus et al., 1995); the same spatial location was represented by two different states depending on navigation direction (inbound versus outbound). In contrast, in the 2D open field environment, hippocampal pyramidal cells exhibit place-specific firing that is statistically independent of the direction of traversal through the place field; therefore, navigation direction was not incorporated into the open field analysis.

Table 1: Summary of Experimental Data.

Data set	Number of (Used) Units	Total (RUN) epoch	Environment
R1 (rat A)	74 (30)	30.5 (4.7) min	Linear track (3.0 m)
R2 (mouse)	107 (51)	39.6 (4.0) min	Linear track (1.5 m)
R3 (rat B)	23 (20)	34.0 (0.9) min	Linear track (2.1 m)
R4 (rat B)	22 (18)	31.5 (2.5) min	Linear track (2.1 m)
R5 (rat C)	39 (25)	16.0 (4.0) min	T-maze (2.05 m)
R6 (rat D)	27 (27)	63.7 (13.2) min	H-maze (5.3 m)
R7 (rat E)	49 (47)	24.3 (4.9) min	Open field (radius ∼0.5 m)
R8 (rat F)	37 (36)	22.9 (7.5) min	Open field (radius ∼0.5 m)

3.2 Spatial Topology of One-dimensional Tracks During Navigation.
In the 1D linear track data sets from rats and mice, we observed qualitatively
similar ensemble spiking patterns on a single lap basis (see Figure 2A)
despite the behavioral variability (e.g., run velocity) and differences in the
size of the environment (track length). Provided the data are binned and
viewed in space (rather than in time), each panel in Figure 2A can be treated
as single-lap place fields for all hippocampal units. The similarities among
those ensemble spike patterns can be characterized by a cross-correlation
matrix (below Figure 2A).

3.2.1 Selection of the Number of Hidden States. Our probabilistic modeling
required us to select a priori the number of hidden states m for the HMM. It is
important to select an appropriate model size for fitting a small data sample.
Selecting a model size that is too small or too large results in underfitting
or overfitting of the spike data. In our experimental observations, when the
selected m was too small, we often found nonunique mappings between
the state and the animal's position in the state-space map (e.g., one state
representing more than two positions in the space; see Figure 3A), indicating
that the state dimensionality was insufficient. When the selected m was too
large, we often found redundant states, which meant that some states had
not been used for mapping the space (result not shown; but simulation
results were shown in Chen et al., 2012).

To carry out model selection, we used the Bayesian DIC (see equation 2.4).
Under each selected model size, we ran a number of Monte Carlo exper-
iments and obtained the mean and SD of the DIC values (see Figure 3E).
In one 1D linear track example (R1 data set), we found that among the
evaluated models, the optimal model size was around 60. In this case, we
obtained the state-space map, the state transition probability matrix, and
the spatial topology map (see Figures 3A–3D). We also recovered the tun-
ing curves for all units in the state space (result not shown; see an example
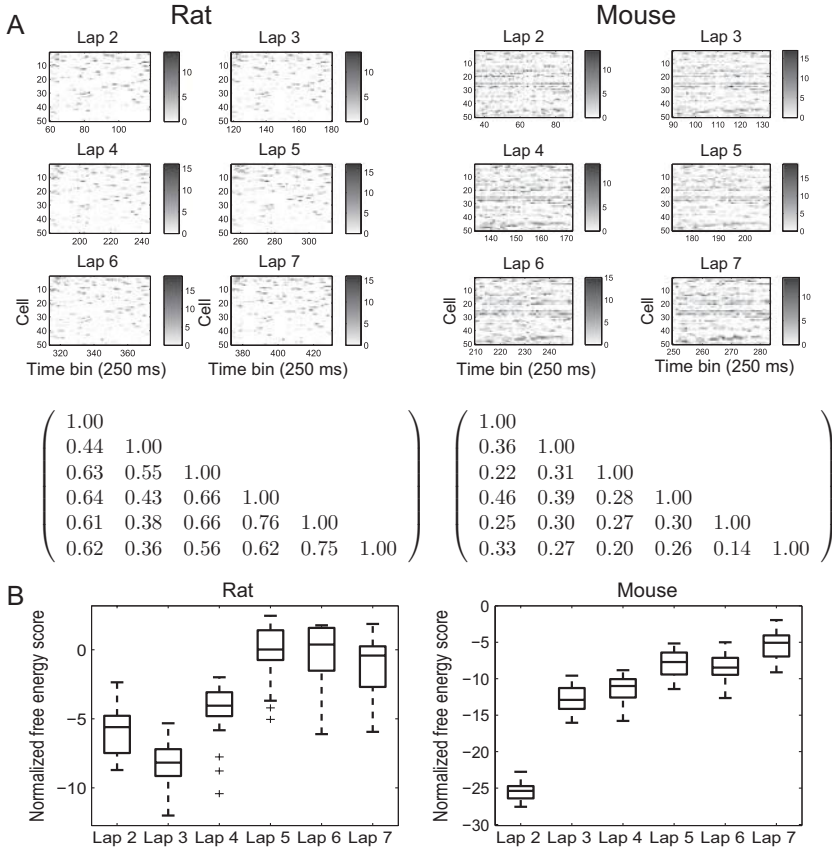
Figure 2: (A) Variability of rodent hippocampal population neuronal responses (in the form of spike count matrix) in a lap-by-lap comparison (Left: rat, linear track R1 data set; right: mouse, linear track R2 data set). All six laps consisted of one cycle of an outbound-inbound trajectory that contained varying number of time bins. Each panel can be viewed as single-lap place fields for all hippocampal units. The grayscale bar represents the spike count. The similarities between six-lap ensemble spike patterns are quantified by a cross-correlation matrix shown below. The average running velocities for the rat during the six laps were 42.8, 40.7, 38.8, 42.8, 42.0, and 44.3 cm/s, respectively; the average running velocities for the mouse during the six laps were 22.2, 27.9, 31.6, 33.3, 30.8, and 36.4 cm/s, respectively. (B) Quantifying the variability of the ensemble spike data show above. In each box plot, the normalized free energy score was computed from each single lap based on the inferred model parameters. The statistics were obtained from 50 independent Monte Carlo trials.

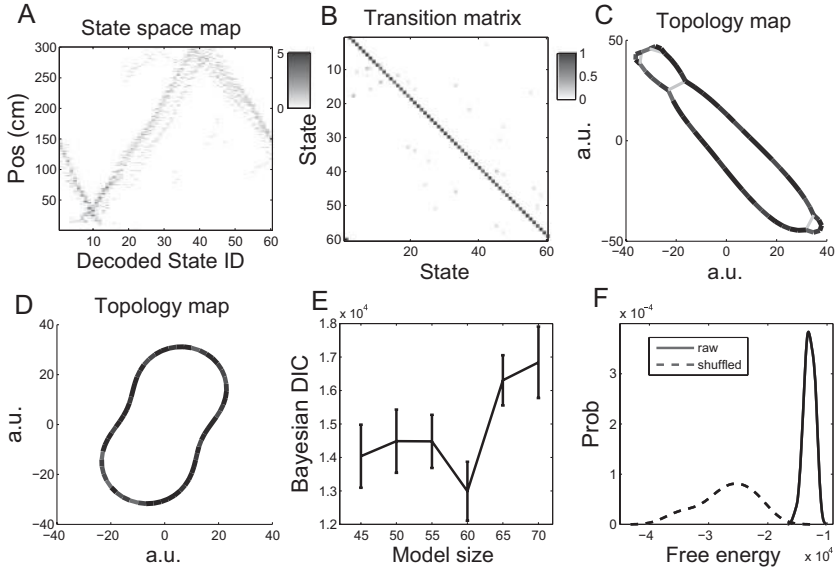Figure 3: Linear track (R1 data set) result. (A) State-space map ($m = 60$ from model selection). Grayscale bar represents the repeated frequency of mapping between the state and position. (B) One example of the estimated state transition probability matrix. (C) Inferred spatial topology graph from the transition matrix. Dark-colored edges imply that two nodes are strongly connected or two states are linked by high transition probability values. Note that the axes of the topology map are in arbitrary unit (au). (D) Inferred topology graph derived from a thresholded transition matrix (threshold 0.2). Despite the deformation of the orientation and width of the graph, the ring topology structure remains the same due to topological invariance. (E) Model selection based on the Bayesian deviance information criterion (DIC). Error bar statistics were computed from the results of 50 independent Monte Carlo trials. (F) Distributions of the derived free energy value from the raw (solid curve) and shuffled (dashed curve) spike count data for the same R1 data set ($m = 60$). In each condition, probability distributions were estimated from the derived free energy values based on 100 independent Monte Carlo trials. The median free energy obtained from the raw spike data was significantly greater than that obtained from the shuffled data (Mann-Whitney test, $P < 0.001$), and these two empirical distributions were also significantly different (two-sample KS test, $P < 0.001$).

shown in Chen et al., 2012). In this example, the spatial topology had a ring structure, reflecting the linear track topology. There were also a few light gray color-coded shortcuts inside the ring structure (see Figure 3C), which reflected the animal turning before the ends of the track (see Figure 1A; three turning points were around 75 cm, 200 cm, and 250 cm). However,

the edges of these shortcuts had relatively small strengths because of their infrequency. Provided that we thresholded the state transition probability matrix by 0.2, these two shortcuts disappeared in the final topology graph (see Figure 3D).

The overfitting problem was not an issue in our data analysis. To test the generalization, we also tried using the first half of the data to estimate the unknown parameters θ and then used that to decode the unknown state sequence from the second half of the data. In this testing setup, we always succeeded in recovering consistent state trajectories (result not shown). The inferred spatial topology reflects the internal representation of hippocampal ensemble spike activity.

3.2.2 Sensitivity to the Number of Place Cells, Threshold, and Bin Size. The outcome of derived neural representations certainly depends on the number of neurons and the coverage of their place fields. However, as far as there are sufficient numbers of place cells with receptive fields that cover the spatial environment, the ensemble representation shall remain similar (see more discussion in section 5.1). In the linear track example, we randomly selected a subset of neurons (25%–80%) and observed slightly degraded estimate in the transition matrix, but the derived topology graph was qualitatively similar (for more results, see section 4.2).

The exact shape of the derived topology graph also depends on the threshold. Thresholding is considering a necessary denoising step prior to data visualization. The purpose of thresholding is to discard the small state transition probability values (rare events) that may not be essential in the topology. In practice, a heuristic guideline for threshold selection is to examine the distribution (or histogram) of the state transition probability values. Typically that will be a bi- or multimodal distribution, and we often discard the long tails with small probability values. However, in the presence of small or noisy samples (e.g., irregular animal behavior, insufficient place field coverage of the environment), such a heuristic procedure may not be sufficient. (See section 4 for further discussion.)

We also note that the neural representation of spatial topology is rather robust with respect to varying temporal bin sizes (between 150 and 300 ms; results not shown). When the bin size is too small to capture the timescale of the run behavior, the correlation structure between place cells may be difficult to detect, and therefore the derived neural representation may be poor.

