Offline memory replay in recurrent neuronal networks emerges from constraints on online dynamics

Aaron D. Milstein^{1,2,†}, Sarah Tran^{1*}, Grace Ng^{1*} and Ivan Soltesz¹

¹ Neurosurgery Dept. and Wu Tsai Neurosciences Institute, Stanford University School of Medicine, Stanford CA

² Dept. of Neuroscience and Cell Biology, Robert Wood Johnson Medical School and Center for Advanced Biotechnology and Medicine, Rutgers University, Piscataway, NJ

* These authors contributed equally to this work.

†Corresponding author. Please send correspondence to Aaron D. Milstein:

milstein@cabm.rutgers.edu

Abstract

During spatial exploration, neural circuits in the hippocampus store memories of sequences of sensory events encountered in the environment. When sensory information is absent during “offline” resting periods, brief neuronal population bursts can “replay” sequences of activity that resemble bouts of sensory experience. These sequences can occur in either forward or reverse order, and can even include spatial trajectories that have not been experienced, but are consistent with the topology of the environment. The neural circuit mechanisms underlying this variable and flexible sequence generation are unknown. Here we demonstrate in a recurrent spiking network model of hippocampal area CA3 that experimental constraints on network dynamics such as spike rate adaptation, population sparsity, stimulus selectivity, and rhythmicity enable additional emergent properties, including variable offline memory replay. In an online stimulus-driven state, we observed the emergence of neuronal sequences that swept from representations of past to future stimuli on the timescale of the theta rhythm. In an offline state driven only by noise, the network generated both forward and reverse neuronal sequences, and recapitulated the experimental observation that offline memory replay events tend to include salient locations like the site of a reward. These results demonstrate that biological constraints on the dynamics of recurrent neural circuits are sufficient to enable memories of sensory events stored in the strengths of synaptic connections to be flexibly read out during rest and sleep, which is thought to be important for memory consolidation and planning of future behavior.

Introduction

In mammals, the hippocampus is a brain region involved in the storage and recall of spatial and episodic memories. As an animal explores a spatial environment, different subpopulations of hippocampal neurons known as “place cells” are selectively activated at different positions in space, resulting in sequences of neuronal spiking that are on the seconds-long timescale of locomotor behavior (O’Keefe and Conway, 1978). The synchronous firing of these sparse neuronal ensembles is coordinated by population-wide oscillations referred to as theta (~4-10 Hz) and gamma (~30-100 Hz)

rhythms (Colgin, 2016). Within each cycle of the theta rhythm (~125 ms), the spiking of active neurons is organized into fast timescale sequences such that neurons selective for just-visited positions spike first, then neurons selective for the current position, and finally neurons selective for the next and future positions spike last in the sequence (Drieu and Zugaro, 2019; Foster and Wilson, 2007). These order-preserving fast timescale “theta sequences” are thought to be involved in planning and learning of event order through associative synaptic plasticity (Jensen et al., 1996; Kay et al., 2020).

When an animal stops running, theta and gamma oscillations decrease, and neuronal circuits in the hippocampus instead emit intermittent synchronous bursts of activity that are associated with high-frequency oscillatory activity detectable in local field potential recordings in hippocampal area CA1. These ~100-200 ms long events are referred to as “sharp wave-ripples” (SWRs) (Colgin, 2016), and they occur both during non-locomotor periods of quiet wakefulness, during reward consumption, and during sleep, when sensory information about the spatial environment is reduced or absent. During SWRs, sparse subsets of neurons are co-activated, with a tendency for neurons that fire sequentially during exploratory behavior to also fire sequentially during SWRs, either in the same order, in reverse, or with a mixture of both directions (Davidson et al., 2009; Pfeiffer, 2020; Stella et al., 2019; Wu and Foster, 2014). The hippocampus can also activate sequences of place cells during SWRs that correspond to possible paths through the environment that have not actually been experienced, suggesting a possible role for offline sequence generation in behavioral planning (Gupta et al., 2010; Igata et al., 2021; Ólafsdóttir et al., 2015, 2018; Wu and Foster, 2014). Manipulations that disrupt neuronal activity during SWRs result in deficits in memory recall (Girardeau et al., 2009), supporting a hypothesis that offline reactivation of neuronal ensembles during SWRs is important for the maintenance and consolidation of long-term memories (Buzsáki, 1989; Joo and Frank, 2018).

Hippocampal SWRs are thought to be generated by the synchronous firing of subpopulations of neurons in the CA2 and CA3 regions of the hippocampus (Csicsvari et al., 2000; Oliva et al., 2016), which are characterized by substantial recurrent feedback connectivity (Duigou et al., 2014; Guzman et al., 2016; Okamoto and Ikegaya, 2019). Recurrent networks have long been appreciated for their ability to generate rich internal dynamics (Amit and Brunel, 1997), including oscillations (Ermentrout, 1992). It has also been shown that associative plasticity at recurrent connections between excitatory neurons can enable robust reconstruction of complete memory representations from incomplete or noisy sensory cues (Amit and Brunel, 1997; Guzman et al., 2016; Hopfield, 1982; Treves and Rolls, 1994). However, this “pattern completion” function of recurrent networks requires that similarly tuned neurons activate each other via strong synaptic connections, resulting in sustained self-activation, rather than sequential activation of neurons that are selective for distinct stimuli (Lisman et al., 2005; Pfeiffer and Foster, 2015; Renno-Costa et al., 2014). Previous work has shown that, in order for recurrent networks to generate sequential activity, some mechanism must be in place to “break the symmetry” and enable spread of activity from one ensemble of cells to another (Tsodyks et al., 1996). During spatial navigation, feedforward sensory inputs carrying information about the changing

environment can provide the momentum necessary for sequence generation. However, during hippocampal SWRs, sensory inputs are reduced, and activity patterns are thought to be primarily internally generated by the recurrent connections within the hippocampus. In this study we use a computational model of hippocampal area CA3 to investigate the synaptic, cellular, and network mechanisms that enable flexible offline generation of memory-related neuronal sequences in the absence of ordered sensory information.

A number of possible mechanisms for sequence generation in recurrent networks have been proposed:

- 1) Winner-take-all network mechanism (Almeida et al., 2007, 2009a; Lisman and Jensen, 2013): Within this framework, the subset of excitatory neurons receiving the most strongly weighted synaptic inputs responds first upon presentation of a stimulus. This active ensemble of cells then recruits feedback inhibition via local interneurons, which in turn prevents other neurons from firing for a brief time window (e.g. the ~15 ms duration of a single gamma cycle). This highlights the important roles that inhibitory neurons play in regulating sparsity (how many cells are co-active), selectivity (which cells are active), and rhythmicity (when cells fire) in recurrent networks (Almeida et al., 2009b; Rennó-Costa et al., 2019; Stark et al., 2014; Stefanelli et al., 2016). However, while oscillatory feedback inhibition provides a network mechanism for parsing neuronal sequences into discrete elements, additional mechanisms are still required to ensure that distinct subsets of excitatory neurons are activated in a particular order across successive cycles of a rhythm (Lisman et al., 2005; Ramirez-Villegas et al., 2018).
- 2) Heterogeneous cellular excitability (Luczak et al., 2007; Stark et al., 2015): If the intrinsic properties of neurons in a network are variable and heterogeneous, when a stimulus is presented, those neurons that are the most excitable will fire early, while neurons with progressively lower excitability will fire later, resulting in sequence generation. This mechanism can explain the offline generation of stereotyped, unidirectional sequences, but cannot account for variable generation of sequences in both forward and reverse directions.
- 3) Asymmetric distributions of synaptic weights (Tsodyks et al., 1996): During learning, if changes in synaptic weights are controlled by a temporally asymmetric learning rule, recurrent connections can become biased such that neurons activated early in a sequence have stronger connections onto neurons activated later in a sequence (Reifenstein et al., 2021). This enables internally generated activity to flow along the direction of the bias in synaptic weights. While this mechanism accounts for offline replay of specific sequences in the same order experienced during learning, it cannot account for reverse replay or the flexible generation of non-experienced sequences (Gupta et al., 2010; Igata et al., 2021; Ólafsdóttir et al., 2015, 2018; Wu and Foster, 2014).
- 4) Cellular or synaptic adaptation: It has also been proposed that short-term adaptation of either neuronal firing rate (Itskov et al., 2011) or synaptic efficacy (Romani and Tsodyks, 2015) can enable neuronal sequence generation in recurrent networks without asymmetric synaptic weights. According to this scheme, recently-activated neurons initially recruit connected partners with high

efficacy, but continued spiking results in either a decrease in firing rate, or a decrease in the probability of neurotransmitter release. This causes connections to fatigue over time, and favors the sequential propagation of activity to more recently activated cells. These mechanisms do allow for the stochastic generation of neuronal sequences in either the forward or reverse direction, though they do not prescribe which or how many neurons will participate in a given replay event.

In this study, we sought to understand how neuronal sequence generation in hippocampal area CA3 depends on the structure and function of the underlying network. To do this, we constructed a computational neuronal network model comprised of recurrently connected excitatory and inhibitory spiking neurons, and tuned it to match experimental constraints on the spiking dynamics of CA3 during spatial navigation, including sparsity, selectivity, rhythmicity, and adaptation. We then analyzed the direction and content of neuronal sequences generated both “online” during simulated navigation, and “offline” during simulated rest. We found that when the network was driven by ordered sensory information in the online state, it generated forward-sweeping “theta sequences” that depended on the structure of recurrent connectivity in the network. In the offline state driven by noise, the network generated heterogeneous memory replay events that moved either in forward, reverse, or mixed directions, and depended on neuronal spike-rate adaptation. Finally, when particular spatial locations were over-represented by the network, as occurs in the hippocampus at sites of reward (Lee et al., 2006; Turi et al., 2019; Zaremba et al., 2017), memory replay events were biased towards trajectories that included those salient positions (Gillespie et al., 2021; Ólafsdóttir et al., 2015; Singer and Frank, 2009).

