Fast and Flexible Sequence Induction In Spiking Neural Networks Via Rapid Excitability Changes

Rich Pang[1,2,3] and Adrienne L. Fairhall[2,3]

1. Neuroscience Graduate Program, University of Washington
2. Physiology and Biophysics Department, University of Washington
3. Computational Neuroscience Center, University of Washington

Abstract

Cognitive flexibility likely depends on modulation of the dynamics underlying how biological neural networks process information. While dynamics can be reshaped by gradually modifying connectivity, less is known about mechanisms operating on faster timescales. A compelling entrypoint to this problem is the observation that exploratory behaviors can rapidly cause selective hippocampal sequences to "replay" during rest. Using a spiking network model, we asked whether simplified replay could arise from three biological components: fixed recurrent connectivity; stochastic "gating" inputs; and rapid gating input scaling via long-term potentiation of intrinsic excitability (LTP-IE). Indeed, these enabled both forward and reverse replay of recent sensorimotor-evoked sequences, despite unchanged recurrent weights. LTP-IE "tags" specific neurons with increased spiking probability under gating input, and ordering is reconstructed from recurrent connectivity. We further show how LTP-IE can implement temporary stimulus-response mappings. This elucidates a novel combination of mechanisms that might play a role in rapid cognitive flexibility.

Introduction

We can rapidly and flexibly adapt how we process incoming information from the environment, a mental faculty known as cognitive flexibility. For example, after being instructed to raise our left hand when one word is heard and our right hand when another is heard, we perform the task with little error. How brain networks quickly induce novel stimulus-response mappings such as this into their underlying neural dynamics, however, remains mysterious. In particular, while extensive prior work has elucidated how stimulus-response mappings might be implemented biologically, the mechanisms for inducing these mappings typically require slow, gradual modifications to network structure, e.g. by incrementally training connection weights to minimize errors between correct and predicted responses (Williams and Zipser, 1989; Sussillo and Abbott, 2009; Laje and Buonomano, 2013; Rajan, Harvey, and Tank, 2016; Nicola and Clopath, 2017) or by allowing local plasticity to reshape network dynamics in response to internal activity (Song and Abbott, 2001; Fiete, et al., 2010; Gilson, Burkitt, and Van Hemmen, 2010; Lee and Buonomano, 2012; Klampfl and Maass, 2013; Rezende, Wierstra, and Gerstner, 2014; Diehl and Cook, 2015). Biologically observed spike-timing-dependent plasticity (STDP) mechanisms, however, typically increase post-synaptic potentials (PSPs) by at most a few percent (Markram, et al., 1997; Bi and Poo, 1998; Sjostrom, Turrigiano, and Nelson, 2001; Wang, et al., 2005; Caporale and Dan, 2008). Furthermore, due to the precise timing requirements of canonical plasticity mechanisms (e.g. in STDP spike pairs must occur in the correct order within tens of milliseconds) and low firing rates of cortex and hippocampus (typically less than a few tens of Hz (Griffith and Horn, 1966; Koulakov, Hromádka, and Zador, 2009; Mizuseki and Buzsáki, 2013), STDP-triggering spike patterns may occur relatively rarely, especially given the asynchronous nature of cortical firing patterns in awake animals (Renart, et al., 2010). As a result, computational models for shaping dynamics that modify network structures via synaptic plasticity typically rely on at least dozens of learning trials over extended time periods (Song and Abbott, 2001; Masquelier, Guyonneau, and Thorpe, 2008; Klampfl and Maass 2013), challenging their suitability for rapid reconfiguration of network dynamics. Computationally, modifying connections to change network function might also interfere

with long-term memories or computations already stored in the network's existing connectivity patterns, leading, for instance, to catastrophic forgetting (McCloskey and Cohen, 1989). Consequently, it is unclear (1) how biologically observed plasticity mechanisms could reshape network function over the timescales of seconds required for rapid cognitive flexibility, and (2) how such restructuring of synaptic connectivity could occur over the short term without degrading existing long-term memories.

One feature of neural dynamics potentially reflecting processes of cognitive flexibility is stereotyped sequential firing patterns (Hahnloser, Kozhevnikov, and Fee, 2002; Ikegaya, et al., 2004; Luczak, et al., 2007; Pastalkova, et al., 2008; Davidson, Kloosterman, and Wilson, 2009; Crowe, Sprekeler, and Kempter, 2010; Harvey, et al., 2012). Functionally, firing sequences are thought to be involved in various cognitive processes, from short-term memory (Davidson, Kloosterman, and Wilson, 2009; Crowe, Sprekeler, and Kempter, 2010) to decision-making (Harvey, et al., 2012). More generally, one can imagine sequential activity as reflecting information propagating from one subnetwork to another, e.g. stimulus S evoking a cascade of activity that eventually triggers motor output M. A compelling empirical example of memory-related firing sequences that arise in a neural network almost immediately after a sensorimotor event, apparently without requiring repeated experience or long-term learning, occurs in awake hippocampal "replay". Here, sequences of spikes in hippocampal regions CA1 and CA3 originally evoked by a rodent traversing its environment along a specific trajectory subsequently replay when the rodent pauses to rest (Foster and Wilson, 2006; Davidson, Kloosterman, and Wilson, 2009; Gupta, et al., 2010; Carr, Jadhav, and Frank, 2011). Such replay events occur at compressed timescales relative to the original trajectory and often in reverse order. Replay has also been observed in primate cortical area V4, where firing sequences evoked by a short movie were immediately reactivated by a cue indicating the movie was about to start again, but without showing the movie (Eagleman and Dragoi, 2012). The functional role of replay has been implicated in memory, planning, and learning (Ego-Stengel and Wilson, 2009; Carr, Jadhav, and Frank, 2011; Eagleman and Dragoi, 2012; Jadhav, et al., 2012; Ólafsdóttir, Bush, and Barry, 2018), but little is known about the mechanisms enabling the underlying sequential activity to be induced in the network dynamics in the first place. Elucidating biological mechanisms for rapidly inducing sequential firing patterns in network dynamics may not only illuminate the processes enabling replay but may also shed light on principles for fast and flexible reconfiguration of computations and information flow in neural networks more generally.

An intriguing fast-acting cellular mechanism whose role in shaping network dynamics has not been investigated is the rapid, outsize, and activity-dependent modulation of cortical inputs onto pyramidal cells (PCs) in hippocampal region CA3 (Hyun, et al., 2013; Hyun, et al., 2015; Rebola, Carta, and Mulle, 2017). Specifically, following a 1-2 s train of 20 action potentials in a given CA3 PC, excitatory postsynaptic potentials (EPSPs) from medial entorhinal cortex (MEC) more than doubled in magnitude within eight seconds (the first time point in the experiment). Thought to arise through inactivation of K^+-channels colocalized with MEC projections onto CA3 PC dendrites and deemed "long-term potentiation of intrinsic excitability" (LTP-IE), potentiation occurred regardless of whether the CA3 PC spikes were evoked via current injection or by upstream physiological inputs, indicating its heterosynaptic nature, since only MEC EPSPs, and not others, exhibited potentiation (Hyun, et al., 2013; Hyun, et al., 2015). Notably, the 10-20 Hz spike rate required for potentiation matches the range of *in vivo* spike rates in hippocampal "place cells" when rodents pass through specific locations (Moser, Kropff, and Moser, 2008; Mizuseki, et al., 2012), suggesting LTP-IE may occur in natural contexts. Furthermore, EPSP potentiation persisted throughout the multi-minute course of the experiment (Hyun, et al., 2015), suggesting that in addition to fast onset, the modulation could extend significantly into the future. Although inducing sequential activation patterns in neural networks is typically associated with *homosynaptic* plasticity (e.g. in STDP a postsynaptic following a presynaptic spike strengthens the activated synapse, thereby incrementally increasing the probability of subsequent presynaptic spikes triggering postsynaptic spikes), the rapidity, strength, and duration of this heterosynaptic potentiation

mechanism suggest it might significantly modulate network dynamics in natural conditions, warranting further investigation within the context of neural sequences. Moreover, while this mechanism has so far been observed in hippocampus only, an intriguing possibility is that functionally similar mechanisms exist in cortex also, enabling rapid effective changes in excitability of recently active cells, with potentially similar computational consequences (see Discussion).

To explore this idea we first develop a computational model for the effect of LTP-IE on EPSPs from upstream inputs as a function of their spiking responses to physiological inputs. Next, we demonstrate how LTP-IE combined with recurrent PC connectivity in a spiking network can yield spike sequences reflecting recent sensorimotor sequences that replay in both forward and reverse and which are gated by an upstream gating signal. We subsequently identify parameter regimes allowing and prohibiting LTP-IE-based sequence propagation and examine the effect of specific parameters on replay frequency and propagation speed. We next show how LTP-IE-based sequences can be used to induce temporary stimulus-response mappings in an otherwise untrained recurrent network. Finally, using a reduced model we give proof-of-concept demonstrations of how LTP-IE might support more general computations. We discuss implications for cognitive flexibility and rapid memory storage that does not require modification of recurrent network weights.

Results

To investigate its consequences on neuronal spiking dynamics we implemented LTP-IE in a network of leaky integrate-and-fire (LIF) neurons (Gerstner, et al., 2014). Neurons received three types of inputs: sensorimotor inputs S, carrying tuned “external” information; recurrent input R from other cells in the network, comprising excitation and inhibition (developed further, **Fig 2**); and “gating” inputs G, assumed to be random, with a homogeneous rate across the network but independent to each cell. We provide an overview and intuition of the model before elaborating our results.

First, we demonstrate our implementation of LTP-IE in spiking neurons. We apply LTP-IE to a neuron, potentiating its G inputs, if the neuron fires within a physiological range (~10-20 Hz in our model) for 1 second in response to input S. Thus, recently active cells get “tagged” by LTP-IE, so that they exhibit augmented EPSPs in response to G inputs. In the presence of inputs G, recently active cells will therefore show larger positive membrane voltage deflections, on average, because their EPSPs from G are larger, and they will sit closer to spiking threshold. Unless otherwise specified, in all simulations that follow we assume the continuous presence of inputs G, as LTP-IE would not affect membrane voltages without them.

Second, in analogy to hippocampal “place cells” (Moser, Kropff, and Moser, 2008; Mizuseki, et al., 2012) we consider the case where different excitatory pyramidal cells (PCs) maximally fire (up 20 Hz) when a simulated animal is in different positions in the environment. When the simulated animal moves along a trajectory, LTP-IE thus tags the cells with place fields along the trajectory, storing the trajectory memory as a set of LTP-IE-tagged neurons with specific position tunings; order and timing information, however, are not stored by LTP-IE. When G inputs are turned on post-navigation (representing an awake resting state), cells activated by the original trajectory thus sit closer to spiking threshold than cells not activated by the trajectory.

Third, inspired by excitatory recurrence in CA3 (Lisman, 1999) we assume recurrent excitatory connections in the network, which allows activity to propagate through the network. We assume that cell pairs with more similar tuning (i.e. with nearby place fields) have stronger connections (Lisman, 1999; Káli and Dayan, 2000; Giusti, et al., 2015). Thus, activity from cell A is more likely to propagate to cell B if (1) cell B has been LTP-IE-tagged after being activated by the trajectory, and (2) cell B has a nearby place field to cell A. Consequently, given the right model parameters, activity that begins at one point in the network (e.g. in cells

tuned to the animal's resting position) should propagate among LTP-IE-tagged cells, in an order reflective of the original trajectory.

Finally, we include a population of inhibitory cells (INH) one-tenth the size of the excitatory PC population and randomly recurrently connected with them. Inhibition limits the number of active PCs and prevents activity from propagating in more than one direction along the LTP-IE-defined trajectory. We discuss additional model components such as LTP-IE extinction and higher-order structure in G inputs in the Discussion.

Note that our model accords two meanings of "excitability" to the LTP-IE acronym. As coined by (Hyun, et al., 2013; Hyun, et al., 2015), LTP-IE's "excitability" originally refers to the augmented EPSPs of G inputs (MEC inputs, in [Hyun, et al., 2013; Hyun, et al., 2015]) arriving to distal dendrites. Here, however, "excitability" also refers to the increased membrane potential of LTP-IE-tagged cells (in the presence of G inputs), which makes them more likely to spike in response to recurrent inputs.

LTP-IE increases membrane voltages and spike rates under random gating inputs

We first simulate the expected effects of LTP-IE on neuronal membrane voltages and spike rates. To model LTP-IE, when a cell spikes sufficiently fast (~10-20 Hz) over 1 second, we scale that cell's G inputs by a factor $\sigma(r)$ that depends on the rate r at which it fired during that 1 second. For such cells, future G inputs will thus yield augmented EPSPs (***Fig 1A, B***). Consequently, cells that have recently emitted several spikes in quick succession end up with high LTP-IE levels, which causes them to sit at higher average voltages given a steady random stream of G inputs (***Fig 1C-E***); additionally, cells with higher LTP-IE exhibit increased variability in their membrane voltages (***Fig 1D, E***). (Note: we chose to model the LTP-IE activation function [***Fig 1B***] as a sigmoid based on the saturation-like relationship between PC spike count and the membrane conductance change measured in Hyun, et al., 2015 [Fig 3H], but evidence also exists for an optimal firing rate between 10 and 20 Hz [Hyun 2013, *J Physiol*]; modeling LTP-IE activation with such an optimum should not affect our results, however, as our goal was simply to apply LTP-IE to PCs firing at physiological rates.)

