Cell Assembly Sequences Arising from Spike Threshold Adaptation Keep Track of Time in the Hippocampus

Vladimir Itskov,^{1,2*} Carina Curto,^{1,3*} Eva Pastalkova,^{4,5} and György Buzsáki⁵

¹Department of Mathematics, University of Nebraska–Lincoln, Lincoln, Nebraska 68588-0130, ²Center for Theoretical Neuroscience, Columbia University, New York, New York, 10032, ³Courant Institute of Mathematical Sciences, New York University, New York, New York 10012, ⁴Janelia Farm Research Campus, Howard Hughes Medical Institute, Ashburn, Virginia 20147, and ⁵Center for Molecular and Behavioral Neuroscience, Rutgers University, Newark, New Jersey 07102-1814

Hippocampal neurons can display reliable and long-lasting sequences of transient firing patterns, even in the absence of changing external stimuli. We suggest that time-keeping is an important function of these sequences, and propose a network mechanism for their generation. We show that sequences of neuronal assemblies recorded from rat hippocampal CA1 pyramidal cells can reliably predict elapsed time (15–20 s) during wheel running with a precision of 0.5 s. In addition, we demonstrate the generation of multiple reliable, long-lasting sequences in a recurrent network model. These sequences are generated in the presence of noisy, unstructured inputs to the network, mimicking stationary sensory input. Identical initial conditions generate similar sequences, whereas different initial conditions give rise to distinct sequences. The key ingredients responsible for sequence generation in the model are threshold-adaptation and a Mexican-hat-like pattern of connectivity among pyramidal cells. This pattern may arise from recurrent systems such as the hippocampal CA3 region or the entorhinal cortex. We hypothesize that mechanisms that evolved for spatial navigation also support tracking of elapsed time in behaviorally relevant contexts.

Introduction

Tracking time is of fundamental importance in a wide range of brain operations, including sensory perception and motor actions, learning, memory, planning, decision-making and language (Gibbon et al., 1997; Buonomano and Karmarkar, 2002; Ivry and Spencer, 2004; Mauk and Buonomano, 2004; Buhusi and Meck, 2005). Despite the central importance of temporal processing, its underlying neural mechanisms remain unknown. At the systems level, two competing ideas have been put forward: timing is generated by a central mechanism and distributed to various brain regions (Church, 1984), or each subsystem produces its own timing (Mauk and Buonomano, 2004). With regard to timing and duration, a distinction is made between subsecond (perceptual-motor) and suprasedond (cognitively mediated) scales (Michon, 1985; Lewis et al., 2003). At the level of mechanisms, two models are typically distinguished; clocks and ramping time keepers, with neuronal substrates in the cerebellum, basal ganglia, prefrontal, motor, and parietal cortical regions (cf. Mauk and Buonomano, 2004; Buhusi et al., 2005). The hippocampus has also been implicated in timing (Clark and

Isaacson, 1965; Thompson and Krupa, 1994; Young and McNaughton, 2000), although the mechanism has remained elusive.

We report here on a novel form of time-tracking mechanism, which is manifested by evolving transiently active cell assemblies and is accurate for periods of tens of seconds. First, we examine the ability of evolving neuronal sequences to predict elapsed time in a memory task. Second, we propose a simple network model with Mexican-hat-type connectivity and adaptation of the membrane potential thresholds for action potential generation that is similar to what has been observed in rodent hippocampus (Henze and Buzsáki, 2001) and in fish (Chacron et al., 2007). The threshold-adaptation model reproduces the key properties of the observed sequences, suggesting that time-keeping in the hippocampus may arise from the same cellular and network mechanisms that support spatial navigation.

Materials and Methods

The experimental data used in this paper were adopted from Pastalkova et al. (2008), where all relevant experimental methods and protocols are described. The animals were all male rats.

Time prediction from experimental data. Two versions of time prediction models were fit from experimental data: a “rate-only” model and a “phase-only” model. These models used firing rates or theta phases of spikes to fit a probability distribution for population spiking activity in 0.5 s time windows. Maximum-likelihood estimation was then used on single trials to predict the most likely time reflected by the population activity at each time bin. In all cases, the models were fit from complementary trials, never including the trial on which time was subsequently inferred. The time estimates corresponding to each time bin were obtained independently, so that only knowledge of the current population activity was needed to estimate elapsed time. See supplemental Text

Received Oct. 31, 2010; revised Dec. 6, 2010; accepted Dec. 14, 2010.