Results

To investigate how sequential activity in the hippocampus generated “online” during spatial exploration can be recapitulated “offline” in the absence of sensory cues, we constructed a simple spiking neuronal network model of rodent hippocampal area CA3 (Materials and Methods). Neural circuits in the hippocampus and cortex typically comprise a majority of excitatory neurons that project information to downstream circuits, and a minority of primarily locally connected inhibitory interneurons. We included populations of excitatory (1000) and inhibitory (200) neurons in proportion to experimental observations (Pelkey et al., 2017; Tremblay et al., 2016) (Figure 1A). Cell models were single-compartment, integrate-and-fire neurons with saturable, conductance-based excitatory and inhibitory synapses (Carnevale and Hines, 2006; Izhikevich, 2007; Izhikevich and Edelman, 2008). Excitatory neurons were endowed with spike-rate adaptation to support punctuated bursting behavior during theta oscillations (O’Keefe and Recce, 1993; Scharfman, 1993), and inhibitory neurons exhibited fast-spiking dynamics to sustain continuous high frequency firing during gamma oscillations (Csicsvari et al., 2003; Ylinen et al., 1995) (Supplementary Figure S1A, Materials and Methods).

To simulate the sensory experience of locomotion in a spatial environment, we provided both excitatory and inhibitory neurons with external afferent inputs from a

Figure 1. Sparsity, selectivity, and rhythmicity in a recurrent spiking neuronal network model of hippocampal area CA3. (A) Diagram illustrates connectivity of model network. Feedforward (FF) external excitatory inputs contact excitatory (E) and inhibitory (I) neurons. E and I neurons are recurrently connected to other E and I neurons. (B) Simulations of rodent “online exploration” emulated the response of the hippocampus during unidirectional locomotion along a circular linear track that takes 3 seconds to traverse at constant run velocity. (C) Population sparsity (active fraction of neurons) vs. time shown for each cell population. (D) Mean firing rate of active neurons vs. time shown for each cell population. (E) Firing rates vs. time of all neurons in each cell population are shown (average of 5 trials from one example network instance). Cells in each population are sorted by the location of maximum firing. (F) Average stimulus selectivity of each cell population. Trial-averaged activity of each cell was centered around the location of maximum firing, and then averaged across cells. (G) The average activity of each population on a single trial (top row) was bandpass filtered in the theta (middle row) and gamma (bottom row) frequency bands. Colored traces show filtered signals (theta: green, gamma: purple). Traces derived from one example network instance. (H) Power spectrum of average population activity indicates dominant frequency components in the theta and gamma bands (one-sided paired t-tests: theta: E vs. FF, $p=0.00001$; I vs. FF, $p=0.00001$; gamma: E vs. FF, $p<0.00001$; I vs. FF, $p<0.00001$). In (C), (D), (F), and (H), data were first averaged across 5 trials per network instance, and then averaged across 5 independent instances of the network. Mean (solid) $\pm$ SEM (shading) were computed across network instances. p-values reflect Bonferroni correction for multiple comparisons.

population of 1000 excitatory neurons, each of which was selectively activated at distinct but overlapping positions within a simulated circular track that took 3 seconds to traverse (Figures 1A, 1B and 1E) (Materials and Methods). Recurrent connections

within and between excitatory and inhibitory cell populations were also included (Figure 1A), as they are hallmark features of hippocampal area CA3, and have been shown to support rich network dynamics (Renno-Costa et al., 2014; Stark et al., 2014). Specifically, inhibitory feedback connections have been shown to regulate the number of simultaneously active neurons (sparsity) (Stefanelli et al., 2016), and to contribute to the generation of theta and gamma network oscillations (Bezaire et al., 2016; Geisler et al., 2005; Rennó-Costa et al., 2019; Stark et al., 2014; Wang, 2010). Plastic excitatory connections between excitatory neurons have long been implicated in stimulus selectivity and the storage and recall of memories (Almeida et al., 2007; Hopfield, 1982; Lisman and Jensen, 2013). It has been proposed that strong connections between ensembles of co-active neurons could arise through a combination of biased connectivity during brain development (Buzsáki et al., 2021; Dragoi and Tonegawa, 2013; Farooq and Dragoi, 2019; Grosmark and Buzsáki, 2016), and experience-driven synaptic plasticity during learning (Bittner et al., 2015, 2017; Brunel and Trullier, 1998; Káli and Dayan, 2000; Milstein et al., 2020; O'Neill et al., 2008). While here we did not simulate these dynamic processes explicitly, we implemented the structured connectivity that is the end result of these processes by increasing the strengths of synaptic connections between excitatory cells that share overlapping selectivity for spatial positions the environment (Supplementary Figure S1B) (Materials and Methods) (Arkhipov et al., 2018).

Despite the relatively simple architecture of this network model, a wide range of networks with distinct dynamics could be produced by varying a number of parameters, including 1) the probabilities of connections between cell types (Káli and Dayan, 2000), 2) the kinetics and strengths of synaptic connections between cell types (Brunel and Wang, 2003), and 3) the magnitude of the above-mentioned increase in synaptic strengths between neurons with shared selectivity (Brunel, 2016; Dorkenwald et al., 2019). To calibrate the network model to produce dynamics that matched experimentally-derived targets, we performed an iterative stochastic search over these parameters, and optimized the following features of the activity of the model network: 1) population sparsity - the fraction of active neurons of each cell type, 2) the mean firing rates of active neurons of each cell type, 3) stimulus-selective firing of excitatory cells, and 4) the frequency and amplitude of theta and gamma oscillations in the synchronous spiking activity of each cell population (Materials and Methods).

This procedure identified a model with dynamics that met all of the above constraints. Given sparse and selective feedforward inputs during simulated navigation (Figure 1B), the excitatory neurons in the network responded with a fraction of active cells (Figure 1C) and with average firing rates comparable to the those of the feedforward input population (Figure 1D). The majority of inhibitory neurons were activated continuously (Figure 1C) at high firing rates (Figure 1D). While excitatory neurons received random connections from feedforward afferents and from other excitatory neurons with heterogeneous spatial tuning, excitatory cells exhibited a high degree of spatial selectivity (Figures 1E and 1F). This selective increase in firing rate at specific spatial locations within the “place field” of each excitatory neuron was supported by enhanced synaptic connection strengths between excitatory neurons with overlapping tuning (Supplementary Figure S1B). While substantial background

excitation occurred in all cells at all spatial positions, firing outside the place field of each cell was suppressed by sufficiently strong inhibitory input (Bittner et al., 2015; Grienberger et al., 2017). Interestingly, inhibitory neurons also exhibited spatial selectivity, albeit to a weaker degree and with a higher background firing rate (Figures 1E and 1F). This feature of the network dynamics was an emergent property that was not explicitly designed or optimized. While excitatory connections onto inhibitory cells were random and not weighted according to shared selectivity (Supplementary Figure S1B), the total amount of excitatory input arriving onto individual inhibitory cells fluctuated across spatial positions, and predicted a small degree of spatial selectivity (Supplementary Figure S1C). Inhibitory inputs received by inhibitory cells reduced their average activity, effectively enabling fluctuations in excitation above the mean to stand out from the background excitation (Supplementary Figure S1C and S1D). This mechanism of background subtraction by inhibitory synaptic input may explain the partial spatial selectivity previously observed in subpopulations of hippocampal inhibitory neurons (Ego-Stengel and Wilson, 2007; Grienberger et al., 2017; Hangya et al., 2010; Marshall et al., 2002; Wilent and Nitz, 2007).

The tuned network model also exhibited oscillatory synchrony in the firing of the excitatory and inhibitory neuron populations, despite being driven by an asynchronous external input (Figures 1G and 1H). The requirement that the network self-generate rhythmic activity in the theta band constrained recurrent excitatory connections to be relatively strong, as this input provided the only source of rhythmic excitation within the network (Supplementary Figure S1E). Interestingly, as the firing rates of inhibitory cells increased within each cycle of the theta rhythm, their synchrony in the gamma band increased, resulting in an amplitude modulation of gamma paced at the theta frequency (Figure 1G and Supplementary Figure S1F). This “theta-nested gamma” is a well-known feature of oscillations in the hippocampus (Soltesz and Deschenes, 1993; Ylinen et al., 1995), and here emerged from fundamental constraints on dual band rhythmicity without requiring additional mechanisms or tuning.

Position decoding reveals “theta sequences” during simulated navigation

Next, we analyzed neuronal sequence generation within the network during simulated navigation. First, we simulated multiple trials and computed trial-averaged spatial firing rate maps for all neurons in the network (Figure 1E). We then used these rate maps to perform Bayesian decoding of spatial position given the spiking activity of all cells in the network from individual held-out trials not used in constructing the decoding template (Figure 2A) (Davidson et al., 2009; Zhang et al., 1998) (Materials and Methods). For the population of feedforward excitatory inputs, the underlying spatial firing rates were imposed, and the spikes of each cell were generated by sampling from an inhomogeneous Poisson process. Thus, decoding position from the activity of this population served to validate our decoding method, and indeed simulated position could be decoded from the spiking activity of the feedforward input population with very low reconstruction error (Figures 2B and 2C). When we applied this method to the population of excitatory neurons within the network, reconstruction error was increased (Figure 2C). This reflected an increased fraction of temporal bins (20 ms)

Figure 2. Online neuronal sequence generation depends on recurrent excitatory synaptic connectivity. (A) Spike times of all neurons in each cell population on a single trial of simulated “online exploration” are marked. A separate set of 5 trials was used to construct a spatial firing rate template for each neuron (e.g. Figure 1E). Here, cells in each population are sorted by the location of maximum average spatial firing rate. (B) The spatial firing rate templates for all neurons were used to perform Bayesian decoding of spatial position from the single trial spiking data shown in (A). The spatial position of maximum likelihood vs. time is shown for each cell population. (C) Position decoding error is quantified for each cell population as a fraction of the track length (one-sided paired t-tests: E vs. FF, $p < 0.00001$; I vs. FF, $p < 0.00001$). (D) Decoded positions of E and I cell populations swept from past to future positions. The lengths of these theta timescale neuronal sequences are quantified as a fraction of the track length (one-sided paired t-tests: E vs. FF, $p < 0.00001$; I vs. FF, $p < 0.00001$). (E – H) Same as (A – D) for an alternative model network with random synaptic strengths at recurrent excitatory connections between E cells. (F) Decoded position error (one-sided paired t-tests: E vs. FF, $p = 0.00003$; I vs. FF, $p = 0.00002$; two-sided t-tests vs. data in (C): E, $p < 0.00001$; I, $p = 0.00003$). (H) Theta sequence length (one-sided paired t-tests: E vs. FF, $p = 0.00003$; I vs. FF, $p = 0.00005$; two-sided t-tests vs. data in (D): E, $p < 0.00001$; I, $p = 0.00006$). In (C), (D), (F), and (H), data were first averaged across 5 trials per network

instance, and then averaged across 5 independent instances of the network. Mean (solid) $\pm$ SEM (shading) were computed across network instances. p-values reflect Bonferroni correction for multiple comparisons.