3.2.3 Comparison with Randomly Shuffled Spike Data. Based on the observed hippocampal neuronal population spiking activity, we recovered the embedded spatial topology and underlying state transition matrix that were imposed by both the spatial environment and the animal's behavior. This showed that our probabilistic model was capable of capturing the underlying structure of population codes. Presumably if the structure of the

population spiking activity was destroyed, either temporally or spatially or both, the consistent temporal firing patterns would be lost, and the inferred spatial topology would be misidentified.

To pursue the hypothesis testing, we randomly shuffled the population spike data and reran the analysis previously applied to the recorded spike data. Because different random initializations led to quantitatively and qualitatively different results, we repeated the procedure 100 times with independent Monte Carlo simulations. In the end, we obtained the empirical statistics from all Monte Carlo trials for both raw and shuffled spike data. We compared the results by examining their empirical free energy distributions (see Figure 3F). The distributions of free energy derived from the raw and shuffled data were significantly different (two-sample KS test, $P < 0.001$), and the median free energy value derived from the raw spike data was significantly greater than the one derived from the shuffled spike data (Mann-Whitney test, $P < 0.001$). Therefore, our model-based computational analysis paradigm provided an effective approach to assess the significance of the embedded data structure.

3.2.4 Illustration on the Mouse Data. For the 1D linear track experiment, we extended the same analysis to the population spiking activity recorded from mice. Generally place fields from the mouse hippocampus were more variable and nonlocalized (i.e., nonsparse) compared with those from the rat hippocampus (see Figure 2A). In other words, the mouse hippocampal neuronal activity was less sparse on the track. Even with a large number of recorded mouse hippocampal units (R2 data set), we found that the reconstructed state-space map was more ambiguous and had a sigmoidal shape at the ends of the track (see Figure 4A1). Increasing or decreasing the model size did not improve the result. This phenomenon was consistently observed in other mouse data sets (data not shown). Consequently, the inferred state transition matrix contained more errors in off-diagonal P_{ij} (where $|i - j| > 2$) entries (see Figure 4A2). After properly thresholding small P_{ij} values ("denoising"), the inferred topology graph reflected the topological nature of the 1D linear track: a ring structure with two shortcut loops (see Figure 4A3).

3.2.5 Quantifying the Variability of Hippocampal Population Codes in Single-Lap Analysis. The variability of animal foraging behavior would be expected to lead to variability in the ensemble spiking activity on a lap-by-lap basis. Therefore, it is important to quantify the variability of population codes with single-lap analyses. We applied our approach to compute the free energy score of the ensemble spike data from each single lap given the learned statistical model. To account for behavioral variability, we normalized the free energy score by the total number of temporal bins. As an illustration using the ensemble spike data from the linear track (see Figure 2A),
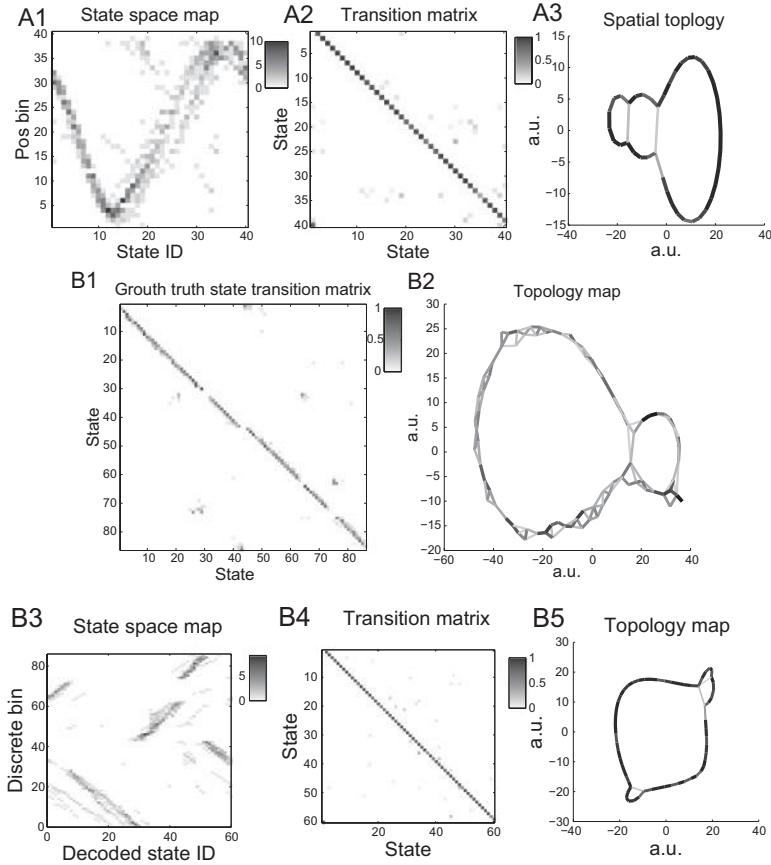
Figure 4: (A) Linear track (R2 data set) result. (A1) State-space map ($m = 40$). Grayscale bar represents the repeated frequency of mapping between the discrete state and discretized position bin of the linear track, and darker shading represents high frequency. (A2) One example of the estimated state transition matrix. (A3) Inferred spatial topology graph from the thresholded transition matrix. The estimated state-space map and topology graph appeared less satisfactory, as compared to the ones estimated from the rat data shown in Figures 3A and 3D. (B) T-maze (R5 data set) result. (B1) Ground truth state transition matrix inferred from the animal's behavior ($m = 86$ from linearization of T-maze shown in Figure 1B). (B2) Inferred spatial topology graph from the ground truth transition matrix. (B3) State-space map ($m = 60$ from model selection). Grayscale bar represents the repeated frequency of mapping between the discrete state and discretized position bin of the T-maze. (B4) One example of the estimated state transition probability matrix. (B5) Inferred spatial topology graph from the thresholded transition matrix. Due to thresholding of low-transition-probability values, the estimated topology graph appeared smoother compared to one shown in panel B2.

from laps 2 to 7, we computed the normalized free energy scores based on 50 independent Monte Carlo trials and computed the box plots and Monte Carlo median statistics (see Figure 2B). We then used the standard deviation (SD) statistic of the Monte Carlo median scores at single laps to evaluate the variability of the population codes. In the rat's case, the SD of the median normalized free energy scores from 6 laps was 3.52 ($m = 60$), whereas in the mouse's case, the SD of the median normalized free energy scores from 6 laps was 7.20 ($m = 40$). In both the rat and mouse data analyzed here, there was a statistically significant increase (paired Mann-Whitney test, $P < 0.01$) in the median normalized free energy scores from the first three laps (laps 2–4) to the consecutive three laps (laps 5–7), suggesting that there might be a nonspatial component (such as the velocity, acceleration, and novelty factor) in population codes. Of note, in the rat's behavior, the average running velocities during the six laps were 42.8, 40.7, 38.8, 42.8, 42.0, and 44.3 cm/s, respectively; in the mouse's behavior, the average running velocities during the six laps were 22.2, 27.9, 31.6, 33.3, 30.8, and 36.4 cm/s, respectively.

3.2.6 One-Dimensional Environment with Junctions. We also examined the experimental data recorded from the rat hippocampus while animals navigated in a 1D environment that consisted of one or more junctions in the track, such as the T-maze and H-maze. The higher behavioral diversity of navigation in such mazes would be expected to increase connectivity in their spatial topology graphs. In the context of topological space coding, the junctions play a similar role to spatial singularities or marked points (Dabaghian et al., 2011). Consequently, at the state that represents the junction point, the row of the state transition probability matrix would have two dominant P_{ij} values (see Figure 4B1, assuming roughly equal behavioral sampling of choices at the junction point). As expected, the task of recovering the spatial topology of such environments became more challenging. This was reflected by an increase in the number of local maxima (of the optimized free energy defined in equation 2.3), the frequency of nonunique state-space mapping, and the difficulty with graph interpretation.

In the case of the T-maze data (R5 data set), among a few evaluated models, we determined the optimal model size as $m = 60$ based on the Bayesian DIC. In one illustrative example, the state-space map between the decoded state and the linearized binned T-maze position (86 bins; see Figure 1B) showed five disjoint line stripes (states 1–8 for bins 63–74, states 9–30 for bins 1–31, states 31–43 for bins 43–62, states 44–51 for bins 75–86, and states 52–60 for bins 32–42; the darkness of the shading was linearly proportional to the count statistic) (see Figure 4B3). Note that there was also error or inconsistency in state-space mapping (overlapping mapping for discrete bins 1–20). The estimated state transition probability matrix (see Figure 4B4) had a few dominant off-diagonal P_{ij} elements that reflected the presence of one junction and the rat's behavior, specifically turns in the
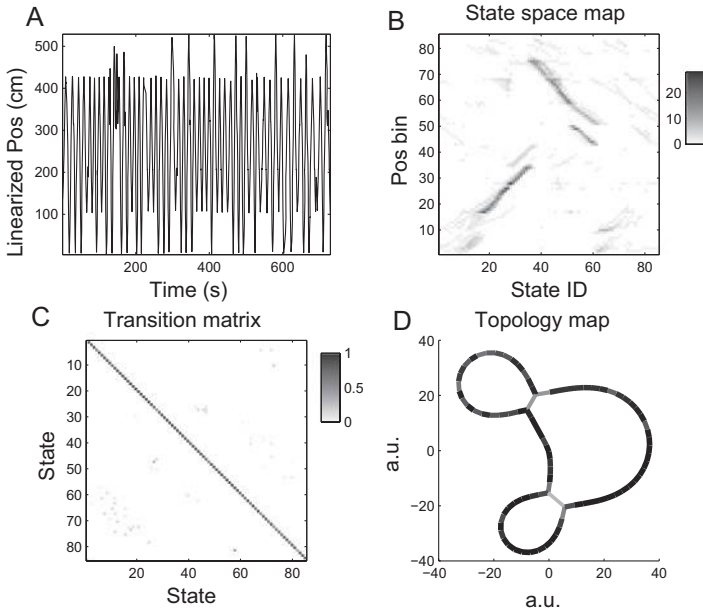
Figure 5: H-maze (R6 data set) result. (A) Linearized position of the H-maze. (B) State-space map ($m = 85$ from model selection). Grayscale bar represents the repeated frequency of mapping between the discrete state and discretized position bin of the H-maze. (C) One example of the estimated state transition probability matrix. (D) Inferred spatial topology graph from the thresholded transition matrix.

middle of the track. Consequently, the inferred topology map exhibited a ringlike structure with several shortcut edges (see Figure 4B5), bearing a resemblance to the ground truth topology map estimated from the animal's true behavior (see Figure 4B2).