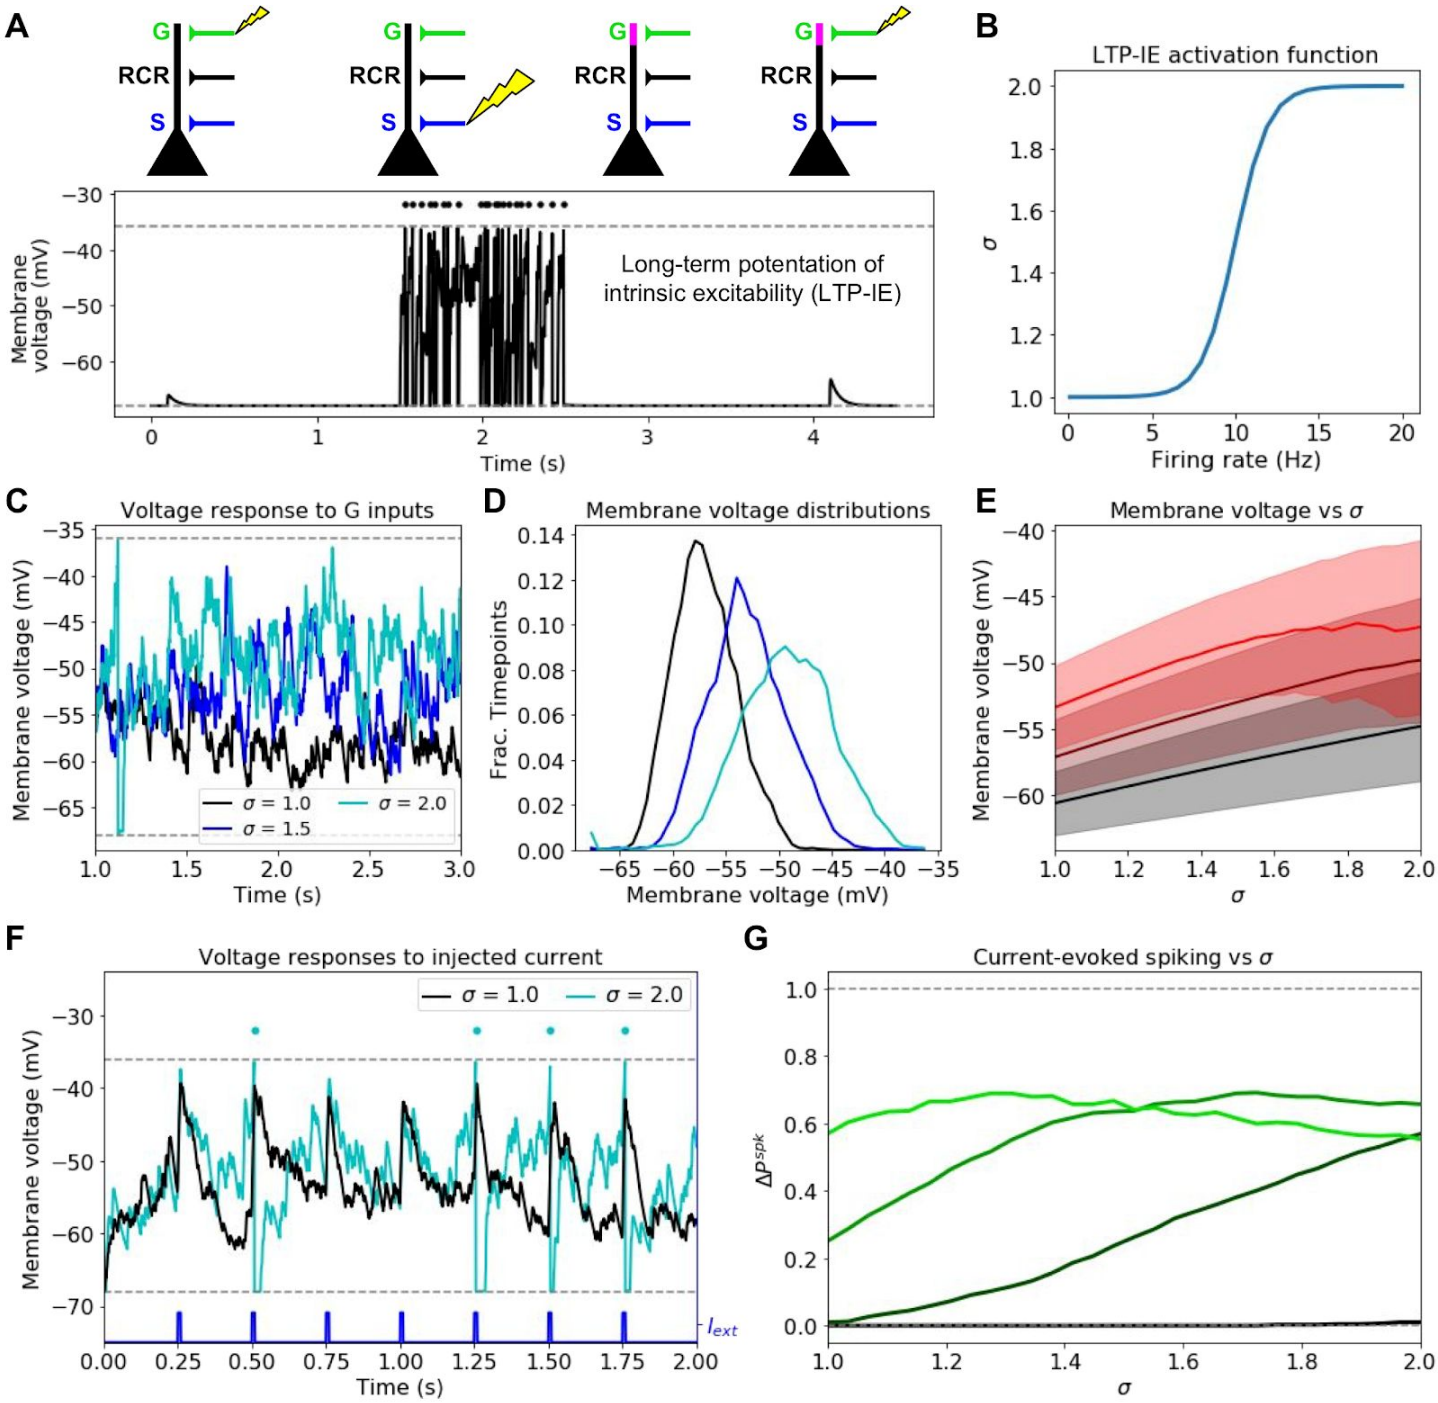

Figure 1: Mechanism and consequences of LTP-IE*. A. Demonstration using leaky integrate-and-fire (LIF) neuron model of fast activity-driven LTP-IE, which doubles membrane voltage responses to gate inputs as described in* (Hyun, et al., 2013; Hyun, et al., 2015)*. A spike in an upstream gate neuron (G) first elicits a small EPSP in the pyramidal cell (PC); a 1-second spike train (dots) at approximately 20 Hz is evoked by strong stimulation of sensory inputs S; when G spikes again, the EPSP has doubled in size. "RCR" refers to recurrent inputs from other PCs (not used in this figure). Dashed lines show leak and spike-threshold voltages. B. Shifted logistic function for LTP-IE strength (effective weight scaling factor) σ vs. PC firing rate over 1 second. C. Example membrane voltages of PCs with different σ receiving stochastic but statistically identical gating input spikes. D. Distribution of PC membrane voltage for σ values shown in C. E. Mean (thick) and standard deviation (shading) of V_m as a function of σ for three gate firing rates. Black: r_G = 75 Hz; dark red: r_G = 125 Hz; red: r_G = 175 Hz. F. Example differential sensitivities of PC spike responses to injected current input (blue) for two different σ. Dashed lines show leak and spike threshold potential; dots indicate spikes (which only occur*

for the σ = 2 case [cyan]). G. Difference between current-evoked spike probability and spontaneous spike probability as a function of σ for four initial gate input weights. Color code, in order of increasing lightness: $w^{PG,G}$ *= .4, .8, 1.2, 1.6 (see methods for units).*

Due to their increased membrane voltage, cells with high LTP-IE also exhibit higher spiking probability, both spontaneously (***Figure 1-figure supplement 1***) and in response to depolarizing input currents (***Figure 1F,G***), although even with high LTP-IE, current-evoked spiking is not assured since substantial variability remains (***Figure 1F,G***). Nonetheless, LTP-IE can increase a model neuron's probability of transforming an input current into an output spike from near zero to approximately 0.5, when the PC→G synaptic weight $w^{PC,G}$ takes a moderate value (***Figure 1G***). (For overly weak $w^{PC,G}$, LTP-IE does not facilitate spiking since the membrane voltage remains subthreshold; for overly strong $w^{PC,G}$, spontaneous spike rates are already high, and increased spiking leads to more frequent voltage resetting [***Figure 1G, Figure 1-figure supplement 1***] so total spike rate does not substantially increase.) Thus, given moderate $w^{PC,G}$, when a cell undergoes LTP-IE its chance of spiking in response to future inputs (i.e. its excitability) increases, although variability remains. Since LTP-IE-triggering spike rates are only around 10-20 Hz, this suggests LTP-IE may play an active role in modulating firing properties of recently active cells in physiological conditions.

Spike sequences propagate along LTP-IE-defined paths through a network

We next asked how LTP-IE as described in ***Figure 1*** would shape activity when introduced into a recurrently connected model network of 3000 excitatory spiking pyramidal cells (PCs). Intuitively, since LTP-IE increases spiking probability within recently active neurons, we expect recently active neurons to be preferentially recruited during sequence propagation through the network. Inspired by hippocampal place cells (O'Keefe, 1979; Moser, Kropff, and Moser, 2008), we allowed excitatory PCs in our network to be tuned to specific positions ("place fields") within the environment (***Figure 2A-C***). We furthermore assumed stronger connectivity between PCs with nearby place fields (***Figure 2D*** and Methods), such that propagation would be more likely to reflect movement along continuous "virtual" trajectories through the environment, and in particular trajectories containing LTP-IE-tagged cells. Finally, the excitatory PC network was connected to a population of 300 inhibitory (INH) cells, which served to control the number of simultaneously active PCs. PC→INH and INH→PC connection probability was 0.5.