V.I. was supported by the Swartz Foundation and National Science Foundation (NSF) Grant DMS 0967377; C.C. was supported by NSF Grant DMS 0920845 and a Courant Instructorship; E.P. was supported by the Patterson Trust and Howard Hughes Medical Institute; and G.B. was supported by the J.D. McDonnell Foundation, NSF Grant SBE 0542013, National Institutes of Health Grant NS034994, and National Institute of Mental Health Grant MH54671.

*V.I. and C.C. contributed equally to this work.

Correspondence should be addressed to György Buzsáki at the above address. E-mail: buzsa@andromeda.rutgers.edu.

DOI:10.1523/JNEUROSCI.3773-10.2011

Copyright © 2011 the authors 0270-6474/11/312828-07\$15.00/0

Figure 1. Time prediction from sequential neural activity in a memory task. **A**, Average raster over 18 s for a population of simultaneously recorded neurons during wheel runs after preceding left trials in a spontaneous alternation task. The rat was required to make a correct choice for water reward on the basis of its previous choice. Neurons are ordered according to the time of peak firing rate in the PSTH. **B**, Population raster for a single trial, using the same neuron ordering as in **A**. **C**, Time prediction for a single trial (shown in **B**) using time prediction models fit from all other trials. In each time bin, elapsed time in the running wheel is inferred either from the population firing rate vector (red) or the firing phases of active cells with respect to the theta oscillation (purple). In each case, the prediction approximates well the true time (black). **D**, Average errors (in seconds) of the time prediction for three rats (green, blue, and red) as a function of time, using the time prediction model via rate. **E**, A reliability measure was computed for each trial. At each time t , the measure quantifies the extent to which population vectors from single trials resemble trial-averaged population vectors at time t better than trial-averaged population vectors at other times (see Materials and Methods). Both the average reliability across trials (black) and the reliability for the single trial (blue) from **B** are well above 0, the value expected for completely unreliable data. Gray lines denote a SD above and below the mean for the distribution of reliabilities across trials. **F**, Reliability of individual trials in the control, home cage condition. The wheel was placed next to the animal's home cage, and the animal ran in the wheel without memory requirement or reward contingency. Conventions as in **E**; the single trial (blue) has been chosen at random.

(available at www.jneurosci.org as supplemental material) for a more detailed description of the time prediction models.

Threshold adaptation model. We modeled network dynamics using a standard firing rate model, with threshold nonlinearity $[y]_+ = \max(y, 0)$. At any point in time, the vector $\tilde{x}(t) = (x_1(t), \dots, x_N(t))$ represents a population vector of firing rates for each of N neurons. A key

ingredient of the model is the activity-dependent adaptation of the “spike” thresholds of individual neurons, represented by the dynamic variables $\tilde{h}(t) = (h_1(t), \dots, h_N(t))$. The model can thus be described by a system of $2N$ equations:

$$\begin{cases} \tau_m \dot{x}_i = -x_i + \left[\sum_{j=1}^N J_{ij} x_j + I_i - h_i \right]_+ & (1) \\ \tau_a \dot{h}_i = -h_i + c x_i, & (2) \end{cases}$$

where $\tau_m = 30$ ms and $\tau_a = 2$ s are membrane and threshold-adaptation time constants, respectively, J is the matrix of synaptic weights for the recurrent network, and I is a (time-dependent) vector of external inputs. The constant $c = 0.5$ controls the strength of the activity-dependent adaptation, whereas τ_a determines the time scale with which the thresholds recover in the absence of firing.

The connectivity matrix J is constant in time and can be written as a sum of two components, $J = J^0 + J^{\text{het}}$, where J^0 has Mexican-hat-type connectivity on a two-dimensional lattice of neurons with periodic (torus) boundary conditions, and J^{het} is a matrix of heterogeneous weights, sampled randomly and independently from a normal distribution with mean zero. In simulations, we generated three different instances of J^{het} , resulting in three different connectivity matrices J^1, J^2 , and J^3 .

In simulations, two types of behavioral conditions were distinguished: “task” and “home cage.” In all cases, initial conditions consisted of the same “bump” of firing rate activity $\tilde{x}(0) = \tilde{x}_0$ in the center of the two-dimensional lattice of neurons, and differed in the adaptation variables $\tilde{h}(0) = \tilde{h}_0$ only. Neurons with adapted thresholds were chosen to lie at the left, bottom, right, or top of the initial bump of activity, resulting in initial conditions A to D, respectively. For task trials, the initial conditions were consistent from trial to trial, and the model was driven by temporally and spatially unstructured noise $I(t)$; different instances of noise was thus the only difference between trials of the same initial condition type. In the home cage trials, the initial conditions A to D were randomized across trials, and the model was driven by spatially unstructured noise that had temporal correlations on the order of 125 ms (see supplemental Text, available at www.jneurosci.org as supplemental material, for further information).