where the decoded position was either behind or in advance of the actual position (Figure 2B). However, rather than simply reflecting reconstruction noise or poor spatial selectivity of individual cells (Figure 1F), these divergences from actual position resulted from consistent sequential structure in the spiking activity of cells in the excitatory population (Figure 2A and 2B). Ordered neuronal firing resulted in decoded positions that continuously swept from past positions, through the current actual position, to future positions, and then reset to past positions, on the timescale of the ongoing theta rhythm. These “theta sequences” caused decoded position estimates to oscillate around the actual position (Figure 2B), and their sequence length could be quantified from the mean amplitude of this oscillation (Figure 2D) (Materials and Methods). Interestingly, we found that position could also be accurately decoded from the moderately spatially-tuned activity of inhibitory cells in the network (Figure 2C), and that the spiking activity of the inhibitory population was also organized into theta sequences (Figure 2D).

A number of possible mechanisms have been proposed to account for theta sequence generation *in vivo*, including synaptic, cell-intrinsic, and network-level mechanisms (Chadwick et al., 2015, 2016; Drieu and Zugaro, 2019; Foster and Wilson, 2007; Grienberger et al., 2017; Kang and DeWeese, 2019). That theta sequences in the model emerged in both excitatory and inhibitory neuron populations implicates recurrent interactions within the network (Chadwick et al., 2016). To further investigate, we analyzed neuronal sequence generation in a variant of the model in which the strengths of recurrent connections between excitatory neurons were randomized and no longer depended on shared spatial selectivity between connected pairs of cells (Supplementary Figure S2A). This alternative model could still be tuned to match experimental targets, including sparsity, selectivity, and rhythmicity (Supplementary Figures S2B – S2F). In this case the spatial selectivity of excitatory cells was entirely determined by the synaptic weights of the feedforward afferent inputs (Supplementary Figure S2A), while the recurrent excitatory input supported synchronization in the theta and gamma bands (Supplementary Figure S2F). However, in this model, theta timescale neuronal sequence generation in both excitatory and inhibitory cells was suppressed (Figure 2E – 2H). Decoding of position from spikes on single trials produced lower reconstruction error (Figure 2F), as neuronal population activity more faithfully followed the current spatial position provided by the feedforward inputs (Figure 2G), and was not organized into the sweeps from past to future positions characteristic of theta sequences (Figure 2H). These results indicate that structure in the synaptic strengths of recurrent excitatory connections is required for the generation of fast timescale (~125 ms) neuronal sequences when network activity is driven by behavioral timescale (> 1 s) sequences of sensory inputs, as occurs during spatial exploration.

Emergence of offline memory replay

The above results show that the same network structure that enables population dynamics in CA3 to exhibit sparsity, selectivity, and rhythmicity also supports neuronal sequence generation in the online state when ordered sensory information is present. We next sought to understand how neuronal sequences consistent with the sensory environment are generated offline when sensory inputs are reduced. To mimic the transient (~150 ms) increase in population activity that occurs during a hippocampal SWR, we transiently stimulated the network by randomly choosing sparse subsets of cells from the feedforward input population to emit spikes (Figures 3A and 3B). We then used the same decoding templates as above, constructed from the trial-averaged activity during simulated run, to decode spatial position from spiking activity during these transient offline events (Materials and Methods).

Given that the place field locations of the stimulated neurons in the feedforward input population were heterogeneous and unordered, the spatial positions decoded from their spiking were typically discontinuous across adjacent temporal bins (Figures 3A and 3B). This input pattern evoked spiking in sparse subsets of both the excitatory and inhibitory populations in the network. In contrast with the feedforward population, the activity evoked in excitatory neurons was structured such that neurons with nearby place field locations spiked in adjacent temporal bins, resulting in decoded spatial trajectories that were continuous (Figures 3A and 3B). Inhibitory neuron activity during these events was organized into high-frequency oscillations (Figures 3A and 3B). This procedure was repeated to produce thousands of offline events evoked by stimulation of different random ensembles of inputs (Materials and Methods). Across these events, each position along the track was decoded with equal probability (Figure 3C). For each event, the length and mean velocity of the decoded trajectory was calculated from the differences in decoded positions between adjacent bins (Figures 3D and 3E). A mean velocity of zero corresponds to events with equal steps in the forward and reverse directions, while positive velocities correspond to net forward-moving trajectories, and negative velocities correspond to net backwards-moving trajectories. While the trajectories decoded from the random feedforward input population were comprised of large, discontinuous steps that traced out large path lengths with an average velocity near zero, the excitatory neuron population generated shorter, more continuous sequences that progressed in either the forward or reverse directions (Figures 3D and 3E). These trajectories on average covered ~0.5 the length of the track in the short (~150 ms) duration of the offline event. Compared to the run trajectory, which took 3 s to cover the full track length, this corresponded to a ~10-fold temporal compression (Figure 3E), similar to experimental data (Davidson et al., 2009). Spatial trajectories decoded from the inhibitory neuron population were intermediate in length, but with little forward or reverse momentum, similar to the feedforward inputs. However, the inhibitory cells exhibited high-frequency synchrony (Figures 3B and 3F), similar to experimentally recorded CA3 interneurons during hippocampal SWRs (Csicsvari et al., 2000; Tukker et al., 2013).

These data demonstrate that random, unstructured input can evoke sequential activity in a CA3-like recurrent spiking network, with sequences corresponding to

Figure 3. Forward and reverse offline memory replay depends on recurrent excitatory synaptic connectivity. (A – B) Top: spike times of all neurons in each cell population on a single trial of simulated “offline rest” are marked. Data from 5 trials of simulated “online exploration” was used to construct a spatial firing rate template for each neuron (e.g. Figure 1E). Here, cells in each population are sorted by

the location of maximum average spatial firing rate. Bottom: the spatial firing rate templates for all neurons were used to perform Bayesian decoding of spatial position from the single trial spiking data shown in the top row. The spatial position of maximum likelihood vs. time is shown for each cell population. (A) and (B) correspond to two example trials from one example network instance. (C – F) This procedure was repeated for 1000 trials with each of 5 instances of the network. (C) Histogram of spatial positions decoded from each cell population across all simulated replay events (two-sided K-S tests: E vs. FF, $p=1.0$; I vs. FF, $p=1.0$). (D) Histogram of the length of spatial trajectories decoded from each cell population (two-sided K-S tests: E vs. FF, $p<0.00001$; I vs. FF, $p<0.00001$). (E) Histogram of the mean velocity of spatial trajectories decoded from each cell population (two-sided K-S tests: E vs. FF, $p=0.00141$; I vs. FF, $p=1.0$). (F) Power spectrum of average population activity indicates high frequency components (one-sided paired t-tests: 75 Hz – 300 Hz frequency band: E vs. FF, $p=1.0$; I vs. FF, $p<0.00001$). (G – L) Same as (A – F) for an alternative model network with random synaptic strengths at recurrent excitatory connections between E cells. (G) and (H) correspond to two example trials from one example network instance. (I) Decoded positions (two-sided K-S tests: E vs. FF, $p=1.0$; I vs. FF, $p=1.0$; two-sided K-S tests vs. data in (C): E, $p=1.0$; I, $p=1.0$). (J) Offline sequence length (two-sided K-S tests: E vs. FF, $p<0.00001$; I vs. FF, $p=1.0$; two-sided K-S tests vs. data in (D): E, $p<0.00001$; I, $p<0.00001$). (K) Offline sequence velocity (two-sided K-S tests: E vs. FF, $p=1.0$; I vs. FF, $p=1.0$; two-sided K-S tests vs. data in (E): E, $p=0.04666$; I, $p=1.0$). (L) Offline high-frequency rhythmicity (one-sided paired t-tests: 75 Hz – 300 Hz frequency band: E vs. FF, $p=1.0$; I vs. FF, $p<0.00001$; two-sided t-tests vs. data in (F): E, $p=1.0$; I, $p<0.00001$). In (C – F) and (I – L), data were averaged across 5 independent instances of the network. Mean (solid) $\pm$ SEM (shading) were computed across network instances. p-values reflect Bonferroni correction for multiple comparisons.

forward, reverse, or mixed direction trajectories through an experienced spatial environment. This self-generated memory-related activity implicates information stored in the synaptic weights of the recurrent connections within the network as being important for offline replay of experience. However, in most previous models, sequence generation was unidirectional, and was enabled by an asymmetric bias in the strengths of recurrent connections such that neurons encoding positions early in a sequence formed stronger synapses onto neurons encoding later positions (Malerba and Bazhenov, 2019; Reifenstein et al., 2021; Tsodyks et al., 1996). In contrast, the current network flexibly generates sequences in forward, reverse, or mixed directions, and has symmetric recurrent connections such that synaptic strengths between pairs of excitatory neurons depend only on overlapping selectivity, not on sequence order. Does sequence generation in the present network still depend on recurrent connectivity? To test this, we first verified that including an asymmetric bias in the strengths of excitatory connections produced offline replay events that were biased towards forward sequences (Supplementary Figure 3). We next analyzed the sequence content of offline events generated in the variant of the network model with random recurrent connection weights (Figures 2E – 2H and Supplementary Figures S2A – S2F). Indeed, without any structure in the recurrent connection weights, spatial trajectories decoded from the activity of excitatory neurons was more similar to those of the feedforward inputs, consisting of large, discontinuous steps without forward or reverse momentum (Figures 3G – 3K). Still, this network exhibited high-frequency oscillatory synchrony during these offline events (Figure 3L). These results confirmed that recurrent connectivity is important for offline sequence generation.