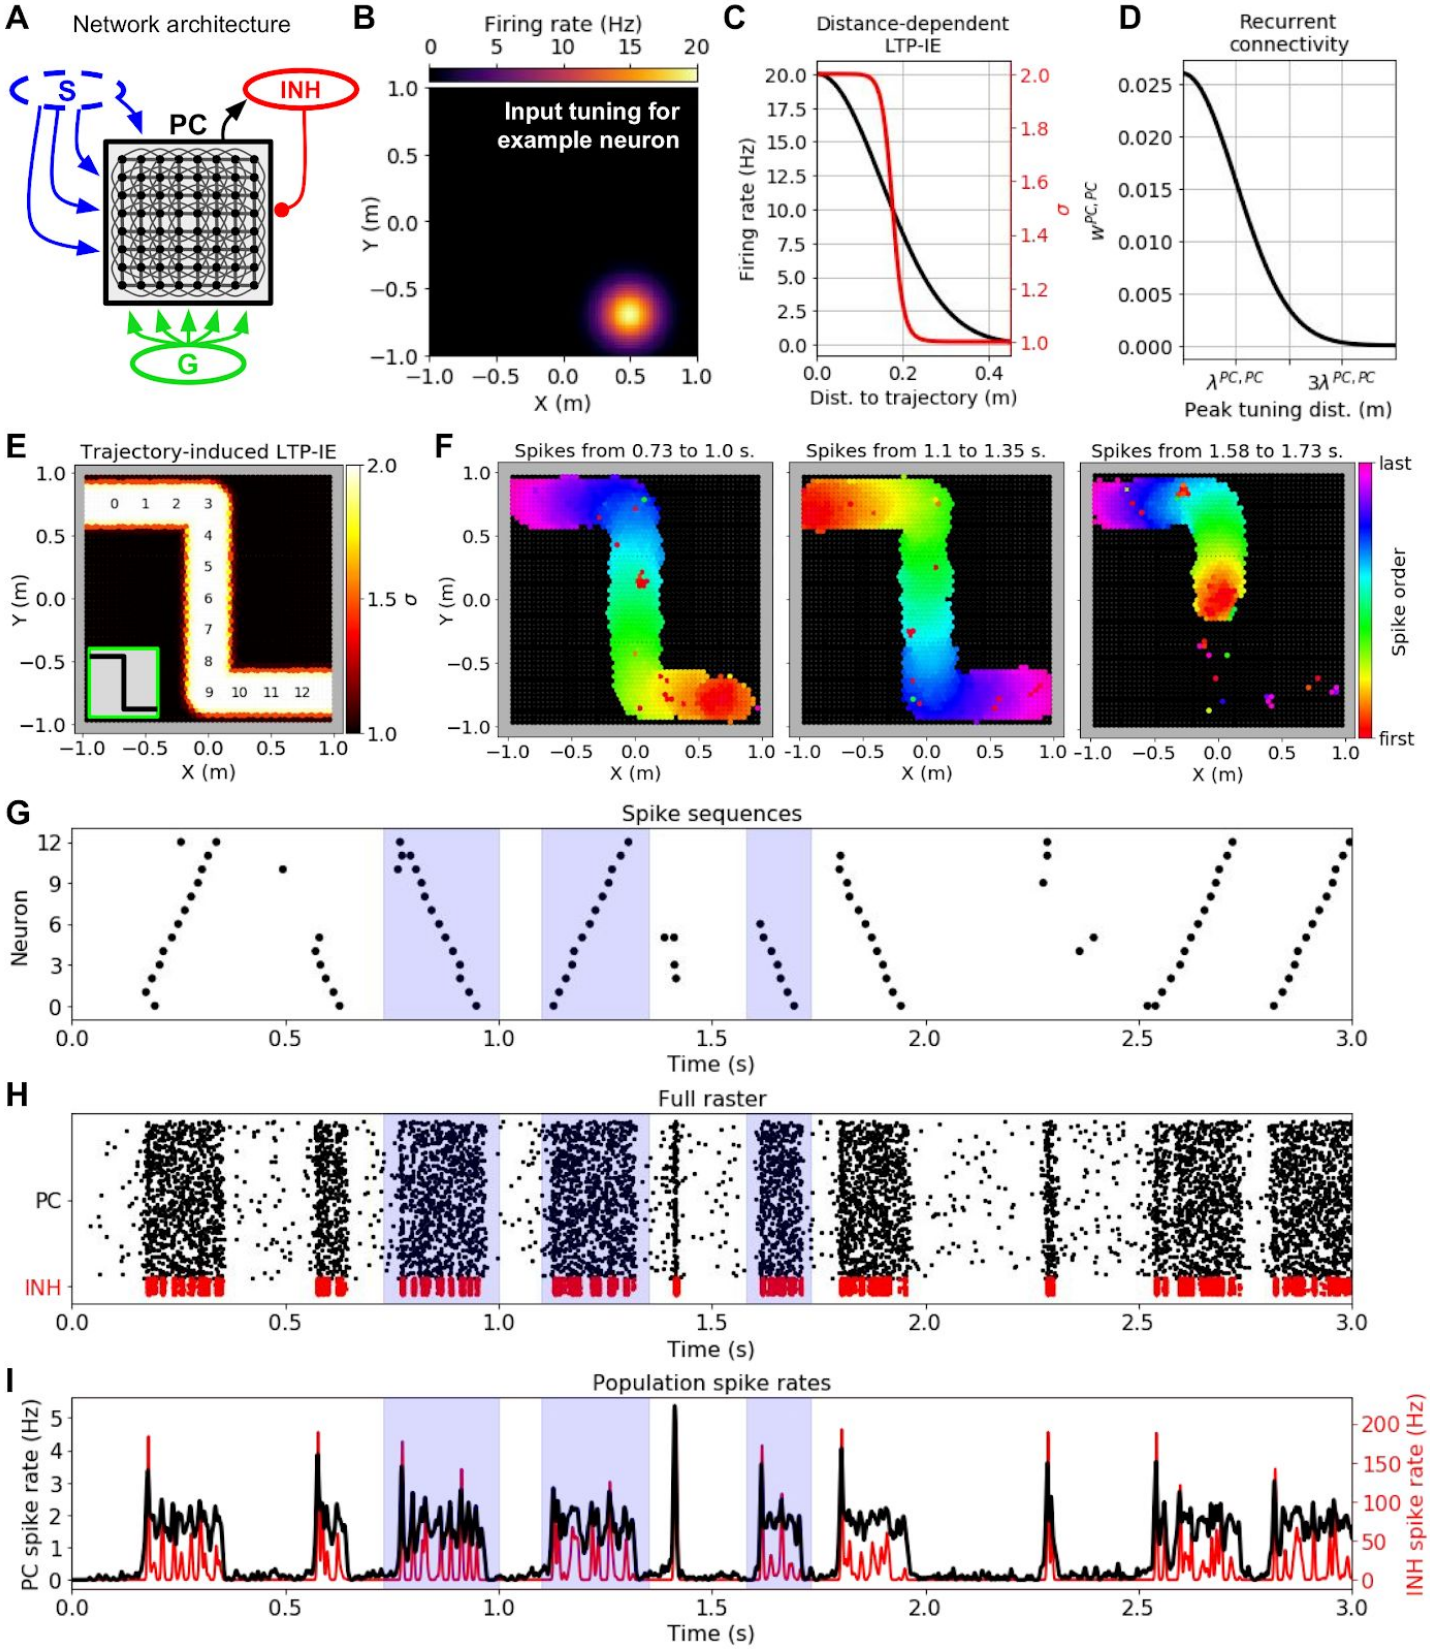

***Figure 2: Demonstration of history-dependent sequence replay via LTP-IE**. A. Network architecture used in simulations. The dashed oval around input S indicates that we did not explicitly simulate its interaction with the pyramidal cell (PC) network, but rather only the LTP-IE profile one would expect following the termination of S input. B. Squared exponential position tuning for an example neuron (20 Hz max. firing rate, 0.15 m length scale). C. Resultant LTP-IE as a function of distance between a cell's maximum tuning and the closest point on the trajectory through the environment (red), computed as a sigmoidal function (**Figure 1B**) of*

position-dependent firing rate (black). D. Recurrent excitatory weights between PCs as squared exponential function of distance between the two cells' peak tuning positions. E. LTP-IE profile induced in PCs by an example Z-shaped trajectory (inset) in a network of 3000 PC and 300 inhibitory (INH) cells. PCs are positioned according to peak tuning and colored by the LTP-IE level (σ) expected to result from the trajectory. The numbers indicate the position tunings of the PC identifiers in G. (Note: individual cells cannot be seen here due to their high density). F. Cells activated during the different spontaneous replay events shaded in blue in G-I, colored by the order of the first spikes each cell emitted during the event. Black cells did not activate during replay. Left and middle: replay in two different directions; right: partial replay. G. Partial raster showing spike times for cells with position tunings marked in E. H. Full raster plot for PC and INH population over a 3-second trial. I. Cell-averaged spike rates for PC and INH populations throughout trial.

To investigate network activity shaped by recent sensorimotor sequences, we considered the following scenario, in analogy with typical experimental setups used to measure neuronal replay in hippocampus (Foster and Wilson, 2006; Davidson, Kloosterman, and Wilson, 2009; Carr, Jadhav, and Frank, 2011): a rodent has just run along a short trajectory through its environment and now rests for several seconds in an awake, quiescent state. How does the neural activity evoked during the trajectory shape the spontaneous activity in the awake, quiescent state? Since our goal was to understand sequential reactivation during this awake, quiescent period, we only used the trajectory geometry to predict the LTP-IE levels one would expect following the trajectory's termination (i.e. the termination of S inputs). To do so, we computed each neuron's maximum expected spike rate during the trajectory as a function of the distance from its place field to the nearest point on the trajectory (***Figure 2C***). We then passed the result through a soft-thresholding LTP-IE activation function (***Figure 1B***) to compute the final expected LTP-IE level σ (***Figure 2C***) for each PC. This allowed us to model the expected LTP-IE profile over the network of neurons as a function of the recent trajectory (***Figure 2E***). As per our design, the LTP-IE profile stores which locations were covered by the original trajectory but bears no explicit information about its speed or direction.

Can replay-like sequences that recapitulate the original trajectory structure emerge from the LTP-IE profile stored in the network? When we let our network run in the presence of random G inputs independent to each PC, replay-like events lasting on the order of 100-200 ms spontaneously arose in the network (***Figure 2F-I***). Spontaneous replay events propagated in both directions (***Figure 2F, left and middle***), recapitulating the forward and reverse replay observed experimentally (Foster and Wilson, 2006; Davidson, Kloosterman, and Wilson, 2009; Carr, Jadhav, and Frank, 2011). Due to the PC refractory period of ~8 ms in our model (see *Methods*), activity propagated in one direction without reversing. Partial replay also occurred, in which replay began in the middle of the trajectory and propagated in one direction to the end (***Figure 2F, right***). When an alternative trajectory was used to induce the LTP-IE profile, triggered replay recapitulated that trajectory instead (***Figure 2-figure supplement 1A-E***), indicating that replay of a specific trajectory was not due *a priori* to the recurrent connectivity; indeed, this should be avoided by the spatial uniformity of the recurrence. The virtual trajectory encoded by the replay event could be decoded by computing the median place field of the PCs spiking across short time windows during the event (***Figure 2-figure supplement 1F***). We additionally note that the replay sequence was able to turn corners, suggesting robustness to nonlinear trajectories. Finally, plotting population firing rates (***Figure 2I***) reveals a potential connection between the replay events in our model and "sharp-wave ripples", rapid oscillatory field potentials experimentally observed to co-occur with hippocampal sequence replay (Carr, Jadhav, and Frank, 2011). Thus, despite sizable variability in membrane voltages arising from the random gating inputs G (***Figure 1C,D,F***), LTP-IE was able to induce spontaneously arising activity sequences into the network that recapitulated recent sensorimotor experiences.

Dependence of LTP-IE-based sequence propagation on network parameters

What conditions must hold for LTP-IE to induce successful sequences in the network? To address this we explored how the frequency of spontaneous replay events changed as we varied different network parameters. To ensure that we only analyzed replay events reflecting recent sensorimotor experiences, we only included events in which average PC activity was sufficiently elevated and in which PC spikes were confined primarily to LTP-IE-tagged PCs.

We observed that spontaneous replay events among LTP-IE-tagged cells arose across a sizeable range of network parameters (***Figure 3A-C***). We first found that replay events occurred with a frequency greater than 1 Hz across a large range of excitatory and inhibitory feedback strengths, $w^{PC,PC}$ and $w^{PC,INH}$, respectively (***Figure 3A***); when $w^{PC,PC}$ was too large without sufficient inhibitory compensation, however, the network entered a "blowup" regime in which activity spread across the whole network instead of being confined to the LTP-IE-tagged PCs (***Figure 3A, black***). A further role for inhibition is demonstrated in ***Figure 3D-E***. In particular, while replay confined to LTP-IE-tagged PCs can occur in the absence of inhibitory feedback (***Figure 3D-right***), in this case one observes an increase in spontaneous "bidirectional" events, in which replay activity begins in the middle of the trajectory and propagates outward in both directions (compare ***Figure 3D-left*** to ***Figure 2F-right***).