Layer 2 simulations. To investigate the ability of a downstream layer to “inherit” the sequence generated by the threshold adaptation model (“layer 1”), we simulated activity in a second layer connected to the first via sparse and random feedforward projections. The dynamics in this layer were governed by the same Equations 1 and 2 in the previous layer, with

two differences. First, the connectivity matrix J in layer 2 represented an overall global inhibition (to ensure sparse firing) and had no spatial structure. Second, the input vector $I(t)$ had two components: temporally and spatially unstructured noise; and feedforward input derived from the activity in the previous layer via random and sparse connections (10% connection

probability between layers) (see supplemental Text, available at www.jneurosci.org as supplemental material for further information).

Reliability measure. A reliability measure was used to compare population vectors from individual trials to the mean. First, we normalized firing rates so that each neuron had the same average firing rate, calculated over all time bins and all trials. Let $\tilde{v}_i(t)$ be the population vector corresponding to the i th trial at time t . For each fixed time bin t_0 , we computed a “reliability score” $R_i(t_0)$ for each trial by computing the squared-distance between the population vector $\tilde{v}_i(t_0)$ and the mean vector across trials $\langle \tilde{v} \rangle(t_0)$, $W_i(t_0) = \|\langle \tilde{v} \rangle(t_0) - \tilde{v}_i(t_0)\|^2$. This distance was then compared with the average squared-distance between $\tilde{v}_i(t_0)$ and the mean population vectors $\langle \tilde{v} \rangle(t)$ at all other times, $B_i(t_0) = \frac{1}{N_t - 1} \sum_{t \neq t_0} \|\langle \tilde{v} \rangle(t) - \tilde{v}_i(t_0)\|^2$, where N_t is the total number of time bins. If a trial is reliable at time t_0 , then $W_i(t_0) < B_i(t_0)$. The reliability was thus defined to be

$$R_i(t_0) = \frac{-W_i(t_0) + B_i(t_0)}{W_i(t_0) + B_i(t_0)},$$

with the denominator chosen so that $-1 \leq R_i(t_0) \leq 1$. A reliability of 1 corresponds to the “best case” scenario of $\tilde{v}_i(t_0) = \langle \tilde{v} \rangle(t_0)$, and -1 corresponds to the “worst case” scenario of $\tilde{v}_i(t_0) = -\langle \tilde{v} \rangle(t_0)$ for all $t \neq t_0$, and $\tilde{v}_i(t_0) \neq \langle \tilde{v} \rangle(t_0)$. If $\tilde{v}_i(t_0)$ is closer to the mean vector $\langle \tilde{v} \rangle(t_0)$ than it is to the mean vectors from other time bins $\langle \tilde{v} \rangle(t)$, then $R_i(t_0)$ will be positive. For random (and thus completely unreliable) population vectors, $R_i(t_0) \approx 0$. The average reliability $R(t)$ was computed as the trial-average of the time series vectors $R_i(t)$.

Results

Elapsed time is well estimated by the sequence of cell assembly activation

A reliable pattern of sequential activation of neuronal activity was observed in the CA1 region of hippocampus during the delay period of a memory task (Pastalkova et al., 2008; cf. Gill et al., 2010; Kraus et al., 2010; Macdonald et al., 2010). Three rats were trained to run for ~ 20 s in a running wheel during the delay period before making a choice (left or right) for the next run through a T maze. Action potentials for pyramidal cells were recorded together with local field potentials. The pattern of sequential activation of simultaneously recorded neurons for a given trial type (Fig. 1*A,B*) was reliable across trials and lasted 10–20 s without repeating itself. Therefore, we hypothesized that the population spiking activity of pyramidal neurons in CA1 at any point in time during a trial could be used to infer elapsed time.

To verify this hypothesis, we designed two probabilistic models for inferring elapsed time from instantaneous neural activity: one based on the cells’ firing rates, and the other using phases of spikes with respect to the theta oscillation (see Materials and Methods). Both models were good predictors of elapsed time on single trials (Fig. 1*C*). The average error of time estimation by the rate model was 1–2 s in rats 2 and 3 and 2–3 s in rat 1 (Fig. 1*D*). Similar errors were observed using the phase model (supplemental Fig. 1, available at www.jneurosci.org as supplemental mate-