If symmetric connectivity based on shared stimulus selectivity enables recurrent networks to generate either forward or backward steps, what “breaks” this symmetry

Figure 4. Neuronal spike rate adaptation supports offline memory replay. (A) Intracellular voltage recordings of three neuronal cell models with distinct spiking dynamics in response to simulated square-shaped intracellular current injections. (B – G) Same as Figures 2A – 2F for an alternative model network in which E cells are regular-spiking cell models without spike rate adaptation. (B) and (C) correspond to two example trials from one example network instance. (D) Decoded positions (two-sided K-S tests: E vs. FF, $p=1.0$; I vs. FF, $p=0.65872$; two-sided K-S tests vs. data in Figure 2C: E, $p=1.0$; I, $p=1.0$). (E) Offline sequence length (two-sided K-S tests: E vs. FF, $p<0.00001$; I vs. FF, $p<0.00001$; two-sided K-S tests vs. data in Figure 2D: E, $p<0.00001$; I, $p=1.0$). (F) Offline sequence velocity (two-sided K-S tests: E vs. FF, $p=1.0$; I vs. FF, $p=0.45269$; two-sided K-S tests vs. data in Figure 2E: E, $p=0.00087$; I, $p=0.46951$). (G) Offline high-frequency rhythmicity (one-sided paired t-tests: 75 Hz – 300 Hz frequency band: E vs. FF, $p=0.00182$; I vs. FF, $p=0.00003$; two-sided t-tests vs. data in Figure 2F: E, $p<0.00001$; I, $p=0.00139$). In (D – G), data were averaged across 5 independent instances of the network. Mean (solid) $\pm$ SEM (shading) were computed across network instances. p-values reflect Bonferroni correction for multiple comparisons.

and generates sequences that make net progress in either the forward or backward direction? Unlike our simulated run condition, sequence direction in the offline

generation in the network was facilitated by our choice of “bursty” excitatory cell model, which exhibited spike-rate adaptation (Figure 4A). As mentioned before, use-dependent decreases in either firing rate or synaptic transmission over time can provide momentum to neuronal sequences by favoring the recruitment of new neurons that have not yet been activated during a network population event (Itskov et al., 2011; Romani and Tsodyks, 2015). To test a possible role for cellular adaptation in sequence generation in our model network, we replaced the “bursting” excitatory cell model with a “regular spiking” model without spike rate adaptation (Figure 4A and Supplementary Figure S4). This cell model does not support the high instantaneous firing rates of the bursting cell model, which compromised the peak firing rates of excitatory cells in the network and their entrainment by higher frequency gamma oscillations during simulated online navigation (Supplementary Figures S4D and S4E). Otherwise, this variant of the network did meet criterion for sparsity, selectivity, and rhythmicity (Supplementary Figure S4A – S4C). However, during simulated offline rest, random feedforward inputs evoked a truncated response from the network (Figures 4B – 4E), with the high frequency rhythmic activity of the inhibitory neurons diminishing before the end of the stimulus period (Figures 4B, 4C and 4G). Spatial trajectories decoded from the activity of excitatory neurons in the network were comprised of large steps that did not progress in either forward or reverse directions, similar to the random feedforward inputs (Figure 4F). Overall, these data show that adaptation of neuronal spiking provides a cellular-level mechanism for flexible and reversible sequence generation in recurrent spiking networks.

Preferential replay of reward location

Thus far, we have simulated network activity during spatial navigation, and identified components of the network that enable offline replay of behavioral sequences stored in memory. However, in these simulations all spatial positions were visited with equal occupancy, and considered to be of equal salience or relevance to the virtual animal. This resulted in all positions being replayed with equal probability offline (Figure 3C), mimicking experimental conditions where all spatial positions contain discriminable sensory cues, and opportunities for reward are provided at random times and positions (Turi et al., 2019; Zaremba et al., 2017). However, it has been shown that when reward is provided at a fixed location that the animal must discover through learning, offline memory replay events become biased towards sequences of place cells that encode positions nearby and including the site of reward (Gillespie et al., 2021; Ólafsdóttir et al., 2018; Pfeiffer, 2020; Singer and Frank, 2009). In parallel with the development of this bias in offline memory replay during learning, it has been shown that the fraction of hippocampal pyramidal cells that selectively fire along the path to reward increases (Lee et al., 2006; Turi et al., 2019; Zaremba et al., 2017). Here we sought to test the hypothesis that this network-level over-representation of reward location is sufficient to bias the content of offline memory replay.

We chose a position along the virtual track to be the fixed location of a simulated reward, and biased the allocation of place field locations such that an increased fraction of excitatory neurons were selectively activated at positions near the

Figure 5. Offline memory replay is biased towards reward positions over-represented by the

network. (A) In this variant of the network, an increased proportion of E cells express place fields with peak locations near the site of a simulated reward. Firing rates vs. time of all neurons in each cell population are shown (average of 5 trials from one example network instance). Cells in each population are sorted by the location of maximum firing. Reward site marked with red dashed line. (B – G) Same as Figures 2A – 2F for an alternative model network with population-level over-representation of reward location in E cells. (B) and (C) correspond to two example trials from one example network instance. Reward site marked with red dashed line. (D) Decoded positions (two-sided K-S tests: E vs. FF, $p < 0.00001$; I vs. FF, $p = 0.13860$; two-sided K-S tests vs. data in Figure 2C: E, $p < 0.00001$; I, $p = 1.0$). (E) Offline sequence length (two-sided K-S tests: E vs. FF, $p < 0.00001$; I vs. FF, $p < 0.00001$; two-sided K-S tests vs. data in Figure 2D: E, $p = 1.0$; I, $p < 0.00001$). (F) Offline sequence velocity (two-sided K-S tests: E vs. FF, $p = 0.02194$; I vs. FF, $p = 1.0$; two-sided K-S tests vs. data in Figure 2E: E, $p = 0.76372$; I, $p = 1.0$). (G) Offline high-frequency rhythmicity (one-sided paired t-tests: 75 Hz – 300 Hz frequency band: E vs. FF, $p = 0.00094$; I vs. FF, $p = 0.00007$; two-sided t-tests vs. data in Figure 2F: E, $p = 0.00011$; I, $p = 0.00003$). In (D – G), data were averaged across 5 independent instances of the network. Mean (solid) $\pm$ SEM

(shading) were computed across network instances. p-values reflect Bonferroni correction for multiple comparisons.

reward (Figure 5A). As before, feedforward and recurrent synaptic connection strengths were increased between neurons with overlapping selectivity (Supplementary Figure S5A). Aside from the enhanced fraction of active excitatory neurons near the reward site (Supplementary Figure S5B), this produced network dynamics during simulated navigation that conformed to experimental constraints for sparsity, selectivity, and rhythmicity (Figure 5A and Supplementary Figures S5B – S5E). During simulated offline rest, the excitatory neurons in the network responded to random feedforward inputs by generating neuronal sequences corresponding to forward, reverse, and mixed direction trajectories through the environment (Figures 5B – 5F), paced by high frequency oscillations in the inhibitory cells (Figure 5G), as before (Figures 3A – 3F). However, now positions near the simulated reward site were replayed in a higher proportion of replay events (Figure 5D). This preferential replay of locations over-represented by the network recapitulated experimental findings and supported our hypothesis that nonuniform place cell allocation and biased memory replay are causally linked.

Discussion

In this study we used a simple recurrent spiking network model of hippocampal area CA3 to investigate the structural and functional requirements for offline replay of spatial memories. We optimized synaptic, cellular, and network parameters of the network to produce population dynamics that match experimentally observed sparsity, selectivity and rhythmicity. We found that networks that fit these constraints exhibit additional emergent properties, including the ability to generate fast timescale memory-related neuronal sequences. During simulated spatial navigation, when ordered sensory information was provided on the seconds-long timescale of locomotion behavior, the network produced neuronal sequences that swept from past to future positions on the faster timescale (~125 ms) of the theta rhythm (“theta sequences”). During simulated offline rest, the network responded to transient noisy activation of random, sparse inputs by generating neuronal sequences that corresponded to forward, reverse, or mixed direction trajectories through the spatial environment.

Both online and offline sequence generation depended on structure in the strengths of excitatory synaptic connections such that pairs of neurons with overlapping spatial tuning were more strongly connected. In the online phase, different sparse subsets of excitatory neurons were activated at different spatial positions due to structure in the strengths of connections from spatially-tuned feedforward afferent inputs. The constraint that recurrent excitation must drive rhythmic synchronization in the theta band resulted in relatively strong recurrent connections. During each cycle of the theta rhythm, when the firing rates of the excitatory neurons were at their maximum, synaptic excitation from recurrent connections exceeded that from the feedforward afferents (Supplementary Figure S1E), favoring activation of cells encoding positions ahead of the current position. This generated forward-sweeping sequences that outpaced the speed of locomotion. However, at the opposite phase of the theta

rhythm, when the firing rates of the excitatory cells reached their minimum, the non-rhythmic feedforward input became greater than recurrent excitation (Supplementary Figure S1E), causing theta sequences to reverse direction and relax back towards the current position encoded by the feedforward inputs.