In the case of the H-maze data (R6 data set), we first linearized the H-maze track (see Figure 5A), and then determined the optimal model size as $m = 85$ based on the Bayesian DIC. In one illustrative estimation result, the state-space map between the decoded state and the linearized and binned H-maze position showed several disjoint line stripes (see Figure 5B). Notice that some parts of the track were less frequently navigated and the corresponding state-space mapping for those regions was more ambiguous. There was also partial overlapping mapping between the state and position. The estimated state transition probability matrix had slightly more high-valued off-diagonal P_{ij} elements (see Figure 5C as compared to Figure 5D) that reflected the presence of the two junctions, as well as the rat turning in the middle of the track. By thresholding the small values of P_{ij}, the inferred

topology map appeared to have a multiple loop structure (see Figure 5D). Lowering the threshold would make the topology graph appear messier and thereby difficult to interpret. Therefore, thresholding was used as a necessary step to "denoise" the transition matrix and "unfold" the basic structure of the topology graph.

Thus far, we have succeeded in extracting the basic topological structure of several 1D environments, either without or with junctions. We found that the topology map of a linear track corresponded to a single ring (see Figure 3D), the topology map of a T-maze corresponded to a double ring (see Figure 4B5), and the topology map of a H-maze corresponded to a triple ring (see Figure 5D). These results were qualitatively consistent with the cartoon illustration shown in Figure 1 of our previous work (Chen et al., 2012).

3.3 Spatial Topology in Two-Dimensional Open Field Navigation. We further analyzed the ensemble spike data of rat hippocampal neurons while the animals navigated in a 2D open field environment. Within about 24 minutes, the freely foraging rat nearly occupied the whole environment (see Figure 1D), yet with nonequal spatial occupancy within the arena (see Figure 6A).

In the 2D environment, there were multiple ways to bin the space, either evenly or nonevenly. The bin size determines the number of states for representing the 2D environment. We used equal spatial bin size varying between 5 cm and 20 cm, which resulted in different numbers of position bins. Doubling the spatial resolution would increase the number of spatial bins by four. In order to seek a good trade-off among the spatial resolution, model complexity, and available data sample size, we also conducted a model selection analysis based on the Bayesian DIC (see Figure 6G). In the case of the R7 data set, although the Bayesian DIC favored the smallest model size (20 cm bin size), we chose the bin size of 15 cm and the model size of $m = 8 \times 9 = 72$ in order to have a reasonably sized state space. To illustrate the typical estimation result, we plotted the state-space map (see Figure 6C, which showed the median value of the spatial position that each state represented), the (sorted) state occupancy curve (see Figure 6D), as well as the estimated animal's spatial trajectories (see Figures 6E and 6F). Note that the state-space map and state occupancy curve were constructed only for the purpose of result assessment. In the illustrated example, the mean (median) estimation errors across the complete run trajectory in the x- and y-axes were 11.70 (9.02) cm and 11.71 (8.85) cm, respectively, which were both smaller than the actual spatial bin size of 15 cm. This result suggested that the decoding performance accuracy was quite reasonable, considering the fact that the decoder used only a pure hippocampal ensemble topological code in the complete absence of spatial information of the environment.

To test against the null hypothesis (H0) that the spike data were random and nonstructured, we shuffled the population spike count data (row and column shuffling as before) and repeated the random shuffling procedure 100 times, each based on independent initial conditions. We compared the Monte Carlo distributions of the free energy values derived from the raw and shuffled spike data. In addition, we tested against another alternative hypothesis (H1) that the raw spike data were observed in a 1D navigation environment (as opposed to a 2D navigation environment). As shown in Figure 6H, nonparametric statistical testing indicated that the median free energy obtained from fitting the raw spike data under the 2D-environment hypothesis (H2) was significantly greater than those obtained from the H0 and H1 hypotheses (Mann-Whitney test, $P < 0.001$). In addition, the H2 hypothesis free energy distribution was significantly different from both the H0 and H1 hypothesis free energy distributions (two-sample KS test, $P < 0.001$).

Given a spatially binned open field environment, we could compute the ground truth state transition matrix directly from the animal's behavior (see Figure 7A), from which the ground truth spatial topology graph was inferred (see Figure 7C). Based on the hippocampal neuronal ensemble spiking activity alone, we also inferred the state transition matrix (see Figure 7B) and the associated topology graph (see Figure 7D). Note that the diagonal elements of the transition matrix were first zeroed out before feeding the transition matrix to infer the topology graph (since the self-transition did not reveal any connectivity information). In comparing the ground truth (see Figures 7A and 7C) with the estimated results (see Figures 7B and 7D), we observed a close qualitative similarity. Basically, the ground truth topology graph reflected the viewpoint of an ideal external observer, which consisted of an evenly spaced grid. Notably, the topology graph inferred from the animal's hippocampal activity ("internal observer") appeared somewhat different in that the grid was nonevenly spaced and there were a few relatively large holes within the graph. This discrepancy might be ascribed to two plausible explanations: First, the externally observed behavior was not necessarily the same as what the animal perceived and represented in its neuronal activity. For example, it is well known that hippocampal neurons encode spatial as well as nonspatial features (Wiener, Paul, & Eichenbaum, 1989; Hampson, Simeral, & Deadwyler, 1999; Lenck-Santini, Rivard, Muller, & Poucet, 2005). Second, because we temporally binned the spiking activity into 250 ms bins, the spike count statistic could vary because of run velocity variability at different spatial locations (see Figure 6B). Generally the firing rates of hippocampal place cells are positively correlated to the run velocity of the animal (Zhang et al., 1998; Lu & Bilkey, 2012). In the open field environment, rats typically ran faster along the edges of the arena and slower across the central area of the arena. It is therefore plausible that in the rat's brain, the 2D environment might not be equally mapped or represented in Euclidean metric. For two running trajectories of equal distance, two spatial locations more quickly navigated would be represented as closer

A Occupany map

B

C

D

E

F

G

H
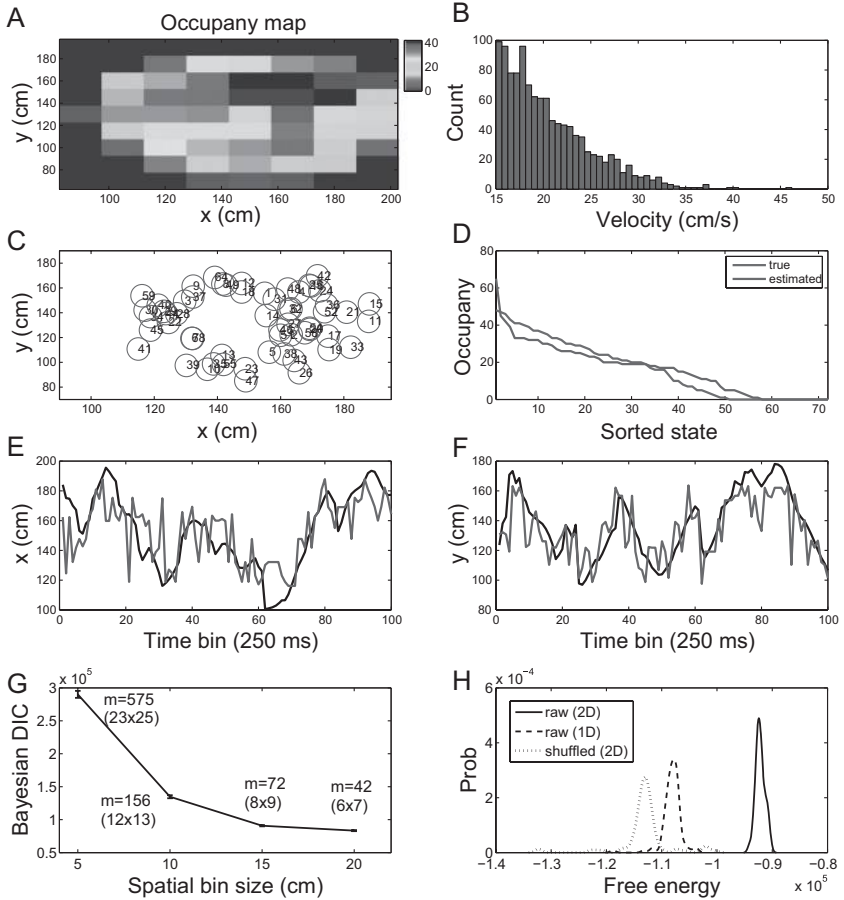
than two spatial locations more slowly navigated. To verify our conjecture, we did two additional tests. First, we increased or decreased the spatial bin size and repeated the analysis, but found that changing the spatial bin size did not affect this observation qualitatively. Second, we simulated a 20 min 2D random walk trajectory (with equal velocity at each step) along with population spike data (based on the real place fields estimated from the R7 data set) and repeated the analysis, and in this case we found that the holes disappeared (data not shown). This suggests that differences in run velocity may underlie inhomogeneous grid spacing represented by the ensemble spike data.

3.4 Quantification Derived from Topology Graphs. Given the inferred state transition matrix and topology graph, we can quantify the animal's

behavior in time even in the absence of spatial correlate. With the inferred state trajectory, we may construct a directed graph that consists of a number of nodes linked together by arrows known as edges. Arrows indicate the direction of the relationship between two nodes. The number of arrows ending at a specific node is called the *in-degree* of the node and the number of arrows leading from it is called the *out-degree*. In time, we can characterize the activeness of the animal's run behavior by counting the number of uniquely visited nodes (normalized by the total number of the nodes). Other quantitative measures, such as the visited frequency of the nodes and the occurrence frequency of state sequences, are also informative to characterize the animal's behavior (e.g., effects of experience or anxiety) or intended behavior (e.g., during the immobile or sleep state) in the absence of spatial correlate. (See Figures 7E to 7H for illustrations of such quantifications on two open-field experimental data sets, R7 and R8.)