Figure 2. The threshold adaptation model. **A**, Neurons are arranged on a two-dimensional “sheet” with periodic boundary conditions, so that the top and bottom (red) edges are identified, and the left and right (black) edges are identified. In the absence of threshold-adaptation, the network activity quickly converges to a bump attractor (bottom left), with the center of the bump marked by the asterisk. In the presence of adaptation, however, the bump of activity moves continuously around the sheet, never stabilizing. The black curve traces the position of the center of the bump; three bump positions (1, 2, and 3) are depicted by shaded circular regions. **B**, Different initial conditions in the adaptation variables lead to different trajectories for the bump of activity. Initial condition A has a bump of firing rate activity centered at *, with adapted thresholds only in the neurons to the left (red shaded region). Initial condition B has the same initial firing rate activity centered at *, but with adapted thresholds only for the neurons immediately below (blue shaded region). Adapted neurons are less likely to fire, so initial condition (init cond) A causes the bump of activity to move initially to the right (red trajectory 1), whereas initial condition B causes the bump of activity to move upward (blue trajectory 1). Due to periodic boundary conditions, the trajectories “wrap around” the sheet; the numbers indicate the temporal ordering of the different components of each (red, blue) trajectory. **C**, Snapshots of the moving bump of activity at different times (0.5, 2.5, and 5 s) following an initial condition, from a single trial simulated using the model. The bump retains its shape, so the population activity over time can be described by tracking only the center of the bump. The bottom right panel shows the trajectory of the center of the bump up until 5 s. **D**, Bump trajectories for 10 s of simulated data. Each panel corresponds to a different initial condition (A to D), differing only in the adapted thresholds of neurons and having the same initial firing rate activity (centered at *). For each initial condition, 3 randomly selected trajectories are shown. The trajectories are reliable, despite noisy (unstructured) input that changes from trial to trial.

rial). The accuracy of time estimation from the models increased with the number of cells used in each animal (data not shown). This observation suggests that by recording from a much larger fraction of hippocampal neurons, the accuracy of time estimation can be improved further. It also suggests that a greater amount of information is available to structures downstream from CA1 than what was available for our method of inference. Since we obtained timing precision on a behaviorally relevant scale, we hypothesize that the brain could use population activity to estimate elapsed time.

To investigate how behavioral relevance might influence timekeeping, we also recorded the activity of CA1 cells during running in a wheel placed adjacent to the home cage of the rat (control condition). The rat could enter the wheel and run at its leisure and was not required to keep track of elapsed time. While the patterns of neuronal activity on individual runs displayed some semblance of sequential activation near the beginning of the wheel run, the overall sequences were not consistent from trial to trial (supplemental Fig. 2, available at www.jneurosci.org as supplemental material); as a result, the time prediction models did not yield any statistically significant prediction of time on the control data (data not shown).

While the reliability $R(t)$ of sequential activity during the memory task was well above 0 (the expected value for random, unreliable data; see Materials and Methods) (Fig. 1*E*), $R(t)$ was not significantly positive for the control (home cage) data (Fig. 1*F*). This suggests that the patterns of neuronal activity during wheel runs reflect timing information only in behaviorally relevant contexts.

Figure 3. The threshold adaptation model produces long-lasting, temporally reliable sequences. **A**, The average raster over 25 trials of simulated data for the task A trial type (same initial conditions). Neurons are ordered according to the timing of their peak firing rates in the average over trials. Trials differ in the particular instance of random noise superimposed on a constant external drive to the network. **B**, Same as **A**, but for the task B trial type (different initial condition from **A**). Neuronal sequences are again present, but the sequential firing of individual neurons is different from that in task A. The same connectivity matrix has been used in both cases. **C**, The task B trial average, ordered according to the task A cell ordering. The sequential activity is no longer apparent, indicating that the sequence in **B** is different from that in **A**. **D**, **E**, A single trial raster, using the same neuron ordering as in **A** and **B**, respectively. **F**, A single trial in the control home cage condition, ordered according to peak firing rates in the average across home cage trials. Except at the very beginning of the run, there is no reliable sequential activity in the single home cage trials. **G**, Average reliability across task A trials (solid black line) gradually declines as a function of time. One SD below the mean is shown (gray lines). The single trial shown in **D** is one of the more reliable ones (blue line). **H**, Same as **G**, but for task B. The same network can generate multiple distinct, equally reliable sequences. **I**, Home cage (control) trials are not reliable. **J**, **K**, Average (black) and single trial (blue) trajectories corresponding to task A and task B conditions. **L**, Average sequence reliability over 25 trials for different strengths of input noise to the neurons. The same trace as in **G** is shown (black), as well as analogous reliability traces for fivefold (red) and tenfold (blue) increases in input noise.