In the offline phase, the feedforward inputs were not activated in a sequence, so momentum had to be entirely internally generated by the network. In this case, the particular subset of active feedforward inputs initially selected a sparse subset of excitatory neurons to begin to fire, which set a starting position for the replayed trajectory. Slight biases in the feedforward input could then influence whether the active ensemble of excitatory neurons next recruited neurons encoding spatial positions in either the forward or reverse direction. Once activity began moving in one direction, spike-rate adaptation facilitated continued sequence movement along that direction. However, depending on fluctuations in the feedforward inputs, sequences were also generated that included changes in direction. Interestingly, this process is akin to interpolation or smoothing – the recurrent connections within the network served to bridge large, discontinuous jumps in position encoded by the noisy feedforward inputs with smaller, more continuous steps. This produced offline sequences that were consistent with the topology of the spatial environment, but did not necessarily replay exact experienced trajectories. These findings are consistent with a recent report that neuronal sequences activated during hippocampal SWRs *in vivo* resembled Brownian motion, or a random walk through the sensory space, rather than precise replay of experience (Stella et al., 2019). This suggests that, rather than serving mainly to consolidate specific episodic memories of ordered sensory experiences, neuronal sequences during SWRs could also explore possible associations within the environment that had not been fully sampled during experience. Our modeling results showing that increased population representations of goal sites bias the content of offline memory replay also corroborate recent findings that previously rewarded locations are replayed more readily than immediate past or immediate future trajectories (Gillespie et al., 2021). Within this framework, synaptic plasticity during offline replay could modify connection strengths to increase the chance that a new path will be taken that is likely to lead to a desired outcome (Ólafsdóttir et al., 2015).

In summary, our modeling results identified a minimal set of elements sufficient to enable flexible and bidirectional memory replay in neuronal networks: spike rate adaptation, and recurrent connectivity between excitatory and inhibitory neuron populations with strengths and kinetics optimized for rhythmogenesis and sparse and selective stimulus representations. In previous models of neuronal sequence generation, additional network components were proposed to enable unidirectional sequences stored in memory to be reversed during offline recall, including neuromodulation (Gauy et al., 2018), excitability of neuronal dendrites (Gauy et al., 2018; Jahnke et al., 2015), coordinated plasticity at both excitatory and inhibitory synapses (Ramirez-Villegas et al., 2018), and functional specialization of diverse subpopulations of inhibitory interneurons (Cutsuridis and Hasselmo, 2011). While these mechanisms may regulate and enhance memory replay, our results suggest that they are not required. This model also makes some experimentally-testable predictions.

First, it implies that ion channel mutations that disrupt neuronal spike rate adaptation may also degrade neuronal sequence generation and memory consolidation (Peters et al., 2005). Secondly, while the direction and content of offline sequences may be largely controlled by internal dynamics and information stored in the synaptic weights within a recurrent neuronal circuit, the model network still required a small amount of random feedforward afferent input to evoke an offline population burst, suggesting that experimental manipulations of afferent projections to hippocampal area CA3 may alter the frequency or content of memory replay events (Chenani et al., 2019; Sasaki et al., 2018). Recent work has also begun to explore the advantages of generative replay for learning in artificial neural networks (Roscow et al., 2021). In addition to better understanding the biological mechanisms of memory consolidation and flexible planning of behavior, characterizing the minimal mechanisms of memory replay could facilitate the engineering of artificial systems with enhanced cognitive functionality.

Acknowledgements

We are grateful to Kristopher Bouchard at LBNL for sharing large-scale computing resources provided by the National Energy Research Scientific Computing Center, a Department of Energy Office of Science User Facility (DE-AC02-05CH11231). This work was also made possible by computing allotments from NSF (XSEDE Comet, NCSA Blue Waters, and TACC Frontera) and supported by NIH (BRAIN U19 award NS104590).

Materials and Methods

Simulations of a recurrent network of excitatory and inhibitory spiking neurons were executed using the python interface for the NEURON simulation software (Hines et al., 2009). Cell models were single-compartment integrate-and-fire neuronal cell models, as defined by Izhikevich (Izhikevich, 2007), and as implemented for the NEURON simulator by Lytton et al (Lytton et al., 2016). Previously calibrated cell models were replicated from those previous reports without modification – the “intrinsically bursting cell” model was used for excitatory neurons (E) with spike rate adaptation, the “regular spiking pyramidal cell” model was used for excitatory neurons without spike rate adaptation, and the “fast-spiking interneuron” model was used for inhibitory neurons (I) (Izhikevich, 2007; Lytton et al., 2016). Individual spikes in presynaptic neurons activated saturable conductance-based synapses with exponential rise and decay kinetics after a constant delay of 1 ms to emulate axonal conduction time (Carnevale and Hines, 2006). Excitatory synapses had a reversal potential of 0 mV (like AMPA-type glutamate receptors), and inhibitory synapses had a reversal potential of -80 mV (like GABA(A)-type receptors). In addition to the excitatory (E) and inhibitory (I) neuron populations, a population of feedforward afferent inputs (FF) provided a source of external excitatory synaptic drive to the model network.

The baseline weights of excitatory synapses onto E cells were sampled from a log-normal distribution (Almeida et al., 2009b; Buzsaki and Mizuseki, 2014), while the weights of excitatory synapses onto I cells, and all inhibitory synapses were sampled

from a normal distribution (Grienberger et al., 2017). In addition to the random baseline synaptic weights assigned to excitatory synapses onto E cells, input strengths were increased by a variable additive factor that depended on the distance between the place fields of cells with overlapping spatial selectivity (Supplementary Figure S1B). The place field locations of the FF and E populations were assigned by distributing locations throughout the circular simulated track at equal intervals, and randomly assigning them to cells within each population. Random connectivity resulted in each E neuron receiving inputs from many FF and E neurons with heterogeneous selectivity, which produced substantial out-of-field excitation at all positions along the track.

For each of six types of connections between the three cell types (E <- FF, E <- E, E <- I, I <- FF, I <- E, I <- I), a number of parameters were varied and explored during optimization to identify model configurations that produced dynamics comparable to experimental observations. These parameters included: the mean and variance of the synaptic weight distribution for each connection type, the decay time constants of the synaptic conductances, the mean number of synapses made by one presynaptic cell onto one postsynaptic cell for each pair of cell types, and the maximum increase in synaptic weight due to shared selectivity, as mentioned above. Self-connections were not permitted.

Optimization was performed using a population-based iterative multi-objective algorithm. During each of 50 iterations, a population of 600 models with different parameters were simulated for one trial of simulated online run, and for 5 trials of simulated offline rest. Different trials were implemented by using a distinct random number stream to sample unique spike times of the feedforward inputs from an inhomogeneous Poisson process. The following features of the network dynamics were evaluated for each model: average minimum and maximum firing rates of E cells during run, average mean firing rates of I cells during run, average fraction of active E and I cells during run, mean firing rates of E cells during rest, average fraction of active E cells during rest, and finally, features related to the frequency and power of theta and gamma band oscillations in E and I cells during run. These features were compared to target values to obtain a set of multiple objective error values. Models within a population were compared to each other and ranked with a non-dominated sorting procedure (Deb, 2011). Then, a new population of models was generated by making small perturbations to the parameter values of the most highly-ranked models from the previous iteration. This algorithm effectively identified model configurations that satisfied multiple objective criterion. In the below tables, the final optimized parameter values and measured network dynamic features are compared for various model configurations discussed in this study:

Table 1. Model parameter values

Parameter	Bounds	Structured E <- E Model	Random E <- E Model	Model with no spike rate adaptation
E <- FF weight mean	0.1 – 5	1.28	0.29	0.55
E <- FF weight st. dev.	0.1 – 5	1.09	0.21	0.41
E <- E weight mean	0.1 – 5	0.75	0.50	2.45
E <- E weight st. dev.	0.1 – 5	0.52	0.48	1.74
E <- I weight mean	0.1 – 5	0.87	0.70	0.54
E <- I weight st. dev.	0.1 – 5	0.47	0.32	0.16
I <- FF weight mean	0.1 – 5	1.50	1.83	1.60
I <- FF weight st. dev.	0.1 – 5	0.68	0.82	0.22
I <- E weight mean	0.1 – 5	1.85	1.31	0.82
I <- E weight st. dev.	0.1 – 5	0.22	0.28	0.79
I <- I weight mean	0.1 – 5	0.16	1.07	0.23
I <- I weight st. dev.	0.1 – 5	0.12	0.68	0.05
E <- FF and E <- E maximum structured Δ weight	1 - 5	3.63	2.80	4.00
E <- FF and E <- E decay (ms)	2 - 20	3.43	10.76	5.91
E <- I decay (ms)	2 - 30	2.32	3.89	27.87
I <- FF and I <- E decay (ms)	2 - 20	15.04	18.53	18.39
I <- I decay (ms)	2 - 30	8.09	5.65	22.35
E <- FF # synapses / pair	0 - 2	0.28	0.31	0.02
E <- E # synapses / pair	0 - 2	0.25	0.18	0.05
E <- I # synapses / pair	0 - 10	2.32	1.07	0.84
I <- FF # synapses / pair	0 - 2	0.08	0.21	0.04
I <- E # synapses / pair	0 - 2	0.10	0.13	0.22
I <- I # synapses / pair	0 - 10	2.66	1.30	2.69

Table 2. Features of model network dynamics

Feature	Target	Structured E <- E Model	Random E <- E Model	Model with no spike rate adaptation
E peak rate (run)	20. (Hz)	17.20	19.96	7.30
E min rate (run)	0. (Hz)	0.24	0.25	0.39
I mean rate (run)	20. (Hz)	19.58	32.05	6.17
E fraction active (run)	0.6	0.59	0.60	0.60
I fraction active (run)	0.95	1.00	1.00	0.87
E theta power (run)	0.5	0.78	0.37	1.29
I theta power (run)	0.5	0.48	0.14	1.15
E gamma power (run)	0.25	0.53	0.40	0.27
I gamma power (run)	0.25	1.19	1.85	1.27
E theta frequency (run)	7. (Hz)	7.09	7.45	6.73
I theta frequency (run)	7. (Hz)	6.91	7.27	6.73
E gamma frequency (run)	70. (Hz)	71.06	72.93	39.29
I gamma frequency (run)	70. (Hz)	71.06	72.93	69.19
E theta frequency tuning index (run)	>5.	6.40	6.30	150.67
I theta frequency tuning index (run)	>5.	9.83	6.10	150.59
E gamma frequency tuning index (run)	>5.	10.67	10.90	-0.01
I gamma frequency tuning index (run)	>5.	16.16	4.14	3.43
E fraction active (rest)	0.06	0.05	0.03	0.06
E mean rate (rest)	50. (Hz)	50.05	50.00	33.95

In the above table, power in the theta and gamma bands were quantified as follows: average population firing rates were band-pass filtered, and the envelopes of the filtered traces were computed from the Hilbert transformation. Then power was expressed as a ratio of the average envelope amplitude to the average population firing rate. To quantify theta and gamma frequency, bandpass filtered traces were subject to frequency decomposition, and the frequency corresponding to the centroid or center-of-mass of the power spectral density distribution was taken as the dominant frequency within the band. The area of the power spectral density distribution was also used to compute a “frequency tuning index” which quantified how concentrated the power distribution was around the centroid frequency. This metric was akin to a signal-to-noise ratio in the frequency domain instead of the time domain, and was computed as follows:

$$1) \quad FTI = \frac{S-N}{2 \cdot \sigma \cdot w}$$

where S is the average power at frequencies within the center of mass quartile containing the centroid frequency (signal), N is the average power at frequencies in the extreme high and low quartiles outside the center of mass quartile (noise), σ is the standard deviation of the power distribution, and w is the half-width of the power distribution in the frequency domain normalized to the width of the bandpass filter. This metric has values near zero when power is distributed uniformly within the filter band, and values larger than one when power is concentrated around the centroid frequency.