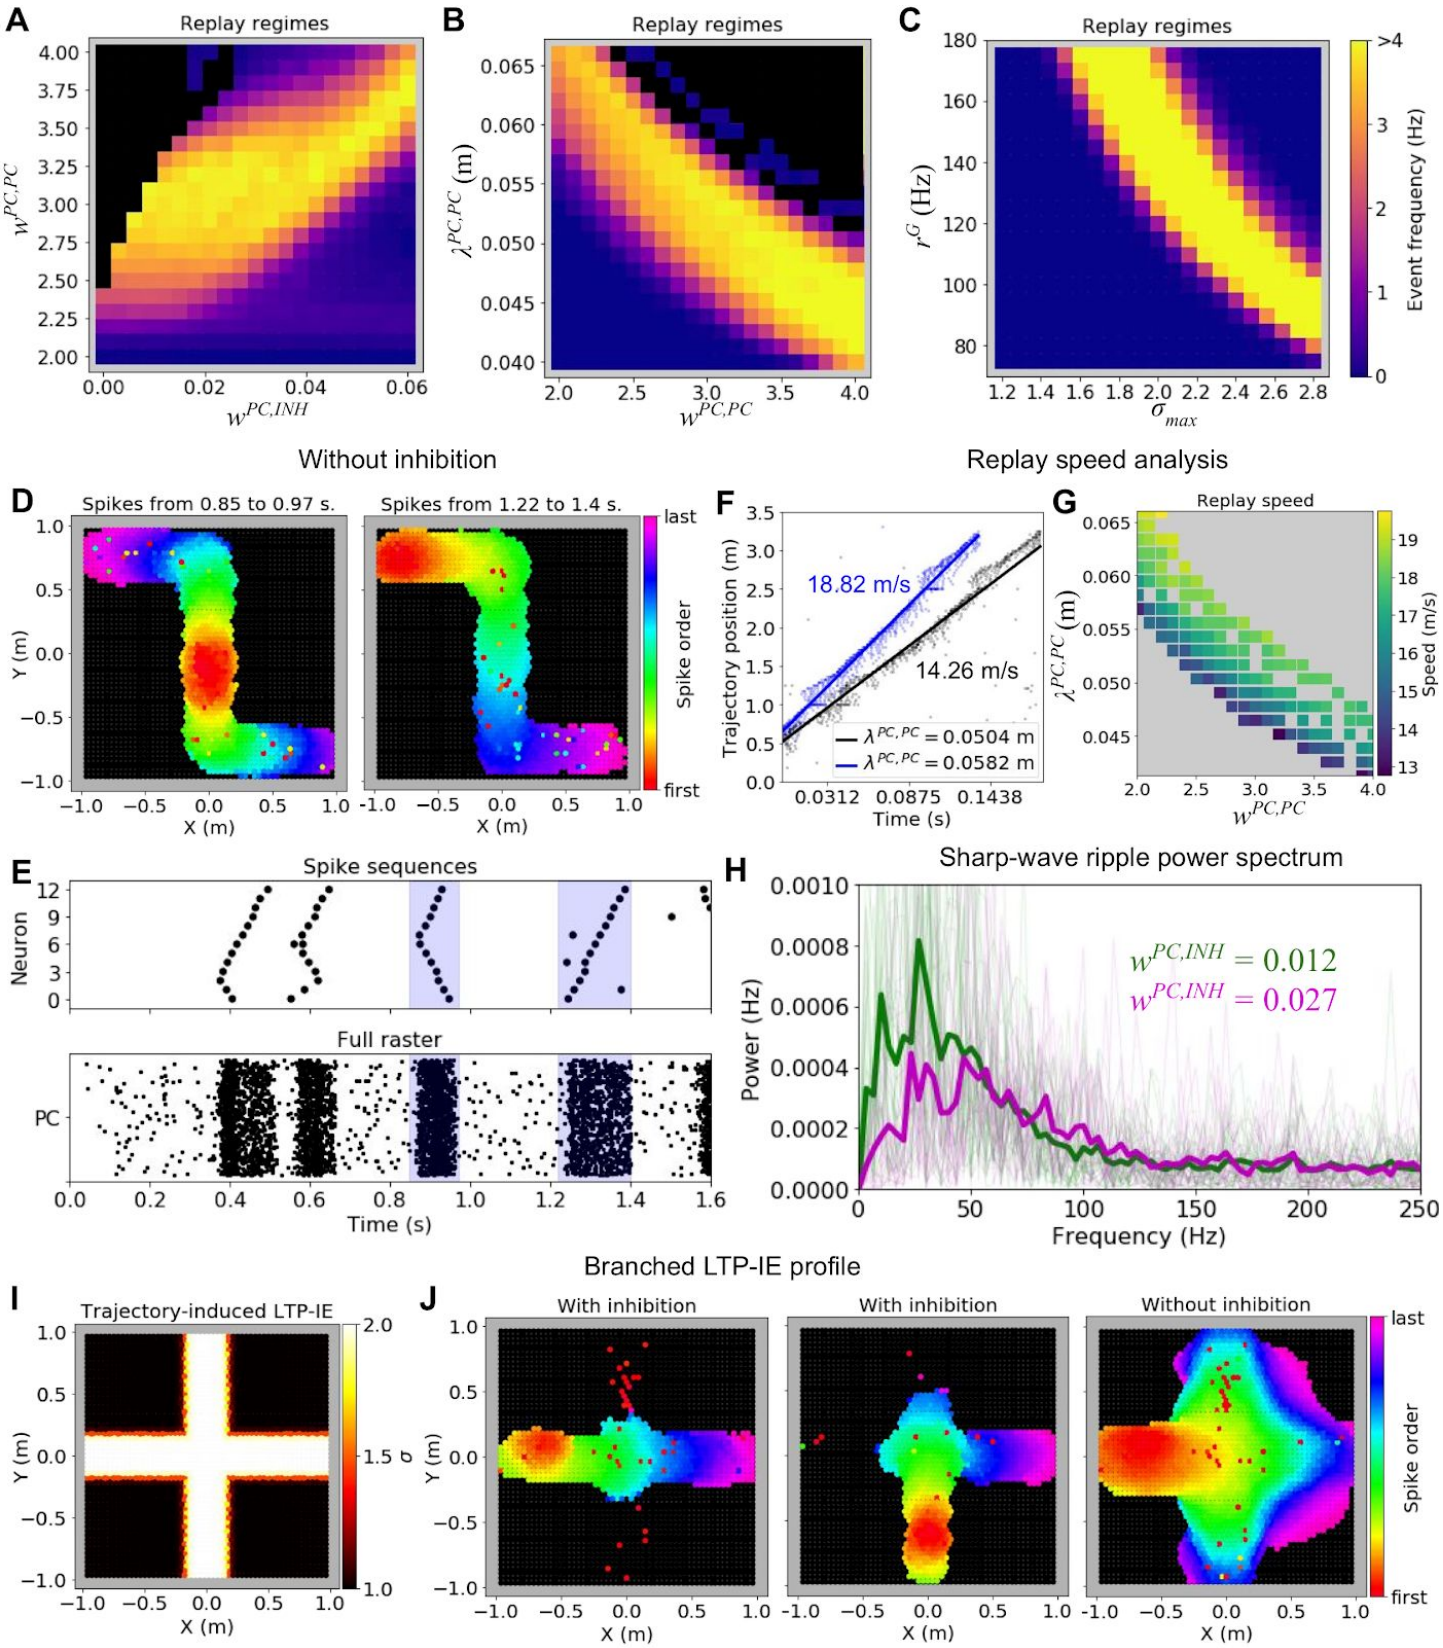

Figure 3: Parameter dependence of LTP-IE-based sequences. *A. Spontaneous replay event frequency as a function of recurrent excitation ($w^{PC,PC}$) and inhibitory feedback ($w^{PC,INH}$) strengths. Black regions indicate regimes of "blowup" replay events, which were not confined to LTP-IE profile. Weights are given in units of the conductance change in response to a presynaptic spike, relative to the leak conductance (inhibitory feedback weights are subsantially smaller because of the high connection probability between PCs and inhibitory cells). B. As in A, but as a function of excitatory connectivity length scale ($\lambda^{PC,PC}$) and $w^{PC,PC}$. C. As in A, B, but as*

function of gate input frequency (r^G) and the maximum LTP-IE level (σ_{max}) a cell could attain. D. Sequential firing of active cells (black cells did not spike) during replay events in a network without inhibition. E. Partial (top) and full (bottom) raster plots showing replay event structure in network without inhibition. Cells in top raster are ordered by tuning positions indicated in 2E. F. Virtual replay speed calculations for spontaneous replay events in networks with two different connectivity length scales $\lambda^{PC,PC}$. Point clouds show linear position of place fields of spiking PCs vs their spike times during replay. Thick lines are best-fit regression lines. G. Estimated replay speed as a function of $\lambda^{PC,PC}$ and $w^{PC,PC}$. H. Power spectral density of PC population activity during replay events for two different levels of inhibitory feedback $w^{PC,INH}$. Thin lines correspond to individual replay events, and thick lines to event averages. I. Example branched LTP-IE profile. J. Spike order during spontaneous replay events given branched LTP-IE profile in I. Left and middle: example events in network with inhibition; right: example event in network without inhibition.

In addition to being robust to changes in $w^{PC,PC}$, spontaneous replay events also arose across a substantial range of recurrent connectivity length scales $\lambda^{PC,PC}$ (***Figure 3B***). Notably, $\lambda^{PC,PC}$ *and* $w^{PC,PC}$ both shaped the "virtual relay speed" of the replayed trajectory, with increases in either parameter generally yielding faster replay (***Figure 3F-G***). This result corroborates the hypothesis that the speed of replay events emerges from internal network structure, as opposed to the temporal structure of the behavioral trajectory (Davidson, Kloosterman, and Wilson, 2009). Spontaneous replay events also occurred reliably across a range of maximal LTP-IE levels (σ_{max}) and random gating input rates r^G (***Figure 3C***), suggesting that neither of these variables would have to be precisely tuned *in vivo* for successful sequence replay.

We next investigated two other aspects of our sequence replay model: sharp-wave ripples (SWRs) and sequence propagation with branched LTP-IE profiles. Inspired by the resemblance of the population firing rates accompanying LTP-IE-induced replay in our network model to experimentally observed sharp-wave ripples (***Figure 2I***) (Csicsvari, et al., 1999; Carr, Jadhav, and Frank, 2011), we computed the power spectral density (PSD) of the population PC firing rates accompanying replay events for networks constructed with two different parameter sets that differed in their level of inhibitory feedback (***Figure 3H***). The resulting PSDs showed salient peaks at tens of Hz, suggesting replay in our network is accompanied by oscillatory SWR-like events. While empirical SWRs in CA3 occur near higher frequencies of 100-130 Hz (Csicsvari, et al., 1999), the two networks we investigated also exhibited different PSD peaks, with the more strongly inhibited PC population displaying a higher peak frequency. This suggests the SWR oscillation frequencies in our model may be determined by the balance of excitatory and inhibitory feedback strengths and that parameter regimes may exist where SWRs alongside replay occur at biological frequencies.

Since realistic locomotor trajectories often intersect with one another, we also investigated how our model behaved when the trajectory reflected in the LTP-IE profile contained multiple pathways branching out from a central intersection (***Figure 3I***). In this scenario we observed that spontaneous replay events in our model could travel through the intersection, selecting a single path along which to continue propagating (***Figure 3J, left and middle***). When inhibition was removed from the model, LTP-IE-induced spike sequences propagated along all branches after the intersection (***Figure 3J, right***), suggesting inhibition likely imposes a winner-take-all-like interaction among competing sequences so that only one sequence continues propagating.

These results show that experience-dependent LTP-IE at biological levels can rapidly and selectively modulate sequence propagation through a network of realistic spiking neurons across a range of parameter regimes.

LTP-IE-based sequences can encode temporary stimulus-response mappings

A popular hypothesis in cognitive neuroscience is the multiplexing and interaction of mnemonic and spatial representations, contributing, for example, to the formation of “memory maps”. Indeed, hippocampus’ experimentally observed roles in both memory and spatial navigation lends compelling evidence to this notion (Schiller, et al., 2015). The potential neural mechanisms underlying this phenomena, however, remain poorly understood. To demonstrate how LTP-IE could shape temporary memory storage and computation more generally, beyond simply replaying recent locomotor trajectories, we considered the simple task of requiring a network to represent one of two possible stimulus-response mappings, and then asked how LTP-IE could induce these mappings in the network (***Figure 4A***). We assume such an induction could be evoked by appropriately transformed sensorimotor or cognitive inputs from upstream areas, but we did not model this explicitly, as our goal was to demonstrate the final storage and recall of the mapping. Note that in the following analysis we have used a larger r^G (200 Hz) and decreased $w^{PC,G}$ (0.5) and σ_{max} (1.84), to reduce spontaneous activity while still allowing propagation of sequences triggered by a brief current injection.