3.4.1 Topological Complexity. We further investigate the question that given a derived topology graph, is there any quantitative measure for topological complexity that is independent of the shape and size of the graph as well as the between-node distance? In other words, the topological complexity shall reveal key information about the mapping between spatial

---

Figure 6: Open field (R7 data set) result. (A) Space occupancy map in an evenly binned 8×9 grid (15-cm spatial bin size). (B) Histogram of the rat's run velocity with a 250 ms temporal bin size. (C) One example of the estimated state-space map, where the median value of the spatial position that each state (circled number) represented was shown. From the numbered state sequence, one can infer the animal's trajectory in the open field. (D) The comparison of true and estimated (sorted) state occupancy curves, which both showed a skewed distribution. (E, F) Snapshot comparison between the true (black) and estimated (blue) animal's run trajectory in the 2D environment (based on the mapping in panel C). The mean estimation errors across the complete run trajectory on the x- and y-axes are 11.70 cm and 11.71 cm, respectively. (G) Model selection for spatial bin size based on the Bayesian deviance information criterion (DIC). Error bar statistics were computed from the results of 50 independent Monte Carlo trials. (H) Distributions of the free energy value derived from the H2 hypothesis with raw (red solid curve) and shuffled (dotted curve) spike count data, and from the H1 hypothesis with raw data (blue dashed curve). All free energy values were obtained based on the same model size ($m = 72$). Probability distributions of the free energy were estimated from 100 random Monte Carlo trials. The median free energy derived from the H2 hypothesis with raw data (red solid curve) was significantly greater than those derived from the H0 and H1 hypotheses (Mann-Whitney test, $P < 0.001$). The free energy distribution of H2 (red) was significantly different from the others (two-sample KS test, $P < 0.001$). A color version of the figure is provided in the online supplementary material.
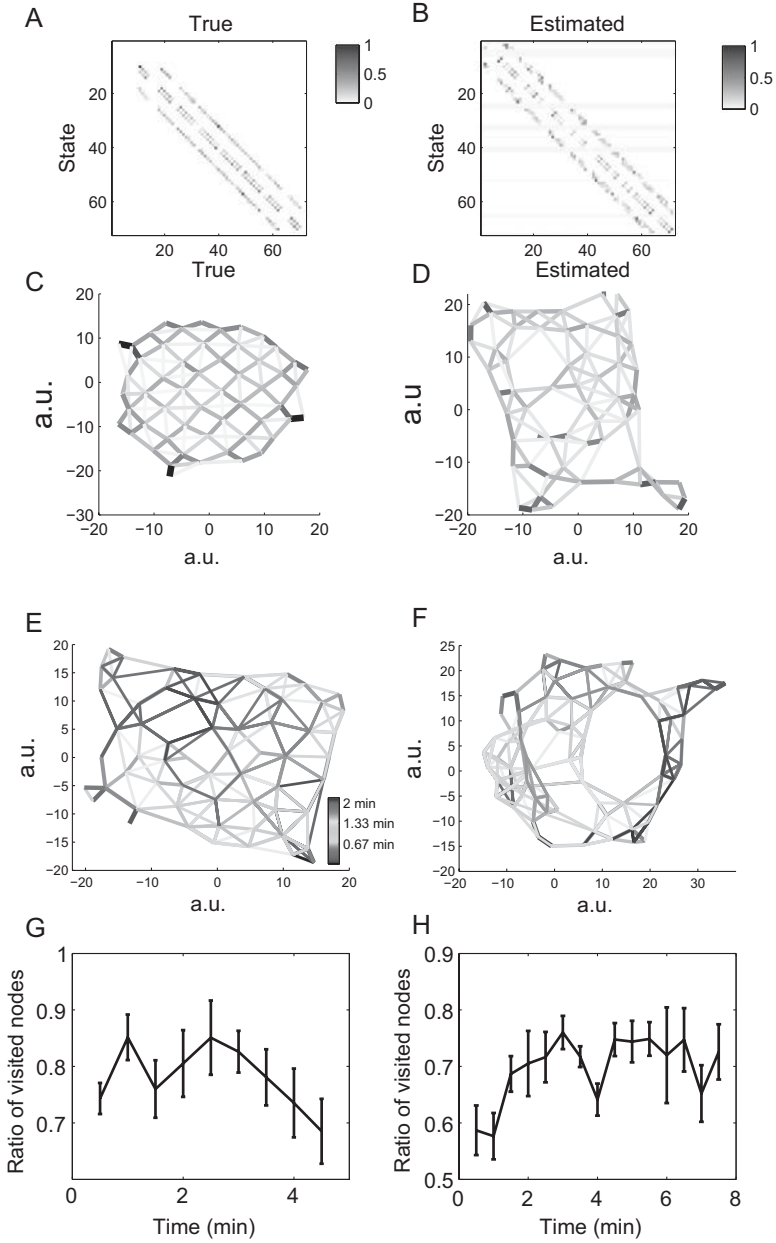
locations and the states (or nodes). Given the inferred state-space map, we can assess the state density for every spatial location in the environment, and we define the topological complexity T_d by the mean state density across all spatial locations. Specifically, the state density measures the degree of ambiguity about the number of unique states required to represent a specific spatial location. The state density is also independent of the spatial sampling and behavioral occupancy. For instance, a topographic map representation would imply $T_d = 1$ since there is only one-to-one mapping between the spatial locations and states. In theory, a linear track would have $T_d = 2$ due to the direction ambiguity (i.e., two states are required to represent the same spatial location). For a similar reason, the junction points of the T-maze or H-maze would have a higher state density due to many-to-one state-space mapping. As expected, the 2D open field environment would have a higher mean state density than the 1D environments.

To confirm our conjecture, we computed the heat maps of state density for four topologically different environments from previous examples (see Figure 8). As expected, the topological complexity T_d is increased from the linear track (R1, $m = 60$), to T-maze (R5, $m = 60$), to H-maze (R6, $m = 85$), and to open field (R7, $m = 72$), with increasing estimated mean state density (or T_d) of 2.02, 2.35, 2.40, and 4.15, respectively. The computation of state density also took into account the size of m by properly adjusting the distance metric. In all heat maps, the spatial bin size is 5 cm; therefore, the state density unit is $5\,\mathrm{cm}^{-1}$ (1D environment) or $25\,\mathrm{cm}^{-2}$ (2D environment). Interestingly, the respective median state density estimates are 2, 2, 2, and 4, revealing the true difference in topological complexity between 1D and 2D environments.

Figure 7: Open field (R7 data set) result. (A) Ground truth state transition probability matrix estimated from the animal's behavior ($m = 72$). (B) One example of the state transition probability matrix estimated from the ensemble spike data, which shared a similar structure as panel A. In panels A and B, the diagonal elements $\{P_{ii}\}$ of the matrices were zeroed out before feeding the transition matrices to infer topology graphs. (C, D) Ground truth topology graph (from the ground truth transition matrix) and inferred topology graph (from the estimated transition matrix), respectively. (E, F) Topology graphs inferred from hippocampal ensemble activity for the open field environment (left column: R7 data set; right column: R8 data set). All axes of the panels are in au. Nodes correspond to states. In each panel, the animal's run trajectory during the first 2 minutes was color-coded. (G, H) Quantification of the animal's behavior in terms of the number of visited nodes within the graph shown above every 30 seconds (normalized by the total number of visited nodes). Error bar shows the SD estimated from 10 Monte Carlo runs. A color version of this figure is provided in the online supplementary material.
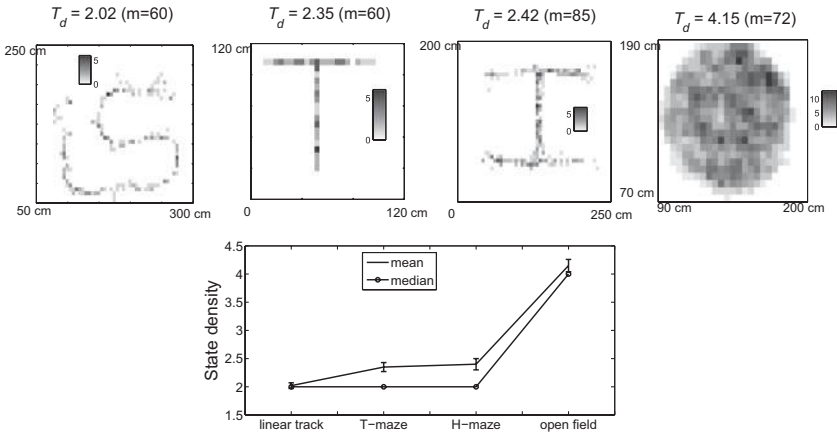
Figure 8: (Top row) Heat maps of inferred state density (in color bar) for four spatial environments (shown in Figure 1), with estimated topological complexity T_d (mean state density) and state dimensionality m shown above. In all heat maps, the spatial bin size is 5 cm; therefore, the state density unit is 5 cm^{-1} or 25 cm^{-2}. Note that the state density is typically higher at the bifurcation points in the T-maze and H-maze. (Bottom panel) Estimated mean and median state density statistics for the above four spatial environments. Error bar shows the SD statistic from 10 Monte Carlo runs.

3.5 Hippocampal Neuronal Representation During Quiet Wakefulness. Next, we examined the ensemble spiking activity from the rat hippocampus in a linear track environment when the animal was awake and stationary (i.e., nonlocomotion, velocity threshold: <2cm/s), which we refer to as the off-line awake state. During the quiet wakefulness, ripple activity (150–250 Hz) was frequently observed in the hippocampal local field potential (LFP), which was associated with brief multi-unit activity (MUA) bursts. We used a previously established criterion (Davidson et al., 2009) to detect these MUA bursts and examined the associated hippocampal ensemble spiking activity. Many groups have reported the compressed reactivation of spatial sequences (replay events) in this state, in association with sharp wave ripples (SPW-R) (Foster & Wilson, 2006; Diba & Buzsáki, 2007; Davidson et al., 2009; Karlsson & Frank, 2008; Carr, Jadhav, & Frank, 2011; Jadhav, Kemere, German, & Frank, 2012).

We adapted our method to detect such replay events in the off-line awake state. We first binned the ensemble spiking activity with a 250 ms bin size during periods of locomotion (velocity threshold >15 cm/s). We then binned the ensemble spiking activity with a 25 ms bin size (compression factor of 10) during periods of hippocampal SPW-R-associated MUA bursts. In the learning mode, we estimated the state transition matrix
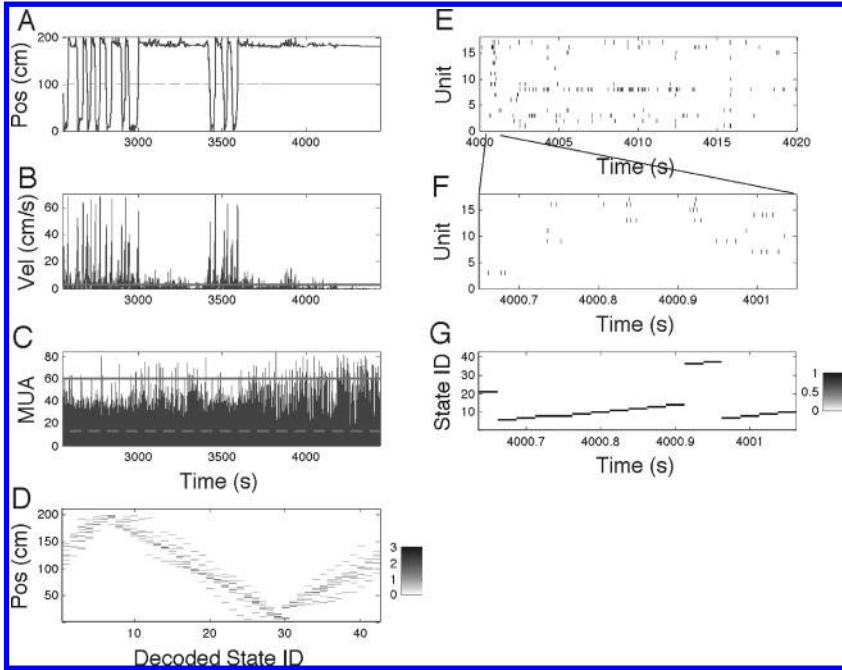
Figure 9: Detection of the reactivation of spatial sequences on a linear track (R4 data set). (A) Animal's position on the track (blue line). The horizontal red line indicates the the stationary awake period where the animal had low velocity (<2 cm/s) and high SPW-R-associated hippocampal multiunit activity bursts (MUA > mean +4 SD). (B) Animal's run velocity, where the horizontal line represents the 2 cm/s velocity threshold. (C) Aggregated MUA across all tetrodes, where the horizontal dashed and solid lines represent the mean and mean + 4 SD of the MUA, respectively. (D) State-space map inferred from the locomotion period. (E, F) Snapshots of the ensemble spiking activity from 18 single units during quiet wakefulness. (G) Reconstructed state trajectory for a replay event from panel *F*. Note that in our replay detection approach, the virtual position can be "filled in" even in the absence of spiking activity in some time bins. A color version of this figure is provided in the online supplementary material.