Following parameter optimization, each variant of the network was evaluated by
 715 simulating 5 trials of online run, and 1000 trials of offline rest for each of 5 independent
 network instances. For a given set of model parameters, independent instances of
 each network variant were constructed by using distinct random number streams to
 assign place field locations, input spike times, synaptic connections, and synaptic
 weights for all cells in the network.

Bayesian decoding of spatial position from spike times recorded during a single
 720 trial (Figures 2 – 5 and Supplementary Figure S3) was performed using the procedure
 described in (Davidson et al., 2009). The spatial firing rates of all cells were averaged
 across 5 trials of simulated online run to compute the spiking probabilities of each
 neuron in 1 ms bins. Then, spiking data was taken from either a heldout set of 5 trials
 725 of simulated run (Figure 2), or offline rest trials (Figures 3 – 5 and Supplementary Figure
 S3). The numbers of spikes emitted by each cell in 20 ms bins were used to determine
 a likelihood distribution over spatial positions. The position with maximum likelihood
 was used as the decoded position estimate for each temporal bin. In Figure 2,
 decoded positions of E and I cells oscillate around the actual position. The amplitude
 of this oscillation was computed using the Hilbert transform. Theta sequence length
 730 was quantified as the time-averaged magnitude of this oscillation envelope, multiplied
 by a factor of two since the peak-to-trough distance of an oscillation is double the
 envelope magnitude.

Data and Code Availability

735 All code necessary to reproduce the data and analysis presented in this work are
 available here:

Network simulation and analysis code:

https://github.com/neurosutras/optimize_simple_network

Network optimization code: <https://github.com/neurosutras/nested>

References

Almeida, L. de, Idiart, M., and Lisman, J.E. (2007). Memory retrieval time and memory
 745 capacity of the CA3 network: role of gamma frequency oscillations. *Learn Memory* 14,
 795–806.

Almeida, L. de, Idiart, M., and Lisman, J.E. (2009a). A second function of gamma
 frequency oscillations: an E%-max winner-take-all mechanism selects which cells fire.
J Neurosci 29, 7497–7503.

Almeida, L. de, Idiart, M., and Lisman, J.E. (2009b). The input-output transformation of
 750 the hippocampal granule cells: from grid cells to place fields. *J Neurosci* 29, 7504–
 7512.

Amit, D.J., and Brunel, N. (1997). Dynamics of a recurrent network of spiking neurons
 before and following learning. *Netw Comput Neural Syst* 8, 373–404.

- Arkhipov, A., Gouwens, N.W., Billeh, Y.N., Gratiy, S., Iyer, R., Wei, Z., Xu, Z., Abbasi-
755 Asl, R., Berg, J., Buice, M., et al. (2018). Visual physiology of the layer 4 cortical circuit in silico. *Plos Comput Biol* 14, e1006535.
- Bezaire, M.J., Raikov, I., Burk, K., Vyas, D., and Soltesz, I. (2016). Interneuronal mechanisms of hippocampal theta oscillation in a full-scale model of the rodent CA1 circuit. *Elife* 5, e18566.
- 760 Bittner, K.C., Grienberger, C., Vaidya, S.P., Milstein, A.D., Macklin, J.J., Suh, J., Tonegawa, S., and Magee, J.C. (2015). Conjunctive input processing drives feature selectivity in hippocampal CA1 neurons. *Nat Neurosci* 18, 1133–1142.
- Bittner, K.C., Milstein, A.D., Grienberger, C., Romani, S., and Magee, J.C. (2017). Behavioral time scale synaptic plasticity underlies CA1 place fields. *Science* 357,
765 1033–1036.
- Brunel, N. (2016). Is cortical connectivity optimized for storing information? *Nat Neurosci* 19, 749–755.
- Brunel, N., and Trullier, O. (1998). Plasticity of directional place fields in a model of rodent CA3. *Hippocampus* 8, 651–665.
- 770 Brunel, N., and Wang, X.-J. (2003). What Determines the Frequency of Fast Network Oscillations With Irregular Neural Discharges? I. Synaptic Dynamics and Excitation-Inhibition Balance. *J Neurophysiol* 90, 415–430.
- Buzsáki, G. (1989). Two-stage model of memory trace formation: A role for “noisy” brain states. *Neuroscience* 31, 551–570.
- 775 Buzsaki, G., and Mizuseki, K. (2014). The log-dynamic brain: how skewed distributions affect network operations. *Nat Rev Neurosci* 15, 264–278.
- Buzsáki, G., McKenzie, S., and Davachi, L. (2021). Neurophysiology of Remembering. *Annu Rev Psychol* 73, 1–29.
- Carnevale, N.T., and Hines, M.L. (2006). *The NEURON Book* (Cambridge University Press).
780
- Chadwick, A., Rossum, M.C. van, and Nolan, M.F. (2015). Independent theta phase coding accounts for CA1 population sequences and enables flexible remapping. *Elife* 4, e03542.
- Chadwick, A., Rossum, M.C. van, and Nolan, M.F. (2016). Flexible theta sequence
785 compression mediated via phase precessing interneurons. *Elife* 5, e20349.

- Chenani, A., Sabariego, M., Schlesiger, M.I., Leutgeb, J.K., Leutgeb, S., and Leibold, C. (2019). Hippocampal CA1 replay becomes less prominent but more rigid without inputs from medial entorhinal cortex. *Nat Commun* 10, 1341.
- 790 Colgin, L.L. (2016). Rhythms of the hippocampal network. *Nat Rev Neurosci* 17, 239–249.
- Csicsvari, J., Hirase, H., Mamiya, A., and Buzsaki, G. (2000). Ensemble patterns of hippocampal CA3-CA1 neurons during sharp wave-associated population events. *Neuron* 28, 585–594.
- 795 Csicsvari, J., Jamieson, B., Wise, K.D., and Buzsaki, G. (2003). Mechanisms of gamma oscillations in the hippocampus of the behaving rat. *Neuron* 37, 311–322.
- Cutsuridis, V., and Hasselmo, M. (2011). Spatial Memory Sequence Encoding and Replay During Modeled Theta and Ripple Oscillations. *Cogn Comput* 3, 554–574.
- Davidson, T.J., Kloosterman, F., and Wilson, M.A. (2009). Hippocampal replay of extended experience. *Neuron* 63, 497–507.
- 800 Deb, K. (2011). Multi-objective Evolutionary Optimisation for Product Design and Manufacturing. 3–34.
- Dorkenwald, S., Turner, N.L., Macrina, T., Lee, K., Lu, R., Wu, J., Bodor, A.L., Bleckert, A.A., Brittain, D., Kemnitz, N., et al. (2019). Binary and analog variation of synapses between cortical pyramidal neurons. *Biorxiv* 2019.12.29.890319.
- 805 Dragoi, G., and Tonegawa, S. (2013). Distinct preplay of multiple novel spatial experiences in the rat. *Proc National Acad Sci* 110, 9100–9105.
- Drieu, C., and Zugaro, M. (2019). Hippocampal Sequences During Exploration: Mechanisms and Functions. *Front Cell Neurosci* 13, 232.
- 810 Duigou, C.L., Simonnet, J., Teleńczuk, M.T., Fricker, D., and Miles, R. (2014). Recurrent synapses and circuits in the CA3 region of the hippocampus: an associative network. *Front Cell Neurosci* 7, 262.
- Ego-Stengel, V., and Wilson, M.A. (2007). Spatial selectivity and theta phase precession in CA1 interneurons. *Hippocampus* 17, 161–174.
- 815 Ermentrout, B. (1992). Complex dynamics in winner-take-all neural nets with slow inhibition. *Neural Networks* 5, 415–431.
- Farooq, U., and Dragoi, G. (2019). Emergence of preconfigured and plastic time-compressed sequences in early postnatal development.