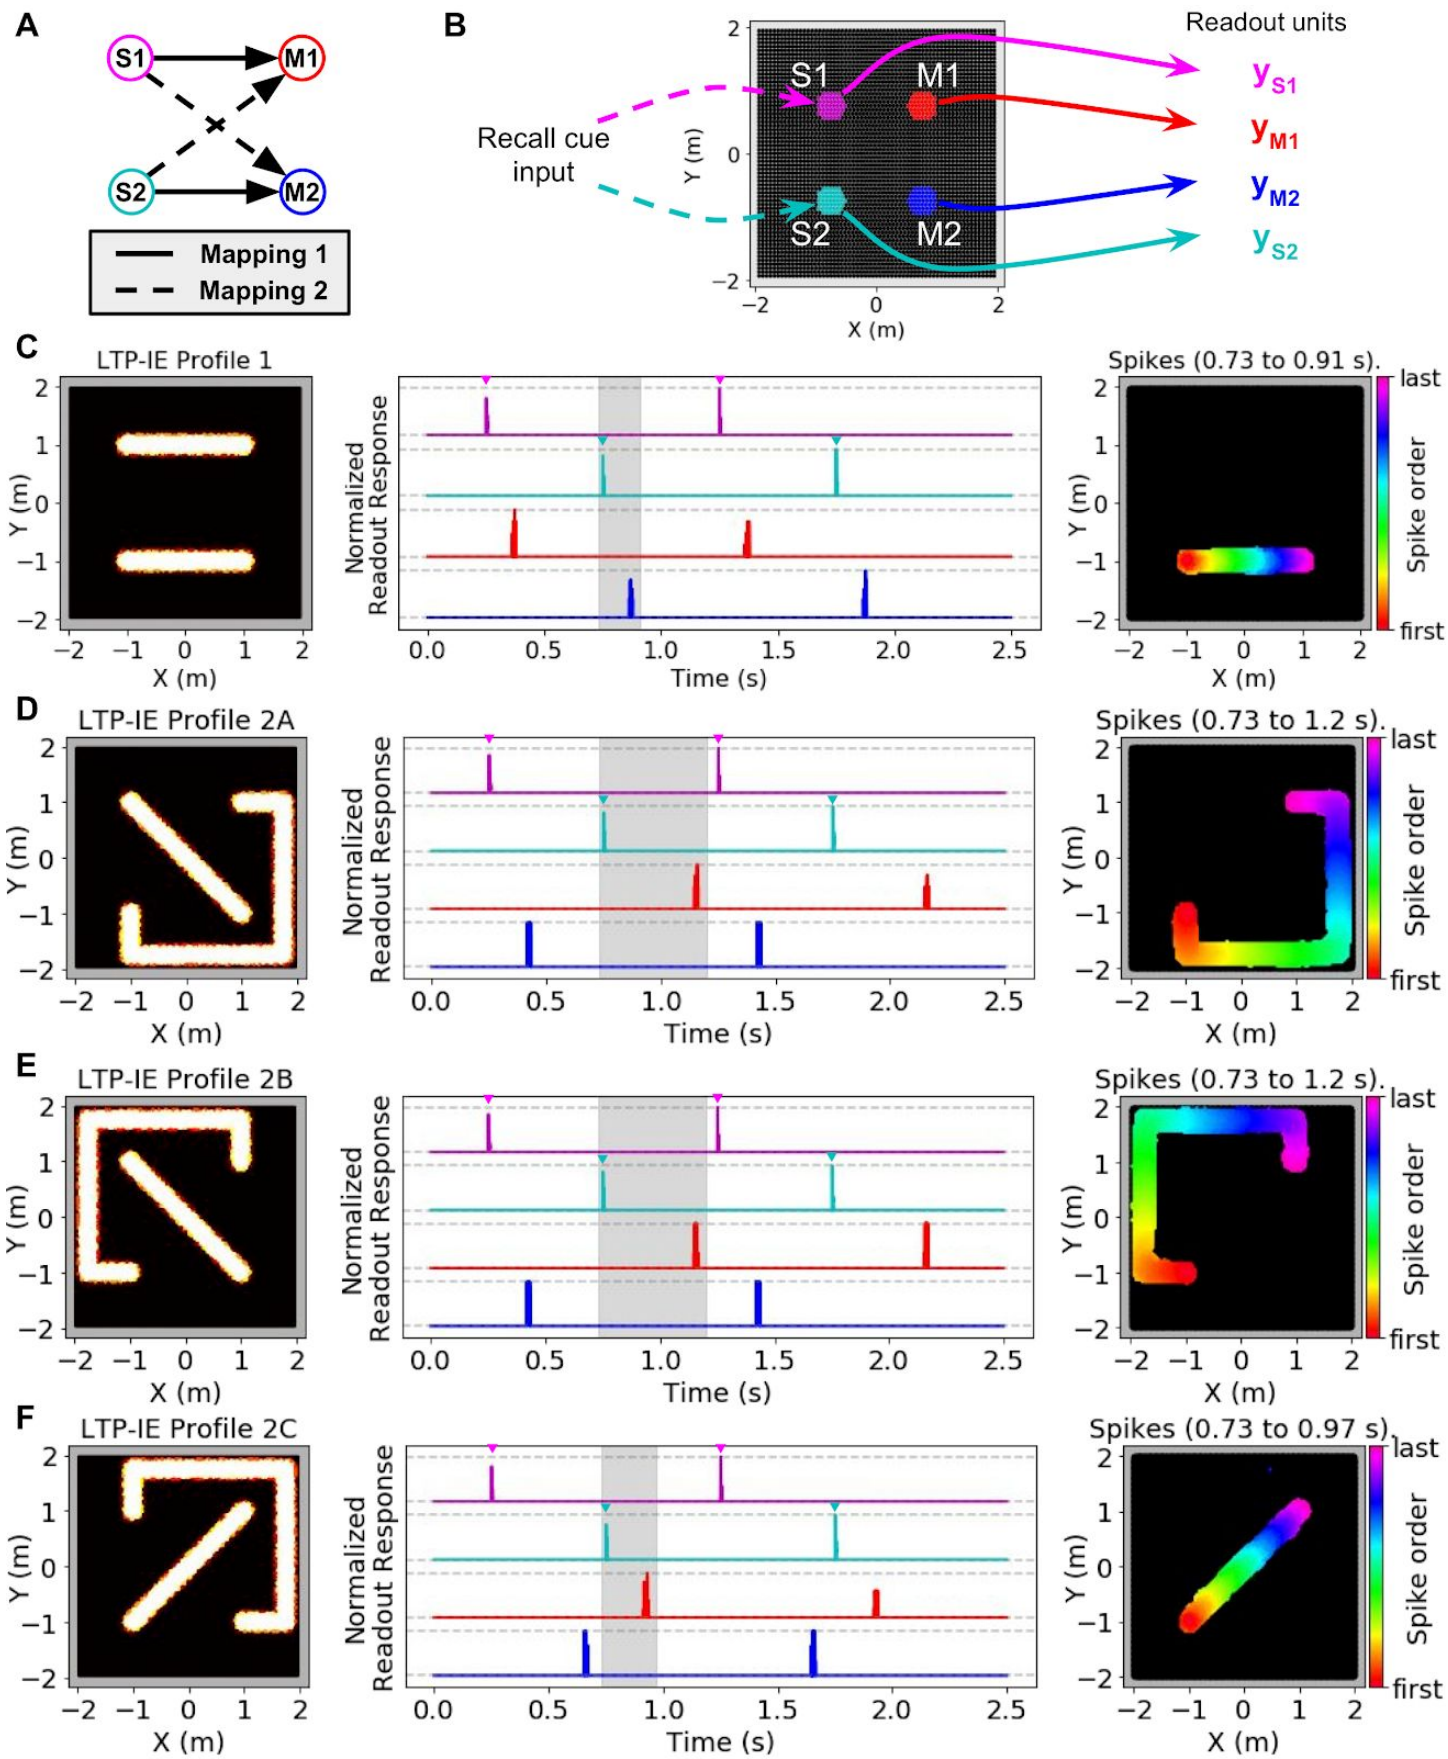

Figure 4: Using LTP-IE-based sequential activity propagation to maintain and decode pairwise associations. *Here we used a network with increased r^G (200 Hz) and decreased $w^{PC,G}$ (0.5) and σ^{max} (1.84) , which decreased spontaneous activity without affecting propagation of sequences triggered by direct current*

injection (see Methods). A. Example of two different mappings between a pair of stimuli (S1, S2) and a pair of responses (M1, M2). In Mapping 1 (solid), activating S1 should activate M1, and activating S2 should activate M2. In Mapping 2 (dashed) activating S1 should activate M2, and activating S2 should activate M1. B. Simulation protocol. Either S1 or S2 is activated with a recall cue (injected current input into colored cells). Readout units average all activity from either S1, S2, M1, or M2 units, indicated by colors. C. Left: Example LTP-IE-based encoding of Mapping 1 with neurons laid out in 2-D space. Middle: Time-dependent readout responses (normalized to maximum readout response over the 2.5 s simulation), with colors corresponding to readout units depicted in B. S1 and S2 were each alternately activated 3 times by appropriate recall cue (direct current injection into the relevant neurons is indicated by the inverted colored triangles), and all readout responses were plotted. In this example, activating S1 (magenta) causes M1 (red) to activate, and activating S2 (cyan) causes M2 (blue) to activate, due to spike propagation along paths defined by LTP-IE profile in C-Left. Right: Order of first spikes of all neurons that spiked during shaded time period in C-Middle. D. Example LTP-IE-based encoding of Mapping 2 with neurons laid out in same 2-D coordinate space as in C, along with readout responses and spike order during shaded epoch. E. Same as C-D but for an LTP-IE profile encoding Mapping 2 that is distinct from that in D. F. Same as D and E but for a third LTP-IE profile encoding Mapping 2.

LTP-IE was able to induce multiple stimulus-response mappings in the network, each in multiple ways. To demonstrate this we first assumed that specific regions in the network corresponded to different stimuli or responses. In other words, we imagined that stimulus regions S1 and S2 might receive inputs containing sensory information from the environment, whereas the response regions M1 and M2 might send outputs controlling different motor commands (***Figure 4B***). To record activity in S1, S2, M1, or M2 we included a unique readout unit for each region representing the region's average activity (***Figure 4B***). We then introduced non-intersecting LTP-IE paths into the network connecting each stimulus region to its associated motor output region, and recorded the ability of an input current trigger into each stimulus region to subsequently activate the correct motor output (***Figure 4C-F***). We first verified our idea with the mapping (S1→M1, S2→M2) in which each stimulus was connected via a one-dimensional LTP-IE profile to its appropriate motor output (***Figure 4C, left***). Indeed, triggering each stimulus region with a depolarizing current input led to subsequent activation of its corresponding motor output after a short delay (***Figure 4C, middle and right***). We next explored the mapping (S1→M2, S2→M1), which in our setup required at least one LTP-IE path to take a roundabout course through the network (***Figure 4D, left***). As before, triggering each stimulus region led to activation of its corresponding motor output, with the longer LTP-IE path reflected in a longer stimulus-response delay (***Figure 4D, middle and right***). Notably, the same stimulus-response mapping could be implemented via LTP-IE in several different ways (***Figure 4D-F***), with the different LTP-IE paths reflected in the different stimulus-response delays. This contrasts with alternative models of temporarily binding together distinct components of a neural network, which typically suppose a unique structure of the mapping/binding representation (Raffone and Wolters, 2001; Botvinick and Watanabe, 2007; Swan and Wyble, 2014). Thus, LTP-IE combined with recurrent connectivity might serve as a biophysically plausible and highly flexible substrate for inducing temporary stimulus-response mappings in a recurrent spiking network. This may be a potential key property enabling rapid and flexible induction of temporary information-processing patterns in brain networks underlying cognitive flexibility.

Encoding complex sequences and non-spatial mappings with LTP-IE

What are the algorithmic limitations of LTP-IE-based sequence induction? For instance, while we have discussed the storage and reactivation of sequences corresponding to simple (non-intersecting) trajectories through space, it is natural to ask whether more arbitrary sequences and mappings can be stored by LTP-IE. To this end we developed a reduced network model capturing the core property of LTP-IE-based reshaping of

network dynamics. This allowed us to separate the generic algorithmic properties of LTP-IE- and excitability-based computation from those tied to its specific biophysical implementation.

Our reduced network obeys the following rules. First, it operates in discrete time, with upstream inputs to a neuron persisting for one timestep, and in which a neuron spikes if those inputs surpass a spike threshold. When a neuron spikes, this spike is converted to inputs to downstream neurons, weighted by the connectivity strengths to those neurons (only excitatory connections were used in our reduced network model), and the spiking neuron enters a refractory period that prevents it from spiking again for the next several timesteps. To model gating inputs G and LTP-IE in our reduced model, we provided all neurons with a constant background input representing G, with LTP-IE-tagged neurons receiving up to twice this background input, thus situating them closer to spike threshold. To mimic inhibitory effects we imposed a maximum number of simultaneously spiking cells, chosen to be those receiving the largest inputs.

Using our reduced model we first demonstrate how LTP-IE could support the storage of *complex* trajectories, in which one of the replayed locations is repeated, i.e. trajectories that intersect themselves (***Figure 5A, inset***). Importantly, here we desire replay without random branching at the intersection point, as one would expect given ***Figure 3J***. To overcome this we introduced an additional tuning dimension to our network. Beyond an (x, y) position, we allowed each cell to be Gaussian-tuned to head-direction θ as well, in line with experimentally observed head-direction cells in the hippocampus (Leutgeb 2000, *Neurosci.*), and we only tagged PCs with LTP-IE if they both (1) lay along the trajectory and (2) had a preferred θ aligned to the animal's presumed travel direction during the original locomotor trajectory. Our assumption is that all three tuning conditions would have to be met during the original trajectory in order for these cells to activate at the physiological levels required for LTP-IE. Accordingly, we modified the recurrent connectivity by only preferentially connecting cells if they both encoded nearby regions of space and had a similar preferred θ.