based on the ensemble spike data from the locomotion period (R4 data set, $m = 42$). In the testing mode, we fed in the ensemble spike data during quiet wakefulness using the previously estimated state transition matrix. Specifically, we computed the free energy score of the data and the posterior probability trace of the state trajectory. An illustrative example is shown in Figure 9. In Figures 9A and 9B, the animal's position and run velocity on the track are displayed as a function of time. Figure 9C shows the hippocampal

MUA aggregated across all seven tetrodes, and Figure 9D shows the state-space map inferred from the data during RUN epochs. The state-space mapping had some ambiguity mainly due to very short recording of RUN epochs (2.5 min and 9.5 full laps). Within the period of hippocampal SPW-R-associated MUA bursts (see Figure 9E), we applied the decoding analysis to hand-picked ensemble spike data where many (>8) units were active (see Figure 9F) and estimated the posterior probability of the state trajectory (see Figure 9G). As an illustrated result shown in Figure 9G, we recovered one long (sequence length >4) state trajectory (9 consecutive time bins of state ID: $6 \rightarrow 7 \rightarrow \cdots \rightarrow 14$) that had high posterior probability. From the inferred state-space map (see Figure 9D), the replayed state trajectory corresponded to an inbound-run trajectory from around 175 cm to 100 cm on the track. The virtual replay speed was about 333 cm/s, 10 to 15 times faster than the behavioral velocity (see Figure 9B). The compression factor was consistent with the scale reported in previous findings (Foster & Wilson, 2006; Davidson et al., 2009). For each candidate replay event, we also ran the inference procedure multiple times, each based on a different initial condition. Based on many Monte Carlo trials, we could score the candidate replay event by its repeated frequency (via state-space mapping) and predictive free energy and then determine its statistical significance using a predefined criterion. If the replay event was significant, the repeated frequency and predictive free energy of the replayed spatial sequence would be presumably high; in contrast, the predictive free energy of a random spatial sequence derived from randomly shuffled spike data would be lower. Nevertheless, we did not delve into this issue further since it was beyond the focus of this study.

4 Topological Versus Topographic Coding of Spaces

There has been debate on the role of hippocampal population codes between topographic and topological representations of the spatial environment. In section 3, we focused on the sufficiency of topological representation. However, neither the necessity nor the efficiency of a topological code has been studied as compared to the topographic code. In this section, we investigate a few specific questions. Given the well-known fact that velocity information is encoded in hippocampal rate codes (Zhang et al., 1998; Hirase, Czurkó, Csicsvari, & Buzsáki, 1999; Huxter, Burgess, & O'Keefe, 2003; Lu & Bilkey, 2012), is it possible for hippocampal neurons to utilize such metric information to accommodate a fuzzy and mixed representation or to construct a topological map with topographic constraints? From a computational perspective, what is the optimal spatial representation that achieves a trade-off between decoding accuracy and encoding efficiency? How does the topological representation vary with respect to the number of the neurons? How does the favor of the topological versus topographic coding (in terms of the decoding accuracy) change with respect to the number

of neurons? The answers to these questions are important and relevant not only in neuroscience but also in engineering applications (such as robotics).

4.1 Comparison of Decoding Efficiency Between Topological and Topographic Codes. A full topographical map implies topological information. However, it is unclear whether a full topographical map is necessary for the hippocampus to adequately represent the spatial environment. To address this question, we compared the decoding performance of a pure topological code, a topographically constrained code, and a pure topographic code in an open field environment (R7 data set). In a pure topological code, there are no a priori place fields, and no spatial correlate can be used in encoding. In contrast, in a pure topographic code, the spatial correlate is used in encoding to estimate the place fields of individual units, which are subsequently used for decoding. For place fields, we used a nonparametric model based on the assumption of Poisson firing rate statistics (see section 2.6). A topographically constrained code is intermediate between purely topographic and purely topological, such that some of the states in the topological map are provided by topographic (actual spatial) information and are topographically constrained as landmarks in the environment.

First, we systematically varied the percentage level of topographic constraints in the topological code and compared the decoding accuracy. A code with 0% topographic constraint implies a pure topological code (decoding result illustrated in Figures 6E and 6F), whereas a code with 100% topographic constraints is essentially a pure topographic code (in this case, place fields were estimated and applied to maximum likelihood decoding analysis). In the case of a pure topographic code, by varying the spatial bin size from 15 cm to 10 cm to 5 cm, we obtained the respective average median decoding errors as 11.13 cm, 8.25 cm, and 5.59 cm. To begin to consider a topological code with topographic constraints, we randomly assigned the spatial bins, either uniformly or proportional to the spatial occupancy frequency in space (see Figure 6A), with the actual (known) spatial locations as landmarks. For a spatially-binned 8×9 open field environment (with 15 cm bin size), the percentage of the topographic constraints varied from 0, 10, 20, 40, 60, 80, to 100. Except for the 0 and 100% levels, at each percentage level, we repeated random samplings with different assignments 20 times to compute the mean decoding error. Thus, decoding error was evaluated over all epochs of run behavior based on the animal's actual positions. The result is shown in Figure 10A. As expected, as more topographic constraints were imposed, the decoding accuracy gradually improved to a minimum of 7.5 cm (half bin size at 100% level). In contrast, for the finest pure topographic code (with 5 cm spatial bin size), the mean and median decoding errors are 10.08 cm and 5.59 cm, respectively.

To gain further insight into the topological code without or with topographic constraints, we hypothesize an idealized situation where the animal has sampled uniformly (i.e., equal navigation occupancy) and sufficiently
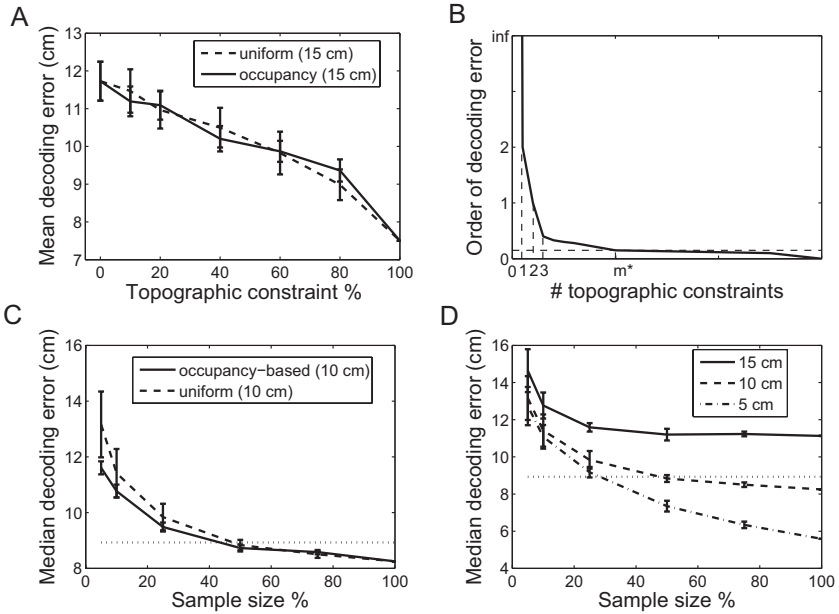
Figure 10: (A) Comparison of decoding accuracy as a function of percentage of topographic constraints, where the error bar shows the SD among 20 Monte Carlo trials with random sampling. The code with 0% topographic constraints implies a pure topological code, whereas the code with 100% topographic constraints is essentially a pure topographic code. (B) An idealized curve between the order of average decoding error and the number of topographic constraints varying from 0 to m. According to certain criteria, an optimal number $3 < m^* < m$ may be identified to achieve a desirable trade-off between the average decoding error and the number of topographic constraints. (C) Decoding accuracy of topographic codes as a function of sample size with strategies of uniform sampling and occupancy-based sampling. (D) Decoding accuracy of topographic codes as a function of sample size with varying spatial bin size. In panels C and D, horizontal lines represent the average decoding error from the pure topological code given 15 cm bin size.

in a circular environment (with radius 1 and area π). Without loss of generality, uniform landmark sampling in the environment is also assumed. When the number of landmarks is 0, there is no spatial correlate, and one cannot establish any correspondence between the abstract spatial topology and actual environment. In this limit, the decoding error is infinity (without relying on occupancy and state-space map). When the number of topographic constraints is 1 (say, the landmark state i is in the center), given any nonzero state transition probability P_{ik} or P_{ki} ($1 \leq \forall k \leq m, k \neq i$) and

the velocity information, one can estimate the between-state distance d_{ik}; however, since there is no direction information, there will be uncertainty (radius of d_{ik} around the center) about the kth state location in the environment. In this case, the average error for mapping the remaining $m - 1$ states is on the scale of the width of environment (diameter 2). When the number of topographic constraints is 2 (say, the landmark states i and j are along the diameter of the circular environment), given any nonzero state transition probability P_{ik}, P_{jk}, P_{ki} or P_{kj} ($1 \leq \forall k \leq m, k \neq i, j$), one can estimate the between-state distances d_{ik} and d_{jk} based on the velocity. Since in this case, there is more metric information, the uncertainty and average decoding error are reduced (roughly to the scale of 1). As three topographic constraints are imposed, given the estimated between-state distances, we can use these three spatial locations to triangulate the remaining $(m - 3)$ state locations. Further addition of landmarks will impose increasing constraints on the environment, and the decreased decoding error will gradually saturate. Ultimately the uncertainty and decoding error will fall to 0, as the number of topographic constraints approaches m. Under certain criteria, one can envision an optimal number $3 < m^* < m$ that achieves an optimal trade-off between the decoding error and the number of topographic constraints. (See Figure 10B for an idealized illustration.) In practice, the value of m^* may vary as a function of the animal's spatial occupancy, behavioral sampling (and the induced state transition statistics), and the size of m.