A possible network mechanism for sequence generation: the threshold adaptation model

The experimentally observed sequences have several important features that make them particularly suitable for timekeeping. First, they are internally generated; that is, the sequences are not brought about by changing, temporally structured environmental or body-derived inputs. Second, the sequences are reliable from trial to trial, which allows for time inference on single trials. Third, the sequences are context-dependent, with “left” and

“right” trials producing different sequences. Fourth, the sequences are long, on the order of 10–20 s, lasting for the entire duration of the required delay period. To explore how the entorhino-hippocampal circuit might be capable of generating these sequences, we considered a minimalistic model of a recurrent network that we call the threshold adaptation model. This model captures all of the above properties.

We began with a firing rate model (Materials and Methods, Eq. 1). Firing rate models provide fairly accurate descriptions of the dynamics of large recurrent networks when exact spike timing is not important. In our model, the recurrent network structure is a superposition of two components: $J = J^0 + J^{het}$, where J^0 is the “correlated” component and J^{het} is a component uncorrelated across neurons. J^0 has connection strengths that depend only on the distance between neurons arranged on a two-dimensional torus-like grid (Fig. 2A) and follow a Mexican-hat pattern of connectivity reflecting short-range excitation and long range effective inhibition. J^{het} is a random matrix that adds heterogeneity to the pattern of connections. Although the synapses represented by J^0 are chosen to be weak relative to J^{het} (see supplemental Text, available at www.jneurosci.org as supplemental material), they strongly influence network dynamics due to their highly correlated structure. The matrix J^{het} disrupts the perfect symmetry of the Mexican-hat connectivity; this ensures that our results do not depend on the fine-tuned symmetry of J^0 , as this symmetry is unrealistic and may produce misleading results (Zhang, 1996; Seung et al., 2000; Renart et al., 2003).

Mexican-hat connectivity, as in J^0 , and the associated continuous attractor dynamics have been hypothesized as an underlying network mechanism of spatial working memory, spatial navigation and path integration (Samsonovich and McNaughton, 1997; Tsodyks, 1999; Constantinidis and Wang, 2004; McNaughton et al., 2006; Burak and Fiete, 2009). For a wide range of perturbed Mexican-hat connectivity matrices $J = J^0 + J^{het}$, the network activity will quickly converge to a “bump attractor” in the presence of constant input (Seung et al., 2000; Renart et al., 2003). If the input stays approximately constant, the bump will not move. Therefore, the dynamics of Equation 1 alone cannot produce self-generated sequences, since the “bump” only moves in response to significant changes in external inputs.

To overcome this limitation, we added an activity-dependent adaptation of the spike thresholds (Materials and Methods, Eq.

2). Threshold adaptation in hippocampal pyramidal neurons was observed experimentally (Henze and Buzsáki, 2001), and evolves on a relatively slow time scale (~ 1 s). In our model, threshold adaptation has the important effect of destabilizing the “bump” attractor states of the network. In the fast-time scale dynamics (Eq. 1), the system still evolves to a bump attractor, but as the firing rates of the neurons in the bump of activity increase, so do the corresponding thresholds, and this in turn decreases each neuron’s ability to continue firing. The threshold adaptation thus forces the bump to move to a new location, at which point the process repeats itself, resulting in a continuously moving bump that never stabilizes (Fig. 2*A,C*; see also supplemental Movie, available at www.jneurosci.org as supplemental material). The moving bump of activity is what produces sequential firing of the neurons.

Cell assembly sequences generated by the threshold-adaptation model are context-dependent, long-lasting, and reliable

To understand how context dependence of sequences may arise in our model, we considered bump trajectories under different initial conditions, mimicking the left versus right trial types in the behavioral task. Similar to the behavioral experiments, we kept the firing rate initial conditions identical in all simulations, initiating them as a bump of activity in the center of the grid of neurons (Fig. 2*C*, top left). Different initial conditions thus differed only in the initial values for the threshold adaptation variables, as these depend on the recent spiking history of the neurons. Initial conditions A to D represent neurons with adapted thresholds to the left, bottom, right, and top of the activity bump, respectively (Fig. 2*B*; see also supplemental Text, available at www.jneurosci.org as supplemental material). In each case, the activity tends to move away from the neurons with adapted thresholds, as it is more difficult for these to fire. Different initial conditions thus produce different center-of-bump trajectories. The precise contours of these trajectories are determined by the synaptic heterogeneities J^{het} . Importantly, the trajectories for the same initial condition are reliable across trials despite being driven by noisy, temporally and spatially unstructured inputs (Fig. 2*D*). Different instances of the heterogeneous connectivity matrix J^{het} also produce reliable, context-dependent bump trajectories, while the trajectories vary significantly between matrices (supplemental Fig. 3, available at www.jneurosci.org as supplemental material).