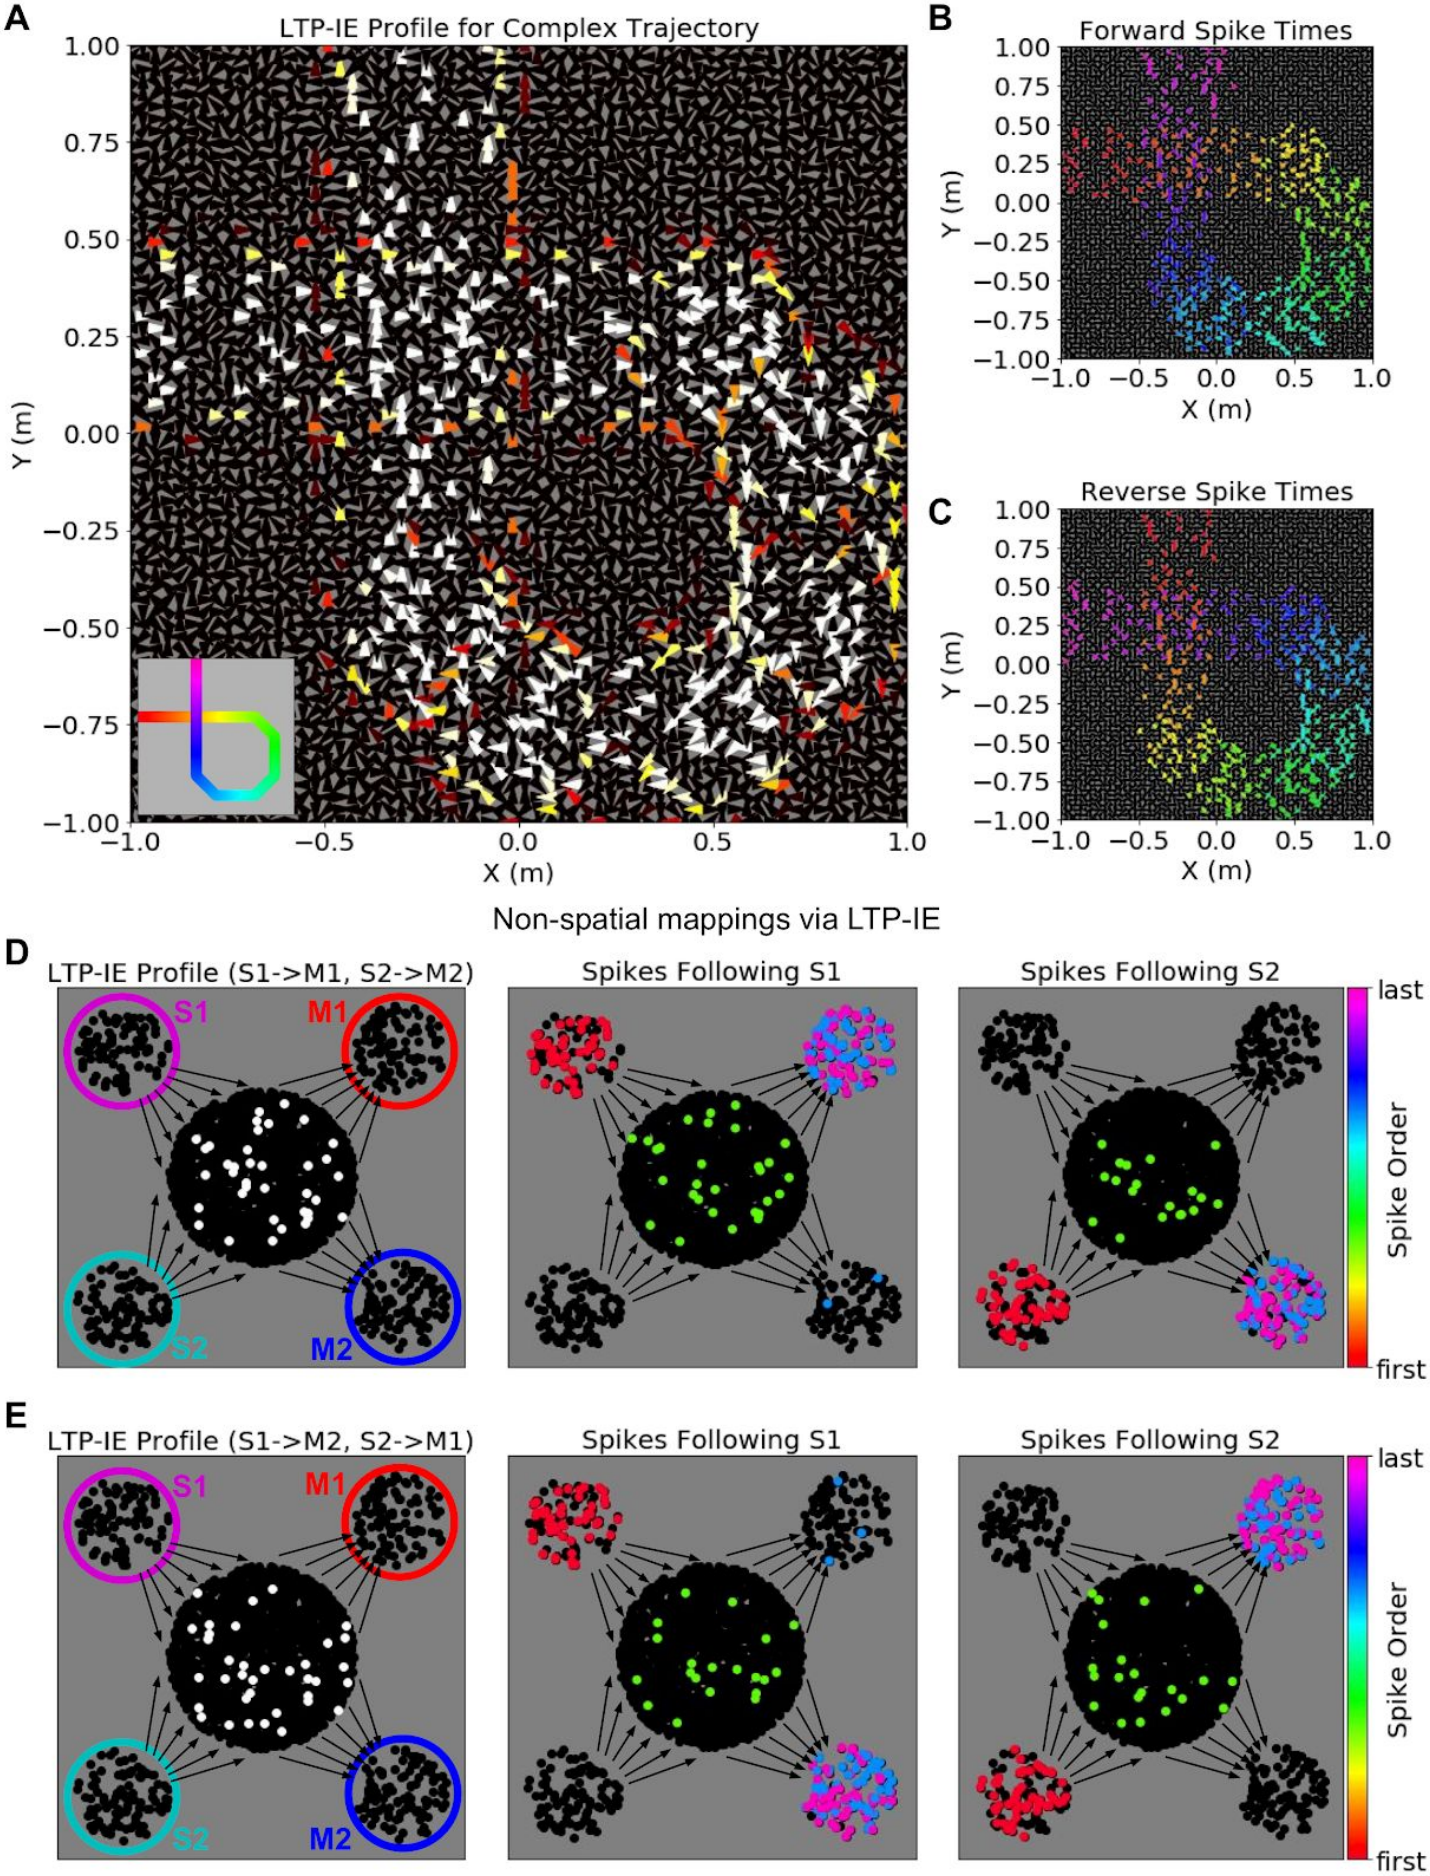
A
LTP-IE Profile for Complex Trajectory
Y (m)
X (m)
B
Forward Spike Times
Y (m)
X (m)
C
Reverse Spike Times
Y (m)
X (m)
Non-spatial mappings via LTP-IE
D
LTP-IE Profile (S1->M1, S2->M2)
S1
M1
S2
M2
Spikes Following S1
Spikes Following S2
last
Spike Order
first
E
LTP-IE Profile (S1->M2, S2->M1)
S1
M1
S2
M2
Spikes Following S1
Spikes Following S2
last
Spike Order
first

Figure 5: Complex sequences and non-spatial mappings via LTP-IE in our reduced network model. *A. LTP-IE profile for self-intersecting trajectory in a network with cells tuned to both position and head direction. Each triangle corresponds to a single cell, with location indicating position tuning and triangle orientation indicating head-direction tuning. This LTP-IE profile corresponds to a "gamma"-shaped trajectory (inset) beginning at (-1, .25) m and ending at (-.25, 1) m, and which intersects itself at (-.25, .25) m. B. "Forward" cell activation sequence given the network and LTP-IE profile in A, triggered by current injection into cells with tuning near (-1, .25) m. C. "Reverse" cell activation sequence given the network and LTP-IE profile in A, triggered by current injection into cells with tuning near (-.25, 1) m. Colors in B, C indicate spike order, with same color code as in 2F. D. Left: LTP-IE profile for one mapping between two stimuli (S1, S2) and two responses (M1, M2) in a network without spatial organization. (The only neurons tagged by LTP-IE were the ones shown in white, which had an LTP-IE level of σ = 2.) Each stimulus (magenta, cyan) or response (red, blue) is represented by an ensemble of 100 neurons, with the "switchboard ensemble" containing 2000 neurons (recurrence with M1, M2 exists but is not depicted here); middle: spike order following current injection into the S1 neurons, leading to subsequent activation of M1; right: spike order following current injection into the S2 neurons, leading to subsequent activation of M2. E. Left: As in D-Left, but with an LTP-IE profile encoding the opposite mapping from D; middle: spike order following current injection into the S1 neurons, leading to subsequent activation of M2; right: spike order following current injection into the S2 neurons, leading to subsequent activation of M1.*

The LTP-IE profile of an example self-intersecting trajectory is shown in ***Figure 5A***. This trajectory begins eastbound at (-1, .25) m, loops around through the lower right quadrant of the environment, crosses itself at (-.25, .25) m, and ends northbound at (-.25, 1) m. Importantly, at the intersection point, only east- and north-tuned cells exhibit strong LTP-IE, and since these cells are not strongly connected, activity propagating through the east-tuned cells should not spread to north-tuned cells, and vice versa. Spike sequence propagation through the network triggered by a brief current injection into cells at the beginning or end of the trajectory is shown in ***Figure 5B,C***, respectively. Activity recapitulating the original trajectory propagated in either the forward or reverse direction, depending on the trigger location, and activity passed through the intersection point in the correct order, without random branching. Note that one would expect random branching probability to increase as the intersection became less orthogonal. This demonstrates how LTP-IE combined with multi-dimensional tuning and corresponding recurrent connectivity can support trajectory sequences containing repeated locations, indicating that LTP-IE-based replay need not be confined to simple sequences.

Finally, we use our reduced model to show how LTP-IE can encode transient stimulus-response mappings without requiring a spatially organized network. While hippocampus has been extensively implicated in spatial computations (Moser, Kropff, and Moser, 2008; Hartley, et al., 2014), exploring how LTP-IE could shape computation without such structure may shed light on how excitability-based computations could arise in other brain networks more generally. To this end we again considered the problem of rapidly inducing one of two sensorimotor mappings (S1→M1, S2→M2) vs (S1→M2, S2→M1) in a network through LTP-IE. Here we represented S1, S2, M1, and M2 as four separate neural ensembles, with the M ensembles containing strong intra-ensemble recurrent excitation but with no recurrence between ensembles (***Figure 5D,E***). Instead, each sensory ensemble sent random projections to a "switchboard" ensemble of neurons, and the switchboard ensemble sent random connections to the motor ensembles (***Figure 5D,E***). For the purposes of this simulation we did not include recurrence within the S or switchboard ensembles.