Next, in order to test the efficiency and robustness of the topographic code while fixing the number of neurons, we systematically varied the percentage of the sample size (i.e., the number of temporal bins along with the spatial correlate during run epochs) available for encoding analysis. A topographic code with 100% sample size would produce the best achievable decoding accuracy, and the decoding accuracy would degrade with decreasing sample size. In contrast, the performance of a pure topological code would remain unchanged because of its encoding-free nature. We varied the sample size percentage (5%, 10%, 25%, 50%, 75%, and 100%) by sampling either uniformly or proportional to the spatial occupancy frequency in time, and computed the median decoding error. Regardless of the sample size, the decoding analysis was always evaluated on the total run epochs based on the animal's actual positions. In this example, $T = 2356$ temporal bins and $m = 72$ spatial bins were used to represent an area of 1.26 m^2. The result is shown in Figure 10C. As fewer samples were used in encoding, the decoding accuracy of a topographic code degraded (first slowly and then dramatically around 50%). This phenomenon was present in both uniform sampling and occupancy-based sampling, but the performance was worse in uniform sampling. This effect also persisted for various choices of spatial bin size (see Figure 10D, assuming uniform sampling). Although the exact statistics could depend on the animal's run behavior, we expect this phenomenon would generalize across behavior. In this example, when the sample size was below 25%, the decoding accuracy of the topological code

(with 15 cm bin size) outperformed that of the topographic code (even with the finest bin size). Note, however, that if the spatial bin size of the topological code was reduced to 10 cm or 5 cm, the decoding performance of the topological code deteriorated, partially due to the insufficient sample size required to infer a large state space. In principle, animals can use an adaptive spatial representation (e.g., fine/low resolution in high/low behavioral occupancy regions) to compensate for nonuniform spatial sampling.

4.2 Spatial Representation with Respect to Number of Neurons. Spatial topological representation of population codes depends on the hippocampal place cells. In order to examine how the topological representation changes with respect to the number of place cells, we randomly selected a subset of neurons (10%, 25%, 50%, 75%) to run the same analysis for both 1D (R1 data set) and 2D (R7 data set) environments. We fixed the state dimensionality $m = 60$ for the 1D environment and $m = 72$ for the 2D environment. As a general observation, as the number of neurons decreased, the quality of the reconstructed state-space map reduced or became more ambiguous (see representative examples in Figure 11A, as compared to Figure 3A), so did the state transition matrix (result not shown). However, the derived topology graph was relatively robust for both environments (see Figure 11B). Specifically, in the 2D open field environment, as fewer number of neurons were included, the detail of the topology graph started to be missing (possibly due to the lack of place field coverage in certain spatial locations); however, the 2D grid structure was still preserved.

In addition, we repeated the same decoding analysis as Figure 10C and compared the decoding accuracy between topological and topographical codes with varying numbers of neurons. The results on the median and median decoding error across 10 random trials are shown in Figure 11C. As expected, the performance of both topological and topographical codes degraded with the decreasing number of neurons. Furthermore, it appeared that the topological code had overall superior performance in decoding accuracy, especially in the presence of small number of neurons. In the 2D example here, the topographic code only obtained a lower average median decoding error with more than 36 neurons (75% out of 47).

To summarize the findings of this section, we have demonstrated the merits and limitations of topological and topographic codes in spatial representation and decoding accuracy. A pure topographic code would imply a precise characterization of the place fields with respect to the full environment. Such a code would be optimal for decoding, but it is not necessarily efficient for encoding, especially in the presence of a large environment and small sample size. In contrast, an abstract topological code with a small number of topographic constraints may achieve a good trade-off between decoding accuracy and encoding efficiency. For the same practical reason, it is more efficient for downstream neurons of the hippocampus to decode an animal's spatial location based on a "fuzzy" semitopographic map without
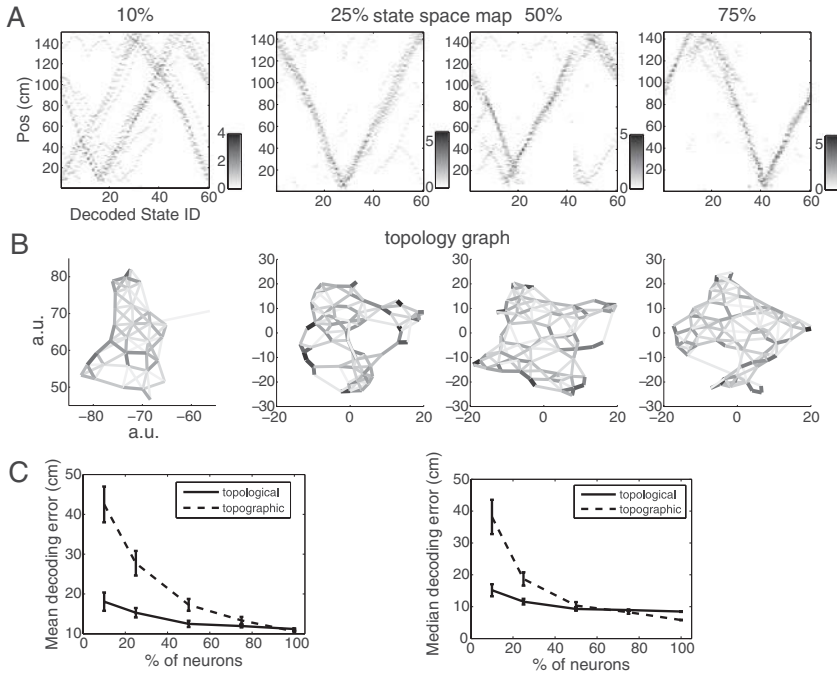
Figure 11: (A) Examples of the derived state-space map for the 1D linear track environment ($m = 60$, R1 data set) with varying numbers of neurons (10%, 25%, 50%, 75% out of 74). (B) Examples of the derived topology graph for the 2D open field environment ($m = 72$, R7 data set) with varying numbers of neurons (10%, 25%, 50%, 75% out of 49). (C) Comparison of the mean and median decoding errors between the topological and topographic codes in the same 2D environment with varying numbers of neurons. Error bar was obtained from 10 random trials.

storing or accessing the place fields of neuronal ensembles. In addition, the topological code is more robust in terms of spatial representation and decoding accuracy, especially with fewer neurons.

5 Discussion

Although the physiology of rodent hippocampal cells has been well studied (Nakazawa, McHugh, Wilson, & Tonegawa, 2004; Ahmed & Meht, 2009), it remains unclear how ensembles of hippocampal place cells could represent abstract spatial topology through spiking activity alone. Understanding the internal representation of rodent hippocampal population codes from the perspective of the internal observer would provide insight into the

population coding mechanism. Here we investigate this problem using a computational approach based on statistical modeling and inference.

Our approach employs unsupervised learning for exploratory analysis of neural ensemble spike data, aiming to discover the inherent temporal structure embedded in population codes. This approach is generalizable beyond the rodent hippocampus, and in principle can be applied to recordings from other neural circuits in both rodents and primates. For modeling other systems, we can employ a parametric neuronal tuning curve model that incorporates covariates (Chen, Vijayan, Barbieri, Wilson, & Brown, 2009; Chen, Putrino, Ghosh, Barbieri, & Brown, 2011). Such a model-based approach allows us to examine ensemble spiking activity in greater detail, such as quantifying the representation variability of single-lap data. We can test hypotheses about the dimensionality (1D versus 2D) of spatial topology embedded in hippocampal population codes, and we can also compare the coding efficiency between topological and topographic codes. Specifically, the quantitative measure of topological complexity may be useful not only in revealing the dimensionality of the spatial topology but also allowing us to compare hippocampal ensemble representations under different behavioral conditions in the same environment.

5.1 Triangulating Embedded Topological Spaces. A general mathematical theory of triangulating topological spaces has been previously discussed in Curto and Itskov (2008) and Dabaghian et al. (2011, 2012). In theory, provided that the spatial environment is covered by a number of overlapping rodent hippocampal place fields that are convex, the space can be triangulated. Consequently, when the place cells fire consistently in time whenever the animal forages in the same spatial location, one can construct cell groups whose spiking activity coverage represents the topological subsets from which the embedded topological space can be triangulated. Accordingly, as long as the space is fully covered by place fields, the topological representation of the spatial environment will remain relatively unchanged even in the presence of plasticity or remapping of individual place fields.

5.2 Relation to Other Methods and Limitation of Our Approach. Recently a few computational approaches have been proposed to uncover the spatial topology maps from hippocampal population codes (Curto & Itskov, 2008; Dabaghian et al., 2012). These approaches are based on two common assumptions: (1) the temporal pattern of hippocampal neuronal firing (especially co-firing) is key to decoding spatial information, and (2) since cofiring implies spatial overlap of place fields, a map encoded by cofiring of place cells will be based on connectivity and adjacency, which will give arise to a topological map.

Specifically, the basic idea of the algebraic analysis approach (Curto & Itskov, 2008) is (1) to define some cell groups consisting of place cells

that collectively fire within a 2 theta cycle (250 ms) window, (2) extract the topological features (homology groups) from the cell groups, (3) reconstruct a topology graph for every cell group and a distance matrix for any pairwise cell groups, and (4) to run a nonmetric multidimensional scaling (MDS) algorithm to visualize the distance matrix, thereby revealing the embedded spatial topology. In their theoretical assumptions, the hippocampal place fields were convex (except for a small percentage of multipeaked place fields), and the union of the place fields covered the entire space. In practice, these assumptions are not always valid; in addition, temporal smoothing of the spike trains and selection of proper threshold values in steps (2) and (3) are crucial in their analysis (K. Marku, personal communication).

In the investigation of two open field experimental data sets (R7 and R8), we have also analyzed the hippocampal ensemble spike data using software of the algebraic approach (http://www.math.unl.edu/~vitskov2/software/SpikeTop/). We found that the results derived from the algebraic analysis were less intuitive and difficult to interpret. In addition, the final results were rather sensitive to the thresholding and smoothing operations (data not shown due to space limitation), whereas the outcomes of our derived topology graph were more robust with respect to different initial conditions (with or without thresholding). This might be due to the fact that our probabilistic model explicitly exploits the Markovian structure of the population codes. Provided that the Markovian structure of $\{S(t)\}$ is discarded and the spike activity at any time bins is marginally independent, then all methods will boil down to estimating the correlations between place cells.

In our statistical approach, the derived topology graph relies on an accurate estimate of the state transition matrix. As discussed in section 3, thresholding is a useful step to denoise the estimated state transition matrix. In addition to the heuristics, another possible solution is to systematically search for a suboptimal threshold along a thresholding path, and the selected threshold would reflect a good trade-off between the interpretability of the graph and the topological complexity. Nevertheless, how to choose sensible threshold value to produce a reliable topology graph remains a subject for future study.

5.3 Neural Representation of Spatial Maps and Temporal Sequences. In rodents, various forms of spatial maps are represented in the hippocampus-entorhinal cortex network. It was suggested that the hippocampus functions as an associative or attractor network (Tsodyks, 1999) that encodes the abstract topological map of the space, which serves as a locus of the animal's spatial awareness (Burgess & O'Keefe, 1996; Muller, Stead, & Pach, 1996), whereas other parts of the circuit (such as the entorhinal cortex) provide metric information with a complementary representation of the environment. Although we have restricted our focus to hippocampal place cells, other types of cells, such as grid cells, head

direction cells, and border cells, also contribute to spatial representation and navigation (Leutgeb, Leutgeb, Moser, & Moser, 2005; McNaughton, Battaglia, Jensen, Moser, & Moser, 2006; Mathis, Herz, & Stemmler, 2012).