To better assess the length and reliability of cell assembly sequences produced by the model, we simulated multiple task trials for each of four initial conditions (A to D). Reliability was quantified using the same reliability measure as in Figure 1, *E* and *F* (see Materials and Methods). Sequential activity on the order of 15–20 s can be seen for different initial conditions (Fig. 3*A,B*), and the sequences are quite different (Fig. 3*C*). As expected from the reliability of bump trajectories (Fig. 2*D*), sequential activity on single trials closely resembled that of the average (Fig. 3*D,E*) and had a high degree of reliability (Fig. 3*G,H*). Average and single-trial bump trajectories also showed reliability in the task conditions (Fig. 3*J,K*). Note that the time scale of the adaptation controls the speed of the moving bump. A shorter time scale would result in a faster-moving trajectory, leading to shorter sequences (data not shown). We also simulated a home cage condition (see Materials and Methods) that resulted in large trial-to-trial variability, which bore little resemblance to the average activity across trials (Fig. 3*F,I*). The lack of sequential structure in the home cage condition resulted from the lack of consistency

Figure 4. Sequential activity is inherited by a second layer via sparse and random connections. *A*, Average sequence over 25 trials of simulated data in a second layer with 1000 neurons. Layer 2 has no recurrent excitation and receives sparse, random feedforward inputs from the torus-like layer simulated in Figure 3. *B*, A single trial, with neuron ordering as in *A*. *C*, The reliability of the sequence as a function of time follows a profile similar to what is seen in the first layer (cf. Fig. 3*G*).

in initial conditions, and the temporally correlated noise (task trials had uncorrelated noise). Finally, we investigated the reliability of sequences as a function of the strength of the input noise to each neuron. As expected, sequences generated in the presence of fivefold and tenfold increases in input noise were less reliable, and the reliability decreased more rapidly as a function of time (Fig. 3*L*). Average and single-trial sequences in these higher noise

conditions are shown in supplemental Figure 6 (available at www.jneurosci.org as supplemental material). In summary, the sequences for all task (but not home cage) initial conditions were long-lasting and reliable, making them suitable for accurate time estimation to tens of seconds.

Although we initially introduced heterogeneities J^{het} to our matrix of synaptic weights to ensure that our results did not depend on the fine-tuned symmetry of the Mexican-hat connectivity matrix J^0 , we found that synaptic heterogeneity had the unexpected benefit of lengthening the sequences (supplemental Fig. 4, available at www.jneurosci.org as supplemental material). This is because the heterogeneities “carved out” irregular bump trajectories, allowing the bump to travel for a longer period of time without repeating itself. To verify that our results did not depend on the particular instance of heterogeneity we chose for the matrix $J^1 = J^0 + J^{het}$, we repeated the analyses in Figure 3 and supplemental Figure 4 (available at www.jneurosci.org as supplemental material) using two more instances of J^{het} to obtain matrices J^2 and J^3 (supplemental Fig. 5, available at www.jneurosci.org as supplemental material).

Cell assembly sequences can be inherited by a downstream layer

We have shown that reliable and context-dependent sequences of neuronal activation similar to what we have observed in CA1 may arise from a recurrent network with torus-like architecture and a weakly correlated pattern of Mexican-hat connectivity. However, the architecture of the CA1 region, with its supersparse recurrent excitation, does not fit with this pattern of connectivity. For this reason, we investigated whether or not reliable sequences generated in one layer can be inherited by a downstream layer. In contrast to the first layer, the second layer we devised had no recurrent excitation and only a global, nonspecific recurrent inhibition. Layer 2 was driven by both the output of the previous layer and noisy, temporally and spatially unstructured inputs. The feedforward connections between the first, torus-like layer and the second layer were random and sparse (see Materials and Methods). Figure 4 shows that despite the lack of structure in layer 2 the sequential activity from the first layer was perfectly inherited by the second layer with a similar reliability profile. The reliability remained unchanged even when the magnitude of the noisy inputs to the second layer was increased 5- or tenfold (see supplemental Fig. 7, available at www.jneurosci.org as supplemental material).

Discussion

Neural correlates of elapsed time on a suprasecond scale have been documented in several cortical regions (Kojima and Goldman-Rakic, 1982; Fuster, 2001; Brody et al., 2003; Janssen and Shadlen, 2005; Lebedev et al., 2008; Mita et al., 2009). Surprisingly, hippocampal circuits have not been considered as timers, despite the critical role of the hippocampus in timing behavior (Clark and Isaacson, 1965; Young and McNaughton, 2000) and the key importance of temporal context in episodic memory (Tulving, 1972) and navigation (McNaughton et al., 1996).