Given this network architecture, we encoded each sensorimotor mapping by applying LTP-IE to different sets of switchboard neurons. For (S1→M1, S2→M2) we applied LTP-IE to all switchboard neurons that either (1) received inputs from S1 *and* sent outputs to M1 or (2) received inputs from S2 *and* sent outputs to M2 (***Figure***

5D). Intuitively, activating S1 should activate the first set of LTP-IE-tagged neurons, causing M1 to activate, whereas activating S2 should activate the second set of LTP-IE-tagged neurons, causing M2 to activate. This is shown in ***Figure 5D*** (***middle and right***), where activating each stimulus indeed activates a subset of LTP-IE-tagged switchboard neurons, causing the appropriate motor ensemble to subsequently activate. Notably, the alternate mapping (S1→M2, S2→M1), in which LTP-IE is applied to switchboard neurons that either (1) receive inputs from S1 *and* send outputs to M2 or (2) receive inputs from S2 *and* send outputs to M1 (***Figure 5E***), yields an LTP-IE profile distinct from the one encoding (S1→M1, S2→M2); instead, this LTP-IE profile allows signal propagation from S1 to M2 and from S2 to M1. We note that the fundamental idea of opening selective communication channels between pairs of ensembles by changing the spiking properties of an intermediary set of switchboard neurons has been explored in substantial detail in the quite similar "binding pool" model of Swan and Wyble (2014). Here we show how the idea also fits naturally into the theory of excitability-based computation. We do not explore how the specific LTP-IE profiles above could be induced in the first place, assuming instead that they could arise via upstream transformations of sensorimotor information. Thus, even in a randomly connected network with no spatial structure, LTP-IE can yield selective stimulus-response mappings, demonstrating the potential of excitability changes as a substrate for cognitive flexibility more generally.

Discussion

Sequential spiking activity is a key feature of recurrent neural network dynamics, potentially reflecting information flow and computations within the network. One noteworthy empirical example of internally generated sequences is the replay of spike sequences representing an animal's recent sensorimotor sequence, such as a locomotor trajectory or sequence of viewed images. Sequence replay has been observed *in vivo* in both hippocampus and cortex during awake quiescent periods in both rodents and primates (Foster and Wilson, 2006; Davidson, Kloosterman, and Wilson, 2009; Karlsson and Frank, 2009; Gupta, et al., 2010; Carr, Jadhav, and Frank, 2011; Eagleman and Dragoi, 2012) and is thought to be involved in memory consolidation (Carr, Jadhav, and Frank, 2011; Jadhav, et al., 2012; Ólafsdóttir, Carpenter, and Barry, 2017; Zielinski, Tang, and Jadhav, 2017) and navigational planning (Foster and Knierim, 2012; Pfeiffer and Foster, 2013). Since replay appears to be strongly modulated by recent experience, it serves as a compelling entrypoint for uncovering the biophysical mechanisms that support rapid and flexible modulation of sequential dynamics and information processing in neural networks.

While substantial work has shown how gradual modifications to a network's recurrent connections can induce sequences into its dynamics repertoire (Sussillo and Abbott, 2009; Klampfl and Maass, 2013; Laje and Buonomano, 2013; Rajan, Harvey, and Tank, 2016], less is known about mechanisms that could act on faster timescales. Here we have demonstrated the sufficiency of an empirical, strong, and fast-acting heterosynaptic plasticity mechanism known as long-term potentiation of intrinsic excitability (LTP-IE) (Hyun, et al., 2013; Hyun, et al., 2015), applied to a spiking network, to coerce the network into generating selective sequences. Briefly, LTP-IE biases recently active cells towards recruitment during replay, with their reactivation order determined by pre-existing but fixed recurrent connectivity. In particular, in our spatially organized spiking network LTP-IE acted to select pathways through the network that guided activity propagation during replay. Similar to "synfire chains", feedforward neural networks supporting stable activity propagation (Abeles, 1982; Abeles, et al., 1994; Hermann, Hertz, and Prügel-Bennett, 1995; Schrader, et al., 2008), the potentiated pathways in our network supported stable and self-sustaining spike sequences, which despite symmetric connectivity propagated without reversing direction due to the moderate refractory period we included (8 ms). One can therefore conceptualize our spatial network as akin to a reservoir of trajectories through space that can be selected for replay in an experience-dependent manner through LTP-IE.

Our spiking network model exhibited several interesting features. First, despite substantial variability in the membrane voltages of LTP-IE-tagged cells and overlap with non-LTP-IE-tagged cells (***Figure 1D***), replay events exhibited stable propagation along LTP-IE-defined network paths, due to averaging effects of populations of LTP-IE-tagged cells. Second, sequences replayed at timescales determined by the network rather than the original trajectory (***Figure 3F***), and did so in both forward and reverse (***Figure 2F,G***), capturing these two key empirical features of hippocampal replay (Foster and Wilson, 2006; Davidson, Kloosterman, and Wilson, 2009; Karlsson and Frank, 2009; Gupta, et al., 2010; Carr, Jadhav, and Frank, 2011). Additionally, inhibition played an important role in our model, promoting replay in one direction only (***Figure 2F***) and inducing competition among multiple sequences in branched trajectories (***Figure 3D,I,J***), even though inhibition bore no spatial organization itself. Inhibition also appeared to shape the power spectra of the sharp-wave-ripple (SWR) events in our model arising in the PC population average alongside replayed sequences; moreover, these events were consistent with the observation that hippocampal SWRs likely arise spontaneously in CA3 and propagate to CA1 (Buzsáki, 1986; Csicsvari, et al., 2000; Carr, Jadhav, and Frank, 2011; Ramirez-Vallegas, et al., 2018), since our model's excitatory recurrence was inspired by that found in CA3 but not CA1 (Lisman, 1999).

Finally, while our biophysical model used a 2-D spatial network and hippocampal LTP-IE mechanism, excitability changes might support more general modulation in dynamics and computation in networks with fixed recurrence. To explore this we used a reduced network and LTP-IE model attempting to capture the core algorithmic features of excitability-modulated dynamics independently from a specific biophysical implementation. For instance, LTP-IE might also occur purely intrinsically, without an explicit correlation with EPSP sizes (Ohtsuki and Hansel, 2018). Using the reduced model we showed how trajectories that intersect themselves in 2-D could be successfully replayed given an additional tuning dimension like head direction (***Figure 5A-C***), and further how LTP-IE could store and recall sensorimotor mappings in a randomly connected network with no spatial structure at all (***Figure 5D-E***), similar to the "binding pool" model of Swan and Wyble (2014). This latter example may be relevant to excitability-based computations in cortex, where LTP-IE-like phenomena have been observed (Sourdet, et al., 2003; Paz, et al., 2009), but whose internal network structures do not necessarily exhibit clear spatial organization. We note, however, that preferential connectivity among similarly tuned cells has been observed in mouse visual cortex (Ko, et al., 2011), suggesting a mechanism similar to the one we have explored might support spontaneous sequence propagation among neurons encoding visual space. Such spontaneous sequences might in turn modulate neural responses to incoming visual stimuli, potentially supporting expectation-modulated sensory processing.

History-dependent excitability changes might also occur not just at the rapid timescales that have motivated our work but over slower timescales as well (Zhang and Linden, 2003; Titley, Brunel, and Hansel, 2017). For example, it was recently found that small neuronal ensembles could be "imprinted" into a mouse cortical network and subsequently reactivated *in vivo* via repeated optogenetic stimulation (Carrillo-Reid, et al., 2016). While this result was suggested to provide evidence of Hebbian (homosynaptic) plasticity, our work suggests that the same result could arise via increasing neural excitabilities instead. Untangling these mechanisms and understanding how they interact presents an exciting avenue for further investigation. Overall, our work demonstrates how excitability changes alone might quickly and selectively reshape network dynamics, implicating their potential role in storing memories and encoding stimulus-response mappings, and more generally in organizing flexible computations over rapid timescales.

Comparison to existing models

While we do not aim to provide a complete account of hippocampal replay, but rather to demonstrate how specific biophysical mechanisms can rapidly generate selective network sequences, we believe it is still

worthwhile to compare our model to existing models for rapidly storing and recalling sequences in neural systems, both in terms of biological plausibility and computational robustness. One family of models for encoding and decoding sequences in hippocampus supposes that sequential information is stored in the timing of spikes relative to theta (4-8 Hz) and gamma (~40 Hz) oscillatory cycles in the hippocampus (Lisman and Idiart, 1995; Lisman 1999). Briefly, gamma cycles are "nested" within theta cycles, and a sequence of stimuli is stored by stimulus-specific neurons spiking at unique gamma cycles within the encompassing theta cycle. While these models provide a persuasive computational account of how oscillations might be used to store temporary information, it is not clear how the sensory input would be appropriately transformed so as to be "entered" into the oscillation cycle at the correct time, how stable these mechanisms are in the face of noise, nor how multiple sequences might be stored simultaneously; in particular, sequence information could be lost if the oscillation were disrupted, and one would require an independent nested oscillation for a second sequence. In our LTP-IE-based spiking-network model, however, memories are directly entered into the network by physiological spiking patterns (***Figure 1***), are stored in effective synaptic weight changes likely driven by ion channel inactivation (as opposed to persistent spiking activity) and might therefore be more robust to noise, and multiple sequences can theoretically be stored in the same network in a manner similar to how we have stored multiple stimulus-response mappings in ***Figure 4***.

The second main family of replay models either explicitly (Molter, Sato, and Yamaguchi, 2007; Veliz-Cuba, et al., 2015; Haga and Fukai, 2018) or implicitly (Chenkov, Sprekeler, and Kempter, 2017) assumes the sequences produced by the network are initially induced in the network via modification of recurrent network weights, e.g. through homosynaptic mechanisms like spike-timing-dependent plasticity (STDP) (Markram, et al., 1997; Bi and Poo, 1998; Fiete, et al., 2010). While such models can indeed selectively bias network sequences, given the hypothesized weak magnitude of STDP, it is unclear whether such a mechanism could account for the awake trajectory replay observed in rodents, which can sometimes occur even when the animal has experienced the original trajectory only once (Foster and Wilson, 2006), nor whether it could be used to immediately induce novel computations that the animal cannot afford to learn over extensive repetitions. In our model we rely on a strong, fast-acting mechanism, LTP-IE, which although heterosynaptic (unlike STDP, LTP-IE follows spiking regardless of which input elicited the spikes [Hyun, et al., 2013; Hyun, et al., 2015]) suffices in biasing a network toward producing highly specific sequences: LTP-IE "tags" cells in the recurrent network with a higher probability of spiking in response to input (***Figure 1***), and sequential ordering is reconstructed by the existing but unmodified recurrent connectivity. While recent research has begun to investigate how small, biophysically plausible recurrent weight changes might encode new memories (Curto, Degeratu, and Itskov, 2012; Yger, Stimberg, and Brette, 2015), to our knowledge this has not been applied in the context of sequences. We would also like to note that experience-dependent short-term plasticity mechanisms, like lingering presynaptic calcium following spiking, can transiently make a neural ensemble more sensitive to subsequent inputs (Mongillo, Barak, and Tsodyks, 2008), suggesting such a mechanism might be able to play a similar role to excitability changes.

In real brain networks, excitability changes likely interact with recurrent synaptic mechanisms like STDP or synaptic long-term potentiation (LTP) to shape computation. For instance, it has been observed that changes in how dendrites integrate EPSPs to trigger spiking (effectively an excitability change since this would affect spiking even if EPSPs remained stable) can co-occur with canonical STDP or LTP (Daoudal, Hanada, and Debanne, 2002; Dauodal and Debanne, 2003; Campanac and Debanne, 2008; Debanne and Poo, 2010). Excitability vs synaptic changes may also coordinate the transference of memories across timescales. For example, as observed *in vivo* and as predicted by our model, hippocampal sequences that replay during awake quiescence are significantly compressed relative to behavioral timescales, suggesting replayed spikes may occur within the appropriate time windows for slower-acting STDP (Markram, et al., 1997; Bi and Poo, 1998; Dan and Poo, 2004]. Indeed, it was shown that replaying place-tuned firing sequences *in vitro* that were

previously recorded *in vivo* could induce long-term connectivity changes between cells with overlapping place fields (Isaac, et al., 2009, *J. Neurosci.*). Thus, LTP-IE may serve as a temporary buffer for sequences that could become embedded in the network as long-term memories through gradual synaptic changes triggered by repeated replay. Additionally, LTP-IE-based replay in our model occurred over a wide range of parameters. While our networks were spatially homogeneous for all parameter sets (i.e. connectivity depended only on *differences* between place fields), this suggests potential robustness to connectivity heterogeneities within a single network, e.g. encoding preferred paths through an environment learned over longer timescales. Intuitively, paths defined by pre-existing stronger connections should bias replay in their favor but might also be overwritten by transient excitability changes encoding paths just traveled. Overall, short-term memory and computational flexibility and their subsequent effect on memory consolidation and long-term network structure likely rely on a host of dynamics and plasticity mechanisms, but we propose LTP-IE may play a significant role.