Hippocampal place fields are well known to be plastic in novel environments (Frank, Stanley, & Brown, 2004; Frank, Brown, & Stanley, 2006) and to remap with manipulation of spatial cues and with progressive morphing of the environment (Gothard, Skaggs, Moore, & McNaughton, 1996; Colgin, Moser, & Moser, 2008; Leutgeb, Leutgeb, Treves, Moser, & Moser, 2004; Leutgeb, Leutgeb, Treves, et al., 2005). These findings suggested the topological coding nature of the rodent hippocampus. One theoretical hypothesis of topological coding is that Voronoi tessellation of many overlapping place fields provides a triangulation of the space and that a topological representation is achieved by chaining these neighboring subspaces together through sequential neuronal activity (Dabaghian et al., 2011).

The rodent hippocampus is also responsible for representing episodic memory in the form of ordered temporal sequences (Hasselmo, 2012). The representation of topological maps can be viewed as a form of ordered spatial sequence, where the place cells encode the "relative" space while animals forage in the environment. When the space is morphed, the place fields also change accordingly. In most rodent behavioral tasks, space and time are strongly coupled and nearly inseparable (Redish, Rosenzweig, Bohanick, McNaughton, & Barnes, 2000). Recently, the question of a hippocampal representation of time has been studied by disentangling the spatial component (Pastalkova, Itskov, Amarasingham, & Buzsáki, 2008; MacDonald, Lepage, Eden, & Eichenbaum, 2011). In a temporal sequence memory task, during the delay periods hippocampal neuronal ensembles temporally organized and disambiguated distinct sequences of events that comprised specific repeated experiences (MacDonald et al., 2011). It would be interesting to use our approach to investigate these "time cells" and "time fields."

5.4 Efficiency of Topological and Topographic Codes. From a computational decoding perspective, a topological code is more efficient than a pure topographic code, especially in the presence of a large environment and sparse sample size. It is impractical to restore every piece of topographic information in the real world. To improve accurate spatial representation, a topological code can be incorporated with some metric information and topographic landmarks, producing a fuzzy semitopographic map. Animals may use different behavioral strategies (such as nonperiodic stopping or orienting) to identify landmarks inside the environment. In addition, other neurons, such as border cells and visual cortical cells, can provide additional cues for the environment, which together may produce a coherent spatial representation. Such a spatial mapping strategy has been widely used in robotic localization and mapping.

5.5 Implications Beyond Hippocampal Population Codes. The probabilistic interpretation of neural population codes has become an active research topic (Zemel, Dayan, & Pouget, 1998; Ma, Beck, Latham, & Pouget, 2006; Beck, Ma, Latham, & Pouget, 2007). An accurate statistical characterization of neural codes would enable us to better understand neural representations at both single cell and population levels. This would also help to develop efficient algorithms to decode neural ensemble spiking activity online in brain-machine interfaces (Kemere et al., 2008). Characterizing the internal representation of ensemble spiking activity in terms of hidden states facilitates our intuitive understanding of data acquired during various tasks, including eye movement, motor planning, associative learning, motor planning, and reaching movement (Seidemann, Meilijson, Abeles, Bergman, & Vaadia, 1996; Rainer & Miller, 2000; Smith et al., 2004; Jones, Fontanini, Sadacca, & Katz, 2007; Czanner et al., 2008; Yu et al., 2009; Chen et al., 2009; Lawhern, Wu, Hatsopoulos, & Paninski, 2010). In addition, the variability in behavior and task-related neural responses at the single cell level would be expected to contribute to variability in population responses. Statistical approaches may provide a quantitative way to examine such noisy ensemble spiking activity across trials.

The concept of topological coding of a stimulus is not limited to the representation of space in the rodent hippocampus. In fact, similar experimental findings have been reported in the primate primary motor cortex (M1) (Georgopoulos, 2002) and primary visual cortex (V1) (Singh et al., 2008). Georgopoulos (2002) discussed motor topology and the neural representations of direction, size, and location of movement in space. Specifically, he reviewed the results of of behavioral studies of drawing geometrical figures and discussed how the representation of topological features can be invariantly extracted from primate M1 neuronal populations irrespective of attributes of size, location, and the muscles effecting the motor trajectory. In another context, Singh et al. (2008) cast the fundamental question about neural representation in terms of the topological structure of primate V1 neuronal ensemble spiking activity. They reported that the topological structure of V1 spiking patterns was similar to those evoked by natural image stimulation and consistent with the topology of a 2-sphere in 3D space. Other brain areas such as the restrosplenial cortex (RSC) (Wyss & Van Groen, 1992; Vann, Aggleton, & Maguire, 2009) might also use topological coding for kinematic stimuli.

6 Conclusion

We have applied a computational approach to examine the neural representation of spatial topology in the rodent hippocampus. Using probabilistic modeling and Bayesian inference without access to spatial correlates of the neuronal activity, we successfully extracted the underlying topological features (e.g., junctions and dimensionality) of the environment, in addition to

certain metric (semitopographic) information (e.g., track length) embedded in the population code. The topological representation of space embedded in rodent hippocampal population codes was robust. Our results raise several testable questions: How would the internal representation of population codes change with manipulation of hippocamal function, such as partial hippocampus lesions (Kesner & Novak, 1982), transgenic manipulations of hippocampal circuitry (McHugh, Blum, Tsien, Tonegawa, & Wilson, 1996; Suh, Rivest, Nakashiba, Tominaga, & Tonegawa, 2011), or Alzheimer's disease models (Cacucci, Yi, Wills, Chapman, & O'Keefe, 2009)? In addition, how can our approach be extended to the analysis of sleep-associated ensemble spike data in which there is no behavioral correlate to compare against? These questions will be the topic of our future investigations.

Appendix

A.1 Electrophysiology and Recording. Male Long-Evans rats or mice were implanted with microdrive arrays containing between 11 and 24 tetrodes. The microdrive arrays were implanted above either the right dorsal hippocampus or bilaterally above both hippocampi in one rat (Yamamoto & Wilson, 2008). The tetrodes were slowly lowered into the brain reaching the cell layer of CA1 two to four weeks following the date of surgery. Each tetrode's signal was split into spike and LFP channels. Spike channels were filtered between 300 Hz and 6 kHz, and the LFP channels were filtered between 1 Hz and 450 Hz. Thresholds were set on each tetrode channel and waveform shape was saved from each tetrode channel whenever an individual channel passed a fixed threshold. A differential reference electrode was left in the white matter immediately dorsal to the hippocampus. Recorded spikes were manually clustered and sorted to obtain single units using a custom software (XClust, M.A.W.). Once stable units were obtained, the rodents were allowed to explore freely in the environment. We used diodes placed on the head of each animal to track position with a sampling rate of 30 frames/s.

A.2 Animal Behavior. Under different experimental protocols, rodents (rats or mice) freely foraged in specific spatial environments, including a linear track, T-maze, end-to-end T-maze ("H-maze"), and open field (see Figure 1). In the linear track, T-maze and open field, there were no behavioral restrictions placed on the animals, and the rodents were encouraged to explore the environment with rewards (e.g., chocolate sprinkles, fruit juice) distributed at the end of the tracks or inside the open field. In the H-maze, the animals performed a spatial working memory task involving paired sample and test epochs, as described in Jones and Wilson (2005). In the sample phase, rats ran in the "force" direction to turn left (L) or right (R) to win reward. In the test phase, rats ran the "choice" direction and found reward by choosing the reward site that required a non-matching behavior

(turning R or L). To identify the period of rodent locomotion during spatial navigation, we used a velocity threshold (>10 cm/s for mice and >15 cm/s for rats) to select the RUN epochs and merged them together.

References

Ahmed, O. J., & Meht, M. R. (2009). The hippocampal rate code: Anatomy, physiology and theory. *Trends in Neurosciences*, *32*, 329–338.

Barbieri, R., Frank, L. M., Nguyen, D. P., Quirk, M. C., Solo, V., Wilson, M. A., et al. (2004). Dynamic analyses of information encoding in neural ensembles. *Neural Computation*, *16*, 277–307.

Battista, G. D., Eades, P., Tamassia, R., & Tollis, I. G. (1998). *Graph drawing: Algorithms for the visualization of graphs*. Upper Saddle River, NJ: Prentice Hall.

Beck, J. M., Ma, W. J., Latham, P. E., & Pouget, A. (2007). Probabilistic population codes and the exponential family of distributions. In P. Cisek, T. Drew, & J. F. Kalaska (Eds.), *Progress in brain research, 165* (pp. 509–519). Amsterdam: Elsevier.

Broome, B. M., Jayaraman, V., & Laurent, G. (2006). Encoding and decoding of overlapping odor sequences. *Neuron*, *51*, 467–482.

Brown, E. N., Frank, L. M., Tang, D., Quirk, M. C., & Wilson, M. A. (1998). A statistical paradigm for neural spike train decoding applied to position prediction from ensemble firing patterns of rat hippocampal place cells. *Journal of Neuroscience*, *18*, 7411–7425.

Burgess, N., & O'Keefe, J. (1996). Cognitive graphs, resistive grids, and the hippocampal representation of space. *Journal of General Physiology*, *107*, 659–662.

Buzsáki, G. (2006). *Rhythms of the brain*. New York: Oxford University Press.

Cacucci, F., Yi, M., Wills, T. J., Chapman, P., & O'Keefe, J. (2009). Place cell firing correlates with memory deficits and amyloid plaque burden in Tg2576 Alzheimer mouse model. *Proceedings of the National Academy of Sciences, USA*, *105*, 7863–7868.

Carr, M. F., Jadhav, S. P., & Frank, L. M. (2011). Hippocampal replay in the awake state: A potential physiological substrate of memory consolidation and retrieval. *Nature Neuroscience*, *14*, 147–153.

Chen, Z., Kloosterman, F., Brown, E. N., & Wilson, M. A. (2012). Uncovering spatial topology represented by rat hippocampal population neuronal codes. *Journal of Computational Neuroscience*, *33*, 227–255.

Chen, Z., Putrino, D., Ghosh, S., Barbieri, R., & Brown, E. N. (2011). Statistical inference for assessing functional connectivity of neuronal ensembles with sparse spiking data. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, *19*, 121–135.

Chen, Z., Vijayan, S., Barbieri, R., Wilson, M. A., & Brown, E. N. (2009). Discrete- and continuous-time probabilistic models and algorithms for inferring neuronal UP and DOWN states. *Neural Computation*, *21*, 1797–1862.

Colgin, L. L., Moser, E. I., & Moser, M. B. (2008). Understanding memory through hippocampal remapping. *Trends in Neurosciences*, *31*, 469–477.

Curto, C., & Itskov, V. (2008). Cell groups reveal structure of stimulus space. *PLoS Computational Biology*, *4*, e1000205.