Our experimental observations and modeling results suggest that in the hippocampus the same cell populations that keep information about past memories and planned travel directions of the animal (Pastalkova et al., 2008) also provide information about elapsed time. Elapsed time was reliably inferred from the population firing rate vector of the recorded neurons at any time point of the memory task. Although the time estimation error

from the neuronal population increased over time, it did not increase proportionally to the duration of elapsed time, in contrast to Weber’s law (cf. Staddon, 2005). In our network model (threshold adaptation model), the sequences emerged as a natural byproduct of a network with a perturbed Mexican-hat connectivity pattern and adaptation of the spike thresholds. The threshold adaptation model does not involve learning of sequences, synfire chains, or a “hidden” feedforward network structure (Abeles, 1991; Levy et al., 2005; Ganguli et al., 2008; Liu and Buonomano, 2009; Fiete et al., 2010). In principle, however, any network that allows for self-sustained, sequential activation of neurons can potentially be a substrate for time keeping. The adapting spike threshold mechanism was adopted because of its simplicity and because dependence of the spike threshold on prior spiking activity has been demonstrated experimentally (Henze and Buzsáki, 2001). We emphasize though that other mechanisms such as short-term synaptic depression may play a similar role (Abbott and Regehr, 2004). Although the mechanism for sequence generation in our model relies on connectivity patterns unlikely to be present in CA1, we have shown that sequences generated in one area with this kind of architecture (potentially in the entorhinal cortex or CA3) can be inherited by another area via sparse, random connections.

The model we have described here may be specific to the hippocampal system, where time keeping is needed on the scale of tens of seconds, and is different from the timing mechanisms in sensory and motor systems. Since evolving neuronal assemblies, or sequences, have been observed in other systems (Luczak et al., 2007; Fujisawa et al., 2008; Long and Fee, 2008; Johnson et al., 2010), it is possible that our modeling results apply to them as well. In general, our findings support the view that each neuronal system generates its own timing, providing temporal frames for its operations (Buonomano and Karmarkar, 2002). In summary, we found that a simple network mechanism can generate long-lasting, reliable sequences that may be used for timekeeping in the hippocampus.

References

- Abbott LF, Regehr WG (2004) Synaptic computation. *Nature* 431:796–803.
- Abeles M (1991) *Corticonics: neural circuits of the cerebral cortex*. Cambridge, UK: Cambridge UP.
- Brody CD, Hernández A, Zainos A, Romo R (2003) Timing and neural encoding of somatosensory parametric working memory in macaque prefrontal cortex. *Cereb Cortex* 13:1196–1207.
- Buhusi CV, Meck WH (2005) What makes us tick? Functional and neural mechanisms of interval timing. *Nat Rev Neurosci* 6:755–765.
- Buhusi CV, Perera D, Meck WH (2005) Memory for timing visual and auditory signals in albino and pigmented rats. *J Exp Psychol Anim Behav Process* 31:18–30.
- Buonomano DV, Karmarkar UR (2002) How do we tell time? *Neuroscientist* 8:42–51.
- Burak Y, Fiete IR (2009) Accurate path integration in continuous attractor network models of grid cells. *PLoS Comput Biol* 5:e1000291.
- Chacron MJ, Lindner B, Longtin A (2007) Threshold fatigue and information transfer. *J Comput Neurosci* 23:301–311.
- Church RM (1984) Properties of the internal clock. *Ann NY Acad Sci* 423:566–582.
- Clark CV, Isaacson RL (1965) Effect of bilateral hippocampal ablation on Drl performance. *J Comp Physiol Psychol* 59:137–140.
- Constantinidis C, Wang XJ (2004) A neural circuit basis for spatial working memory. *Neuroscientist* 10:553–565.
- Fiete IR, Senn W, Wang CZ, Hahnloser RH (2010) Spike-time-dependent plasticity and heterosynaptic competition organize networks to produce long scale-free sequences of neural activity. *Neuron* 65:563–576.
- Fujisawa S, Amarasingham A, Harrison MT, Buzsáki G (2008) Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. *Nat Neurosci* 11:823–833.