- Foster, D.J., and Wilson, M.A. (2007). Hippocampal theta sequences. *Hippocampus* 17, 1093–1099.
- 820 Gauy, M.M., Lengler, J., Einarsson, H., Meier, F., Weissenberger, F., Yanik, M.F., and Steger, A. (2018). A Hippocampal Model for Behavioral Time Acquisition and Fast Bidirectional Replay of Spatio-Temporal Memory Sequences. *Front Neurosci-Switz* 12, 961.
- 825 Geisler, C., Brunel, N., and Wang, X.-J. (2005). Contributions of Intrinsic Membrane Dynamics to Fast Network Oscillations With Irregular Neuronal Discharges. *J Neurophysiol* 94, 4344–4361.
- Gillespie, A.K., Maya, D.A.A., Denovellis, E.L., Liu, D.F., Kastner, D.B., Coulter, M.E., Roumis, D.K., Eden, U.T., and Frank, L.M. (2021). Hippocampal replay reflects specific past experiences rather than a plan for subsequent choice. *Neuron*.
- 830 Girardeau, G., Benchenane, K., Wiener, S.I., Buzsáki, G., and Zugaro, M.B. (2009). Selective suppression of hippocampal ripples impairs spatial memory. *Nat Neurosci* 12, 1222–1223.
- 835 Grienberger, C., Milstein, A.D., Bittner, K.C., Romani, S., and Magee, J.C. (2017). Inhibitory suppression of heterogeneously tuned excitation enhances spatial coding in CA1 place cells. *Nat Neurosci* 20, 417–426.
- Grosmark, A.D., and Buzsáki, G. (2016). Diversity in neural firing dynamics supports both rigid and learned hippocampal sequences. *Science* 351, 1440–1443.
- Gupta, A.S., Meer, M.A.A. van der, Touretzky, D.S., and Redish, A.D. (2010). Hippocampal Replay Is Not a Simple Function of Experience. *Neuron* 65, 695–705.
- 840 Guzman, S.J., Schlogl, A., Frotscher, M., and Jonas, P. (2016). Synaptic mechanisms of pattern completion in the hippocampal CA3 network. *Science* 353, 1117–1123.
- Hangya, B., Li, Y., Muller, R.U., and Czurkó, A. (2010). Complementary spatial firing in place cell–interneuron pairs. *J Physiology* 588, 4165–4175.
- 845 Hines, M.L., Davison, A.P., and Muller, E. (2009). NEURON and Python. *Front Neuroinform* 3, 1.
- Hopfield, J.J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proc National Acad Sci* 79, 2554–2558.
- Igata, H., Ikegaya, Y., and Sasaki, T. (2021). Prioritized experience replays on a hippocampal predictive map for learning. *Proc National Acad Sci* 118, e2011266118.

- 850 Itskov, V., Curto, C., Pastalkova, E., and Buzsáki, G. (2011). Cell Assembly Sequences Arising from Spike Threshold Adaptation Keep Track of Time in the Hippocampus. *J Neurosci* 31, 2828–2834.
- Izhikevich, E.M. (2007). *Dynamical Systems in Neuroscience* (MIT Press).
- 855 Izhikevich, E.M., and Edelman, G.M. (2008). Large-scale model of mammalian thalamocortical systems. *Proc National Acad Sci* 105, 3593–3598.
- Jahnke, S., Timme, M., and Memmesheimer, R.M. (2015). A Unified Dynamic Model for Learning, Replay, and Sharp-Wave/Ripples. *J Neurosci* 35, 16236–16258.
- 860 Jensen, O., Idiart, M.A., and Lisman, J.E. (1996). Physiologically realistic formation of autoassociative memory in networks with theta/gamma oscillations: role of fast NMDA channels.
- Joo, H.R., and Frank, L.M. (2018). The hippocampal sharp wave–ripple in memory retrieval for immediate use and consolidation. *Nat Rev Neurosci* 19, 744–757.
- Káli, S., and Dayan, P. (2000). The Involvement of Recurrent Connections in Area CA3 in Establishing the Properties of Place Fields: a Model. *J Neurosci* 20, 7463–7477.
- 865 Kang, L., and DeWeese, M.R. (2019). Replay as wavefronts and theta sequences as bump oscillations in a grid cell attractor network. *Elife* 8, e46351.
- Kay, K., Chung, J.E., Sosa, M., Schor, J.S., Karlsson, M.P., Larkin, M.C., Liu, D.F., and Frank, L.M. (2020). Constant Sub-second Cycling between Representations of Possible Futures in the Hippocampus. *Cell* 180, 552–567.e25.
- 870 Lee, I., Griffin, A.L., Zilli, E.A., Eichenbaum, H., and Hasselmo, M.E. (2006). Gradual Translocation of Spatial Correlates of Neuronal Firing in the Hippocampus toward Prospective Reward Locations. *Neuron* 51, 639–650.
- Lisman, J.E., and Jensen, O. (2013). The theta-gamma neural code. *Neuron* 77, 1002–1016.
- 875 Lisman, J.E., Talamini, L.M., and Raffone, A. (2005). Recall of memory sequences by interaction of the dentate and CA3: A revised model of the phase precession. *Neural Networks* 18, 1191–1201.
- 880 Luczak, A., Barthó, P., Marguet, S.L., Buzsáki, G., and Harris, K.D. (2007). Sequential structure of neocortical spontaneous activity in vivo. *Proc National Acad Sci* 104, 347–352.

- Lytton, W.W., Seidenstein, A.H., Dura-Bernal, S., McDougal, R.A., Schürmann, F., and Hines, M.L. (2016). Simulation Neurotechnologies for Advancing Brain Research: Parallelizing Large Networks in NEURON. *Neural Comput* 28, 2063–2090.
- 885 Malerba, P., and Bazhenov, M. (2019). Circuit mechanisms of hippocampal reactivation during sleep. *Neurobiol Learn Mem* 160, 98–107.
- Marshall, L., Henze, D.A., Hirase, H., Leinekugel, X., Dragoi, G., and Buzsáki, G. (2002). Hippocampal Pyramidal Cell–Interneuron Spike Transmission Is Frequency Dependent and Responsible for Place Modulation of Interneuron Discharge. *J Neurosci* 22, RC197–RC197.
- 890 Milstein, A.D., Li, Y., Bittner, K.C., Grienberger, C., Soltesz, I., Magee, J.C., and Romani, S. (2020). Bidirectional synaptic plasticity rapidly modifies hippocampal representations independent of correlated activity. *Biorxiv* 2020.02.04.934182.
- Okamoto, K., and Ikegaya, Y. (2019). Recurrent connections between CA2 pyramidal cells. *Hippocampus* 29, 305–312.
- 895 O’Keefe, J., and Conway, D.H. (1978). Hippocampal place units in the freely moving rat: why they fire where they fire. *Exp Brain Res* 31, 573–590.
- O’Keefe, J., and Recce, M.L. (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* 3, 317–330.
- 900 Ólafsdóttir, H.F., Barry, C., Saleem, A.B., Hassabis, D., and Spiers, H.J. (2015). Hippocampal place cells construct reward related sequences through unexplored space. *Elife* 4, e06063.
- Ólafsdóttir, H.F., Bush, D., and Barry, C. (2018). The Role of Hippocampal Replay in Memory and Planning. *Curr Biol* 28, R37–R50.
- 905 Oliva, A., Fernandez-Ruiz, A., Buzsáki, G., and Berenyi, A. (2016). Role of Hippocampal CA2 Region in Triggering Sharp-Wave Ripples. *Neuron* 91, 1342–1355.
- O’Neill, J., Senior, T.J., Allen, K., Huxter, J.R., and Csicsvari, J. (2008). Reactivation of experience-dependent cell assembly patterns in the hippocampus. *Nat Neurosci* 11, 209–215.
- 910 Pelkey, K.A., Chittajallu, R., Craig, M.T., Tricoire, L., Wester, J.C., and McBain, C.J. (2017). Hippocampal GABAergic Inhibitory Interneurons. *Physiol Rev* 97, 1619–1747.
- Peters, H.C., Hu, H., Pongs, O., Storm, J.F., and Isbrandt, D. (2005). Conditional transgenic suppression of M channels in mouse brain reveals functions in neuronal excitability, resonance and behavior. *Nat Neurosci* 8, 51–60.

- Pfeiffer, B.E. (2020). The content of hippocampal “replay.” *Hippocampus* 30, 6–18.
- 915 Pfeiffer, B.E., and Foster, D.J. (2015). Autoassociative dynamics in the generation of sequences of hippocampal place cells. *Science* 349, 180–183.
- Ramirez-Villegas, J.F., Willeke, K.F., Logothetis, N.K., and Besserve, M. (2018). Dissecting the Synapse- and Frequency-Dependent Network Mechanisms of In Vivo Hippocampal Sharp Wave-Ripples. *Neuron* 100, 1224–1240 e13.
- 920 Reifenshtein, E.T., Khalid, I.B., and Kempter, R. (2021). Synaptic learning rules for sequence learning. *Elife* 10, e67171.
- Renno-Costa, C., Lisman, J.E., and Verschure, P.F. (2014). A signature of attractor dynamics in the CA3 region of the hippocampus. *Plos Comput Biol* 10, e1003641.
- 925 Rennó-Costa, C., Teixeira, D.G., and Soltesz, I. (2019). Regulation of gamma-frequency oscillation by feedforward inhibition: A computational modeling study. *Hippocampus* 29, 957–970.
- Romani, S., and Tsodyks, M. (2015). Short-term plasticity based network model of place cells dynamics. *Hippocampus* 25, 94–105.
- 930 Roscow, E.L., Chua, R., Costa, R.P., Jones, M.W., and Lepora, N. (2021). Learning offline: memory replay in biological and artificial reinforcement learning. *Trends Neurosci* 44, 808–821.
- Sasaki, T., Piatti, V.C., Hwaun, E., Ahmadi, S., Lisman, J.E., Leutgeb, S., and Leutgeb, J.K. (2018). DENTATE NETWORK ACTIVITY IS NECESSARY FOR SPATIAL WORKING MEMORY BY SUPPORTING CA3 SHARP-WAVE RIPPLE GENERATION AND
935 PROSPECTIVE FIRING OF CA3 NEURONS. *Nat Neurosci* 21, 258–269.
- Scharfman, H.E. (1993). Spiny neurons of area CA3c in rat hippocampal slices have similar electrophysiological characteristics and synaptic responses despite morphological variation. *Hippocampus* 3, 9–28.
- 940 Singer, A.C., and Frank, L.M. (2009). Rewarded Outcomes Enhance Reactivation of Experience in the Hippocampus. *Neuron* 64, 910–921.
- Soltesz, I., and Deschenes, M. (1993). Low- and high-frequency membrane potential oscillations during theta activity in CA1 and CA3 pyramidal neurons of the rat hippocampus under ketamine-xylazine anesthesia. *J Neurophysiol* 70, 97–116.
- 945 Stark, E., Roux, L., Eichler, R., Senzai, Y., Royer, S., and Buzsaki, G. (2014). Pyramidal cell-interneuron interactions underlie hippocampal ripple oscillations. *Neuron* 83, 467–480.