Czanner, G., Eden, U. T., Wirth, S., Yanike, M., Suzuki, W. A., & Brown, E. N. (2008). Analysis of between-trial and within-trial neural spiking dynamics. *Journal of Neurophysiology*, *99*, 2672–2693.

Dabaghian, Y., Cohn, A. G., & Frank, L. M. (2011). Topological coding in the hippocampus. In *Computational modeling and simulation of intellect: Current state and future prospectives* (pp. 293–320). Hershey, PA: IGI Global.

Dabaghian, Y., Memoli, F., Frank, L. M., & Carlsson, G. (2012). A topological paradigm for hippocampal spatial map formation using persistent homology. *PLoS Computational Biology*, *8*, e1002581.

Davidson, T. J., Kloosterman, F., & Wilson, M. A. (2009). Hippocampal replay of extended experience. *Neuron*, *63*, 497–507.

Diba, K., & Buzsáki, G. (2007). Forward and reverse hippocampal place-cell sequences during ripples. *Nature Neuroscience*, *10*, 1024–1242.

Foster, D. J., & Wilson, M. A. (2006). Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature*, *440*, 680–683.

Frank, L. M., Brown, E. N., & Stanley, G. B. (2006). Hippocampal and cortical place cell plasticity: Implications for episodic memory. *Hippocampus*, *16*, 775–784.

Frank, L. M., Stanley, G. B., & Brown, E. N. (2004). Hippocampal plasticity across multiple days of exposure to novel environments. *Journal of Neuroscience*, *24*, 7681–7689.

Georgopoulos, A. P. (2002). Behavioral and neural aspects of motor topology: Following Bernstein's thread. In M. Latash (Ed.), *Progress in motor control, Vol. 2: Structure-function relations in voluntary movements* (pp. 1–12). Champaign, IL: Human Kinetics.

Georgopoulos, A. P., Schwartz, A. B., & Kettner, R. E. (1986). Neuronal population coding of movement direction. *Science*, *233*, 1416–1419.

Gothard, K. M., Skaggs, W. E., Moore, K. M., & McNaughton, B. L. (1996). Binding of hippocampal CA1 neural activity to multiple reference frames in a landmark-based navigation task. *Journal of Neuroscience*, *16*, 823–835.

Hampson, R. E., Simeral, J. D., & Deadwyler, S. A. (1999). Distribution of spatial and nonspatial information in dorsal hippocampus. *Nature*, *402*, 610–614.

Hasselmo, M. (2012). *How we remember: Brain mechanisms of episodic memory*. Cambridge, MA: MIT Press.

Hirase, H., Czurkó, A., Csicsvari, J., & Buzsáki, G. (1999). Firing rate and theta-phase coding by hippocampal pyramidal neurons during "space clamping." *European Journal of Neuroscience*, *11*, 4373–4380.

Huxter, J., Burgess, N., & O'Keefe, J. (2003). Independent rate and temporal coding in hippocampal pyramidal cells. *Nature*, *425*, 828–832.

Jadhav, S. P., Kemere, C., German, P. W., & Frank, L. M. (2012). Awake hippocampal sharp-wave ripples support spatial memory. *Science*, *336*, 1454–1458.

Jones, L. M., Fontanini, A., Sadacca, B. F., & Katz, D. B. (2007). Natural stimuli evoke analysis dynamic sequences of states in sensory cortical ensembles. *Proceedings of the National Academy of Sciences, USA*, *104*, 18772–18777.

Jones, M. W., & Wilson, M. A. (2005). Theta rhythms coordinate hippocampal-prefrontal interactions in a spatial memory task. *PLoS Biology*, *3*, 2187–2199.

Karlsson, M. P., & Frank, L. M. (2008). Network dynamics underlying the formation of sparse, informative representations in the hippocampus. *Journal of Neuroscience*, *28*, 14271–14281.

Kemere, C., Santhanam, G., Yu, B. M., Afshar, A., Ryu, S. I., Meng, T. H., et al. (2008). Detecting neural-state transition using hidden Markov models for motor cortical prostheses. *Journal of Neurophysiology*, *100*, 2441–2452.

Kesner, R., & Novak, J. (1982). Serial position curve in rats: Role of the dorsal hippocampus. *Science*, *218*, 173–175.

Lawhern, V., Wu, W., Hatsopoulos, N. G., & Paninski, L. (2010). Population decoding of motor cortical activity using a generalized linear model with hidden states. *Journal of Neuroscience Methods*, *189*, 267–280.

Lenck-Santini, P. P., Rivard, B., Muller, R. U., & Poucet, B. (2005). Study of CA1 place cell activity and exploratory behavior following spatial and nonspatial changes in the environment. *Hippocampus*, *15*, 356–369.

Leutgeb, J., Leutgeb, S., Treves, A., Meyer, R., Barnes, C., McNaughton, B. L., et al. (2005). Progressive transformation of hippocampal neuronal representations in morphed environments. *Neuron*, *48*, 345–358.

Leutgeb, S., Leutgeb, J., Moser, M. B., & Moser, E. I. (2005). Place cells, spatial maps and the population code for memory. *Current Opinion in Neurobiology*, *15*, 738–746.

Leutgeb, S., Leutgeb, J., Treves, A., Moser, M. B., & Moser, E. I. (2004). Distinct ensemble codes in hippocampal areas CA3 and CA1. *Science*, *305*, 1295–1298.

Lu, X., & Bilkey, D. K. (2012). The velocity-related firing property of hippocampal place cells is dependent on self-movement. *Hippocampus*, *20*, 573–583.

Ma, W. J., Beck, J. M., Latham, P. E., & Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nature Neuroscience*, *9*, 1432–1438.

MacDonald, C. J., Lepage, K. Q., Eden, U. T., & Eichenbaum, H. (2011). Hippocampal "time cells" bridge the gap in memory for discontiguous events. *Neuron*, *71*, 737–749.

Markus, E., Qin, Y., Leonard, B., Skaggs, W., McNaughton, B., & Barnes, C. (1995). Interactions between location and task affect the spatial and directional firing of hippocampal neurons. *Journal of Neuroscience*, *15*, 7079–7094.

Mathis, A., Herz, A.V.M., & Stemmler, M. (2012). Optimal population codes for space: Grid cells outperform place cells. *Neural Computation*, *24*, 2280–2317.

McGrory, C. A., & Titterington, D. M. (2009). Variational Bayesian analysis for hidden Markov models. *Australian and New Zealand Journal of Statistics*, *51*, 227–244.

McHugh, T. J., Blum, K. I., Tsien, J. Z., Tonegawa, S., & Wilson, M. A. (1996). Impaired hippocampal representation of space in CA1-specific NMDAR1 knockout mice. *Cell*, *87*, 1339–1349.

McNaughton, B. L., Battaglia, F. P., Jensen, O., Moser, E. I., & Moser, M. B. (2006). Path integration and the neural basis of the "cognitive map." *Nature Reviews Neuroscience*, *7*, 663–678.

Muller, R. U., Stead, M., & Pach, J. (1996). The hippocampus as a cognitive graph. *Journal of General Physiology*, *107*, 663–694.

Nakazawa, K., McHugh, T. J., Wilson, M. A., & Tonegawa, S. (2004). NMDA receptors, place cells and hippocampal spatial memory. *Nature Reviews Neuroscience*, *5*, 361–372.

Nirenberg, S., & Latham, P. E. (1998). Population coding in the retina. *Current Opinion in Neurobiology*, *8*, 488–493.

O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map: Preliminary evidence from unit activity in the freely-moving rat. *Brain Research*, *34*, 171–175.

O'Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map*. New York: Oxford University Press.

Pastalkova, E., Itskov, V., Amarasingham, A., & Buzsáki, G. (2008). Internally generated cell assembly sequences in the rat hippocampus. *Science*, *321*, 1322–1327.

Quian Quiroga, R., & Panzeri, S. (2009). Extracting information from neuronal populations: Information theory and decoding approaches. *Nature Reviews Neuroscience*, *10*, 173–185.

Rainer, G., & Miller, E. K. (2000). Neural ensemble states in prefrontal cortex identified using a hidden Markov model with a modified EM algorithm. *Neurocomputing*, *32-33*, 961–966.

Redish, A., Rosenzweig, E., Bohanick, J. D., McNaughton, B. L., & Barnes, C. A. (2000). Dynamics of hippocampal ensemble activity realignment: Time versus space. *Journal of Neuroscience*, *20*, 9298–9309.

Sanger, T. D. (2003). Neural population codes. *Current Opinion in Neurobiology*, *13*, 238–249.

Schwartz, A. B. (1994). Direct cortical representation of drawing. *Science*, *265*, 540–542.

Seidemann, E., Meilijson, I., Abeles, M., Bergman, H., & Vaadia, E. (1996). Simultaneously recorded single units in the frontal cortex go through sequences of discrete and stable states in monkeys performing a delayed localization task. *Journal of Neuroscience*, *16*, 752–768.

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*, 1–18.

Smith, A. C., Frank, L. M., Wirth, S., Yanike, M., Hu, D., Kubota, Y., et al. (2004). Dynamic analysis of learning in behavioral experiments. *Journal of Neuroscience*, *24*, 447–461.

Suh, J., Rivest, A. J., Nakashiba, T., Tominaga, T., & Tonegawa, S. (2011). Entorhinal cortex layer III input to the hippocampus is crucial for temporal association memory. *Science*, *334*, 1415–1420.

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

Vann, S., Aggleton, J., & Maguire, E. (2009). What does the retrosplenial cortex do? *Nature Reviews Neuroscience*, *10*, 792–802.

Wiener, S. I., Paul, C. A., & Eichenbaum, H. (1989). Spatial and behavioral correlates of hippocampal neuronal activity. *Journal of Neuroscience*, *9*, 2737–2763.

Wilson, M. A., & McNaughton, B. L. (1993). Dynamics of the hippocampal ensemble code for space. *Science*, *261*, 1055–1058.

Wilson, M. A., & McNaughton, B. L. (1994). Reactivation of hippocampal ensemble memories during sleep. *Science*, *265*, 676–679.

Wyss, J., & Van Groen, T. (1992). Connections between the retrosplenial cortex and the hippocampal formation in the rat: A review. *Hippocampus*, *2*, 1–12.

Yamamoto, J., & Wilson, M. A. (2008). Large-scale chronically implantable precision motorized microdrive array for freely behaving animals. *Journal of Neurophysiology*, *100*, 2430–2440.

Yu, B. M., Cunningham, J. P., Santhanam, G., Ryu, S. I., Shenoy, K. V., & Sahani, M. (2009). Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *Journal of Neurophysiology*, *102*, 614–635.

Zemel, R., Dayan, P., & Pouget, A. (1998). Probabilistic interpretation of population codes. *Neural Computation*, *10*, 403–430.

Zhang, K., Ginzburg, I., McNaughton, B. L., & Sejnowski, T. J. (1998). Interpreting neuronal population activity by reconstruction: unified framework with application to hippocampal place cells. *Journal of Neurophysiology*, *79*, 1017–1044.

_____