- Fuster JM (2001) The prefrontal cortex—an update: time is of the essence. *Neuron* 30:319–333.
- Ganguli S, Huh D, Sompolinsky H (2008) Memory traces in dynamical systems. *Proc Natl Acad Sci U S A* 105:18970–18975.
- Gibbon J, Malapani C, Dale CL, Gallistel C (1997) Toward a neurobiology of temporal cognition: advances and challenges. *Curr Opin Neurobiol* 7:170–184.
- Gill PR, Mizumori SJ, Smith DM (2010) Hippocampal episode fields develop with learning. *Hippocampus*. Advance online publication. Retrieved September 1, 2010. doi:10.1002/hipo.20832.
- Henze DA, Buzsáki G (2001) Action potential threshold of hippocampal pyramidal cells in vivo is increased by recent spiking activity. *Neuroscience* 105:121–130.
- Ivry RB, Spencer RM (2004) The neural representation of time. *Curr Opin Neurobiol* 14:225–232.
- Janssen P, Shadlen MN (2005) A representation of the hazard rate of elapsed time in macaque area LIP. *Nat Neurosci* 8:234–241.
- Johnson HA, Goel A, Buonomano DV (2010) Neural dynamics of in vitro cortical networks reflects experienced temporal patterns. *Nat Neurosci* 13:917–919.
- Kojima S, Goldman-Rakic PS (1982) Delay-related activity of prefrontal neurons in rhesus monkeys performing delayed response. *Brain Res* 248:43–49.
- Kraus BJ, Robinson RJ, Hasselmo ME, Eichenbaum H, White JA (2010) Time and distance dependence of rat hippocampal neuron responses. *Soc Neurosci Abstr* 36:100.16.
- Lebedev MA, O'Doherty JE, Nicolelis MA (2008) Decoding of temporal intervals from cortical ensemble activity. *J Neurophysiol* 99:166–186.
- Levy WB, Hocking AB, Wu X (2005) Interpreting hippocampal function as recoding and forecasting. *Neural Netw* 18:1242–1264.
- Lewis PA, Miall RC, Daan S, Kacelnik A (2003) Interval timing in mice does not rely upon the circadian pacemaker. *Neurosci Lett* 348:131–134.
- Liu JK, Buonomano DV (2009) Embedding multiple trajectories in simulated recurrent neural networks in a self-organizing manner. *J Neurosci* 29:13172–13181.
- Long MA, Fee MS (2008) Using temperature to analyse temporal dynamics in the songbird motor pathway. *Nature* 456:189–194.
- Luczak A, Barthó P, Marguet SL, Buzsáki G, Harris KD (2007) Sequential structure of neocortical spontaneous activity in vivo. *Proc Natl Acad Sci U S A* 104:347–352.
- Macdonald C, Lepage K, Eden U, Eichenbaum H (2010) Hippocampal neurons encode the temporal organization of non-spatial event sequences. *Soc Neurosci Abstr* 36:100.15.
- Mauk MD, Buonomano DV (2004) The neural basis of temporal processing. *Annu Rev Neurosci* 27:307–340.
- McNaughton BL, Barnes CA, Gerrard JL, Gothard K, Jung MW, Knierim JJ, Kudrimoti H, Qin Y, Skaggs WE, Suster M, Weaver KL (1996) Deciphering the hippocampal polyglot: the hippocampus as a path integration system. *J Exp Biol* 199:173–185.
- McNaughton BL, Battaglia FP, Jensen O, Moser EI, Moser MB (2006) Path integration and the neural basis of the 'cognitive map'. *Nat Rev Neurosci* 7:663–678.
- Michon JA (1985) The complete time experimenter. In: *Time, mind and behavior* (Michon JA, Jackson JL, eds), pp 21–52. Berlin: Springer.
- Mita A, Mushiaki H, Shima K, Matsuzaka Y, Tanji J (2009) Interval time coding by neurons in the presupplementary and supplemental motor areas. *Nat Neurosci* 12:502–507.
- Pastalkova E, Itskov V, Amarasingham A, Buzsáki G (2008) Internally generated cell assembly sequences in the rat hippocampus. *Science* 321:1322–1327.
- Renart A, Song P, Wang XJ (2003) Robust spatial working memory through homeostatic synaptic scaling in heterogeneous cortical networks. *Neuron* 38:473–485.
- Samsonovich A, McNaughton BL (1997) Path integration and cognitive mapping in a continuous attractor neural network model. *J Neurosci* 17:5900–5920.
- Seung HS, Lee DD, Reis BY, Tank DW (2000) Stability of the memory of eye position in a recurrent network of conductance-based model neurons. *Neuron* 26:259–271.
- Staddon JE (2005) Interval timing: memory, not a clock. *Trends Cogn Sci* 9:312–314.
- Thompson RF, Krupa DJ (1994) Organization of memory traces in the mammalian brain. *Annu Rev Neurosci* 17:519–549.
- Tsodyks M (1999) Attractor neural network models of spatial maps in hippocampus. *Hippocampus* 9:481–489.
- Tulving E (1972) Episodic and semantic memory. In: *Organization of memory* (Tulving E, Donaldson W, eds), pp 381–403. New York: Academic.
- Young B, McNaughton N (2000) Common firing patterns of hippocampal cells in a differential reinforcement of low rates of response schedule. *J Neurosci* 20:7043–7051.
- Zhang K (1996) Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: a theory. *J Neurosci* 16:2112–2126.