- Stark, E., Roux, L., Eichler, R., and Buzsaki, G. (2015). Local generation of multineuronal spike sequences in the hippocampal CA1 region. *Proc National Acad Sci* 112, 10521–10526.
- 950 Stefanelli, T., Bertollini, C., Luscher, C., Muller, D., and Mendez, P. (2016). Hippocampal Somatostatin Interneurons Control the Size of Neuronal Memory Ensembles. *Neuron* 89, 1074–1085.
- Stella, F., Baracskey, P., O'Neill, J., and Csicsvari, J. (2019). Hippocampal Reactivation of Random Trajectories Resembling Brownian Diffusion. *Neuron* 102, 450–461 e7.
- 955 Tremblay, R., Lee, S., and Rudy, B. (2016). GABAergic Interneurons in the Neocortex: From Cellular Properties to Circuits. *Neuron* 91, 260–292.
- Treves, A., and Rolls, E.T. (1994). Computational analysis of the role of the hippocampus in memory. *Hippocampus* 4, 374–391.
- 960 Tsodyks, M.V., Skaggs, W.E., Sejnowski, T.J., and McNaughton, B.L. (1996). Population dynamics and theta rhythm phase precession of hippocampal place cell firing: A spiking neuron model. *Hippocampus* 6, 271–280.
- Tukker, J.J., Lasztoczi, B., Katona, L., Roberts, J.D., Pissadaki, E.K., Dalezios, Y., Marton, L., Zhang, L., Klausberger, T., and Somogyi, P. (2013). Distinct dendritic arborization and in vivo firing patterns of parvalbumin-expressing basket cells in the hippocampal area CA3. *J Neurosci* 33, 6809–6825.
- 965 Turi, G.F., Li, W.-K., Chavlis, S., Pandi, I., O'Carroll, J., Priestley, J.B., Grosmark, A.D., Liao, Z., Ladow, M., Zhang, J.F., et al. (2019). Vasoactive Intestinal Polypeptide-Expressing Interneurons in the Hippocampus Support Goal-Oriented Spatial Learning. *Neuron* 101, 1150–1165.e8.
- 970 Wang, X.-J. (2010). Neurophysiological and Computational Principles of Cortical Rhythms in Cognition. *Physiol Rev* 90, 1195–1268.
- Wilent, W.B., and Nitz, D.A. (2007). Discrete Place Fields of Hippocampal Formation Interneurons. *J Neurophysiol* 97, 4152–4161.
- 975 Wu, X., and Foster, D.J. (2014). Hippocampal Replay Captures the Unique Topological Structure of a Novel Environment. *J Neurosci* 34, 6459–6469.
- Ylinen, A., Soltesz, I., Bragin, A., Penttonen, M., Sik, A., and Buzsaki, G. (1995). Intracellular correlates of hippocampal theta rhythm in identified pyramidal cells, granule cells, and basket cells. *Hippocampus* 5, 78–90.

- 980 Zaremba, J.D., Diamantopoulou, A., Danielson, N.B., Grosmark, A.D., Kaifosh, P.W.,
Bowler, J.C., Liao, Z., Sparks, F.T., Gogos, J.A., and Losonczy, A. (2017). Impaired
hippocampal place cell dynamics in a mouse model of the 22q11.2 deletion. *Nat*
Neurosci 20, 1612–1623.
- 985 Zhang, K., Ginzburg, I., McNaughton, B.L., and Sejnowski, T.J. (1998). Interpreting
Neuronal Population Activity by Reconstruction: Unified Framework With Application to
Hippocampal Place Cells. *J Neurophysiol* 79, 1017–1044.

Supplementary Figures

Supplementary Figure S1. Related to Figures 1, 2 and 3. Structure and online dynamics of network model with excitatory synaptic connectivity structured by shared stimulus selectivity. (A)

Intracellular voltage recordings from example E and I cells during simulated “online” spatial exploration.

(B) The strengths of synaptic connections within the model network shown in Figure 1 are indicated by grayscale intensity. Cells in each population are sorted by the location of maximum average spatial firing

rate. Excitatory connections from FF and E cells onto E cells are increased in strength for pairs of cells with overlapping spatial selectivity. **(C)** Spatial modulation for each cell is computed as a ratio of maximum to mean activity. The degree of spatial modulation expected from a linear weighted sum of excitatory inputs is larger for E cells than I cells due to the structure of the weight distributions shown in (B). However, the actual spatial modulation measured from the firing rate outputs of each cell is larger than expected in both E and I cells due to suppression of background activity by synaptic inhibition (one-sided paired t-tests: Expected vs. Actual: E, $p < 0.00001$; I, $p < 0.00001$). Mean $\pm$ SEM shown were computed across cells for one example instance of the network. Statistics were computed across 5 independent instances of the network. **(D)** The distance between the expected location of maximum firing, and the actual location is quantified as a fraction of the circular track. Mean $\pm$ SEM shown were computed across cells for one example instance of the network. **(E)** Traces depict average population firing rates of E and FF cells during an example trial for one instance of the network. E cells and FF cells dominate at different phases of the population theta oscillation. **(F)** In Figure 1G (bottom row), the amplitudes of the gamma-filtered (purple) population firing rates for E and I populations vary in time. Here, the amplitudes or envelopes of the gamma-filtered population firing rates were subject to frequency decomposition. The resulting frequency distributions show peaks in the theta band (one-sided paired t-tests: 4 Hz – 10 Hz frequency band: E vs. FF, $p = 0.00002$; I vs. FF, $p < 0.00002$). Mean (solid) $\pm$ SEM (shading) were computed across network instances. Note that the clipped shading in (F) corresponds to small values that appear expanded by the log scaling. p-values reflect Bonferroni correction for multiple comparisons.

Supplementary Figure S2. Related to Figures 2 and 3. Structure and online dynamics of network model with random recurrent excitatory connectivity. (A) Same as Supplementary Figure S1B for an alternative model network with random synaptic strengths at recurrent excitatory connections between E cells. (B – F) Same as Figures 1E, 1C, 1D, 1F, and 1H for alternative model network. (F) Rhythmicity (one-sided paired t-tests: theta: E vs. FF, $p=0.00004$; I vs. FF, $p<0.00001$; gamma: E vs. FF, $p<0.00001$; I vs. FF, $p<0.00001$; two-sided t-tests vs. data in Figure 1H: theta: E, $p=0.00001$; I, $p<0.00001$; gamma: E, $p<0.00001$; I, $p<0.00001$). In (C – F), data were first averaged across 5 trials per network instance, and then averaged across 5 independent instances of the network. Mean (solid) $\pm$ SEM (shading) were computed across network instances. p-values reflect Bonferroni correction for multiple comparisons.

Supplementary Figure S3. Related to Figure 3. Offline replay is unidirectional in

network model with past-skewed synaptic connectivity. (A) Same as Supplementary Figure S1B for an alternative model network with excitatory synaptic connections onto E cells biased such that neurons encoding the current position of the animal during run preferentially activate neurons encoding future positions. This biases offline replay towards forward-moving unidirectional sequences. (B – G) Same as Figures 3A – 3F for an alternative model network with temporally asymmetric synaptic strengths at excitatory connections onto E cells. (B) and (C) correspond to two example trials from one example network instance. (D) Decoded positions (two-sided K-S tests: E vs. FF, $p=1.0$; I vs. FF, $p=1.0$; two-sided K-S tests vs. data in Figure 2C: E, $p=1.0$; I, $p=1.0$). (E) Offline sequence length (two-sided K-S tests: E vs. FF, $p<0.00001$; I vs. FF, $p<0.00001$; two-sided K-S tests vs. data in Figure 2D: E, $p<0.00001$; I, $p=1.0$). (F) Offline sequence velocity (two-sided K-S tests: E vs. FF, $p<0.00001$; I vs. FF, $p=0.00012$; two-sided K-S tests vs. data in Figure 2E: E, $p<0.00001$; I, $p=0.00049$). (G) Offline high-frequency rhythmicity (one-sided paired t-tests: 75 Hz – 300 Hz frequency band: E vs. FF, $p=0.00153$; I vs. FF, $p<0.00001$; two-sided t-tests vs. data in Figure 2F: E, $p=0.00004$; I, $p=0.16726$). In (D – G), data were averaged across 5 independent instances of the network. Mean (solid) $\pm$ SEM (shading) were computed across network instances. p-values reflect Bonferroni correction for multiple comparisons.

Supplementary Figure S4. Related to Figure 4. Online dynamics in network model without spike rate adaptation. (A – E) Same as Figures 1E, 1C, 1D, 1F, and 1H for alternative model network without spike rate adaptation in E cells. **(E)** Rhythmicity (one-sided paired t-tests: theta: E vs. FF, $p=0.00028$; I vs. FF, $p=0.00011$; gamma: E vs. FF, $p=1.0$; I vs. FF, $p=0.00004$; two-sided t-tests vs. data in Figure 1H: theta: E, $p=0.24889$; I, $p=0.00001$; gamma: E, $p<0.00001$; I, $p<0.00001$). In (B – E), data were first averaged across 5 trials per network instance, and then averaged across 5 independent instances of the network. Mean (solid) $\pm$ SEM (shading) were computed across network instances. p-values reflect Bonferroni correction for multiple comparisons.

Supplementary Figure S5. Related to Figure 5. Online activity in network model with population over-representation of reward. (A) Same as Supplementary Figure S1B for an alternative network with population-level over-representation of reward location in E cells. Reward site marked with red dashed line. (B – E) Same as Figures 1C, 1D, 1F, and 1H for alternative model network. (E) Rhythmicity (one-sided paired t-tests: theta: E vs. FF, $p < 0.00001$; I vs. FF, $p < 0.00001$; gamma: E vs. FF, $p = 0.00005$; I vs. FF, $p < 0.00001$; two-sided t-tests vs. data in Figure 1H: theta: E, $p = 0.03168$; I, $p = 1.0$; gamma: E, $p = 0.00322$; I, $p = 0.13110$). In (B – E), data were first averaged across 5 trials per network instance, and then averaged across 5 independent instances of the network. Mean (solid) $\pm$ SEM (shading) were computed across network instances. p-values reflect Bonferroni correction for multiple comparisons.