Finally, selectively biasing cell recruitment without changing recurrence generalizes beyond increasing spike responses through LTP-IE. An alternative mechanism is to simply send stronger baseline inputs to a subset of neurons, corresponding in our model to increasing the gate input rate r^G for certain cells but not others, yet without applying LTP-IE. Such differential baseline inputs can act as context signals, and if interacting multiplicatively with sensory signals can control the gain of the latter without affecting their tuning; by modulating the gain of different neurons in such a way, one can then selectively and transiently encode context-dependent sensorimotor mappings without modifying synaptic weights (Salinas, 2004, Salinas, 2004a), similar to how one might do so with LTP-IE. Given the functional similarity of selective LTP-IE or differential gain-modulating inputs, one would expect them to support similarly flexible sequential dynamics and signal routing, depending on the pre-existing but fixed recurrence structure. Indeed, neurons in entorhinal cortex, which inspired the gating signals G in our network, have been shown to support history-dependent, graded persistent activity (Egorov, et al., 2002), suggesting the recurrent CA3 network could be subject to differential baseline inputs maintained over short-term memory timescales. Such a mechanism could act in concert with LTP-IE and could potentially act as an additional rapid and controllable gating signal governing effective excitabilities of neurons in the recurrent network.

Model limitations

One question we did not explicitly address is whether and how a decay timescale of LTP-IE would affect sequence storage and replay. Indeed, with no forgetting or selective LTP-IE activation, one would expect all neurons in our model to undergo LTP-IE after the animal had fully explored the environment, thus preventing sequential reactivation along specific trajectories through the network. While this problem would be partially solved by considering additional input dimensions as described above (i.e. so that even if the x-y space were fully explored, only a fraction of the x-y-θ space would have been explored), one might also imagine additional neuromodulatory inputs serving to control the strength of LTP-IE during exploration. Indeed, while Hyun et al. (2013, 2015) did not observe a return of activation-triggered LTP-IE to baseline during the course of their slice experiments, *in vivo* LTP-IE in CA3 might be modulated by additional upstream inputs. For example, the neurotransmitter acetylcholine is known to interact with the voltage-gated potassium channel Kv1.2 (Hattan, et al., 2002), the ion channel hypothesized to underlie LTP-IE (Hyun, et al., 2015), and hippocampus receives state-dependent cholinergic inputs from medial septum (Yoder and Pang, 2005; Mamad, et al., 2017). Furthermore, memory-related excitability changes in dentate gyrus have been shown to decay after two hours (Pignatelli, et al., 2019), suggesting similar decay mechanisms may exist in CA3. Thus, sufficient machinery for erasing or selectively modulating the overall effects of LTP-IE, e.g. as a function of arousal or motivation state, may converge in the brain region that has served as the inspiration for our model. Clinically, one would also expect regulation of the spatiotemporal extent of excitability changes to be crucial in preventing pathological dynamics, such as incomplete signal transmission that might lead to short-term memory deficits (Vallar and

Shallice, 1990), or hyperexcitability without sufficient compensatory inhibition leading to seizure-like events, as exhibited in the "blowup" regime of our spiking network model (***Figure 3***).

Additionally, while we have assumed the upstream "gating" inputs in our model are stochastic with homogeneous firing rates, cortical activity likely contains higher order structure beyond mean firing rates, such as the grid-like representations of space observed in neural firing patterns in entorhinal cortex (O'Keefe and Burgess, 2005; Yamamoto and Tonegawa, 2017), the hippocampal input that inspired our gating inputs G. Our model demonstrates that replay can occur absent additional mnemonic information, but one would imagine such additional information would only increase the capacity and robustness for memory storage and computational flexibility, since this would act as an additional memory buffer. Furthermore, as mentioned above, entorhinal neurons may have the capacity for history-dependent persistent activity persisting over the same timescales as LTP-IE (Egorov, et al., 2002), potentially supporting even more flexible heterogeneous gating inputs to the recurrent network. Investigating how such mechanisms would compete or synergize with LTP-IE could be a fruitful line of future inquiry.

What factors missing from our model might set the duration of replay events in more realistic situations? In our model, replay durations were determined by the length of the trajectory through the environment (4 m total in our simulations) and the replay speed was determined by the network parameters. Indeed, while empirical replay events and their accompanying sharp-wave ripples have typically been observed to occur over approximately 100 ms (Carr, Jadhav, and Frank, 2011), when animals are allowed to explore larger environments replay can last several hundred milliseconds (Davidson, Kloosterman, and Wilson, 2009), suggesting replay durations might be largely determined by the size of the environment, as in our model. Parallel lines of evidence, however, suggest hippocampal sequence replay may be gated by particular phases of slow, theta oscillations (4-8 Hz), which would constrain replay to occur within approximately 100-200 ms time windows (Lisman and Idiart, 1995; Fuentemilla, et al., 2010; Lisman, 2010). While we excluded additional oscillatory components from our model for the sake of parsimony, one can imagine imposing time-dependent gating input rates $r^G(t)$ that, while spatially homogeneous, would oscillate at theta-like frequencies. One would expect such gating oscillations to constrain replay events to explicit phases of the theta cycle, thus preventing sequences of arbitrary duration from replaying.

Finally, an intriguing line of research suggests spatiotemporal propagation of activity in certain neural networks may follow a power-law distribution (Beggs and Plenz, 2003; Beggs and Plenz, 2004; Plenz and Thiagarajan, 2007; Fiete, et al., 2010; Scarpetta and de Candia, 2013). Activity patterns distributed as such would imply scale invariance (scale-freeness) of activity structures and suggest the network state sits at a phase transition between dynamical regimes that may be optimal for processing inputs that vary over a large range (Shew, et al., 2009). Since activity levels and durations in our network were capped by inhibition and refractoriness, and by the limited environment size, we would not expect our network to exhibit scale-free activity patterns; however, exploring what network mechanisms and connectivity statistics could yield power-law distributed sequences through LTP-IE and examining how these might be involved in sequence-based stimulus processing would be an exciting line of future inquiry.

Experimental predictions

Despite its parsimony, our model of LTP-IE-based network sequences makes specific predictions that could be tested in hippocampal replay experiments. First, pharmacological blocking of LTP-IE during exploration should prevent the encoding of new trajectory information into CA3, consequently preventing subsequent replay. This could be achieved by *in vivo* application of nimodipine or PP2, which block intracellular calcium increase, protein tyrosine kinase activation, and Kv1.2 channel endocytosis, all potential mechanisms of LTP-IE, and

which have been shown to block LTP-IE *in vitro* (Hyun, et al., 2015). Second, inactivating medial entorhinal layer II (MEC II), which projects to CA3 and which we associate with gating inputs G in our analogy with hippocampal circuitry, should inhibit CA3 sequential replay events. Indeed, inactivation of MEC III inputs has been shown to inhibit forms of replay in CA1 (Yamamoto and Tonegawa, 2017); we propose that a parallel effect would be seen in MEC II inputs to CA3. Finally, the core prediction of our sequence induction model is that the neurons do not have to be initially activated in the order (or the reverse order) in which they will later reactivate. As the only information explicitly added to the network at the time of encoding is the *set* of neurons activated, speed or ordering information must be reconstructed from existing connectivity. While identifying the position tunings of individual neurons *in vivo* and subsequently artificially activating them out of order currently presents a substantial technical challenge, one could potentially achieve a similar effect using a virtual reality environment with place-specific sensory cues, and in which specific places along a virtual trajectory were presented out of order. Our model predicts that as long as the animal had sufficient knowledge of which sensory cues corresponded to which environment locations, replay would occur in an order corresponding to a path through the environment, rather than the order in which the virtual places were presented.

Computational significance

Besides revealing a potential core mechanism for sequence induction, our model lends two key insights into the structure and control of replay in biological neural networks. First, it suggests a link between isotropy in pre-existing spatial representations in the network (i.e. the spatially organized connectivity) and the experimental observation that replay sometimes occurs in the forward, and sometimes in the reverse order relative to the neurons' initial activation during the trajectory (Foster and Wilson, 2006; Davidson, Kloosterman, and Wilson, 2009; Carr, Jadhav, and Frank, 2011). In particular, since no directionality information is stored by LTP-IE (which only tags which neurons were recently active), replay should occur with equal chance in forward or reverse. (Note however, that as observed in (Davidson, Kloosterman, and Wilson, 2009), replay events may have an increased probability of originating at cells tuned to the animal's current location, as a result of place-tuned inputs to the network during awake quiescence.) One would therefore not expect to observe reverse replay for neural sequences that do not have a natural spatial embedding, e.g. spoken sentences in humans. Second, the upstream gating signal in our model might serve as a substrate for the short-term state-dependence of replay, e.g. with arousal or motivation controlling when the gating signal is present or absent. This could ensure that replay occurs only at appropriate times (e.g. quiescent rest) without interfering with other computations being performed by the network, such as pattern separation in external sensory inputs (Leutgeb, et al., 2007; Bakker, et al., 2008).

Further, while our spiking network was embedded in 2-D space, our results frame the more general question of how a network embedded on a higher dimensional manifold, or with no natural embedding, supports excitability-based sequence induction and information flow. For example, one could imagine the rapid induction of spike sequences corresponding to movement along 3-D trajectories if the recurrent architecture reflected a 3-D space, as might occur within the 3-D hippocampal place-cell network of bats (Yartsev and Ulanovsky, 2013). More generally it remains an open and intriguing question as to which network structures are most suited to rapid excitability-based flexibility among sequences and behaviorally relevant computations. One example of a network model with no obvious spatial embedding, and which could potentially be recast in the framework of excitability-based information flow, is the "binding pool" model of associative short-term memory proposed by Swan & Wyble (2014). Here, associations are encoded in a set of steady, persistently active "binding" units that allow preferential signal transmission among specific sets of units representing object features, to which the binding units are randomly connected. While the role of such persistent spiking in coding mnemonic information is presently undergoing substantial reevaluation (Sreenivasan, Curtis, and D'Esposito, 2014), one could theoretically replace it with excitability increases, with an equivalent consequence for

information transmission. We explored this possibility in the context of LTP-IE using a reduced model of LTP-IE (***Figure 5D,E***), in which we were able to encode stimulus-response mappings in a random network by applying LTP-IE to a specific set of "switchboard" neurons (analogous to Swan & Wyble's binding units [2014]). Notably, since connectivity here required no training, this architecture may be well suited to transiently storing novel mappings through LTP-IE that the network has never seen before, so long as the neural ensembles encoding the relevant stimuli and responses were sufficiently (but randomly) connected to the switchboard/binding pool. The nuances of how such a network might implemented in biological spiking networks remains an open question, but our present work suggests its theoretical plausibility.

Finally, one can draw an equivalence between the modulation of cellular properties to shape sequences and information flow and selective variation in upstream modulatory inputs to a cell. For example, a neuron receiving strong inhibition might be less likely to participate in a computation than one receiving excitation, since the latter would be more likely to spike in response to inputs, similar to if one had increased its excitability through intrinsic cellular mechanisms. Indeed, it was recently shown that external modulatory inputs to a recurrent network could control the speed at which a computation performed by the network unfolded through time, closely recapitulating experimental results (Remington, et al., 2018) More flexible control of neural information transmission via modulatory inputs has also been explored in the framework of gain-modulating context inputs that can effectively move a network into different stimulus-response mappings, also without changing synaptic connectivity (Salinas, 2004; Salinas, 2004a). We propose that the control of information flow and computation in recurrent networks via effective excitability alone, without structural connectivity changes, might be a rich framework for investigating rapid cognitive flexibility and short-term information storage. This may lead to significant insights into how we adapt our mental processes and behaviors in the face of ever-changing environments and task demands.

Methods

Membrane and synaptic dynamics for spiking network simulation

We modeled single neurons as leaky integrate-and-fire neurons with conductance-based synapses according to the equation

$$\tau_m \frac{dV^i}{dt} = -(V^i - E_L) - g_E^i(t)(V^i - E_E) - g_I^i(t)(V^i - E_I) + I_{ext}^i(t)$$

where V^i is the membrane voltage of neuron i; τ_m is the membrane time constant (50 ms for pyramidal cells [PCs], 5 ms for inhibitory cells [INHs]); E_{leak} is the leak/resting potential (-68 mV for PCs, -60 mV for INHs); g^i_E and g^i_I are the time-varying conductances of excitatory and inhibitory synaptic inputs, respectively; E_E = 0 mV and E_I = -80 mV are the excitatory and inhibitory synaptic reversal potentials, respectively; and I^i_{ext} is any externally injected current to cell i (in units of the expected deflection in membrane voltage). If V^i exceeded a threshold v_{th} = -36 mV (or v_{th} = -50 mV for INHs) at any time step, a spike was emitted and V^i was reset to E_{leak} for a refractory period τ_r = 8 ms (or τ_r = 2 ms for INHs). The outsize PC refractory period ensured that the sequential events in our simulations propagated only in one direction and is consistent with refractory periods measured in hippocampal region CA3 during some experiments (Raastad and Shepherd, 2003).

Synaptic dynamics were conductance-based and were modeled as exponentially filtered trains of weighted delta functions representing input spikes arriving from different upstream neurons. Specifically, for an individual neuron i

$$g_E^i(t) = \sum_j \sum_k w_{ij} h(t) * \delta(t - t_k^j)$$

and similarly g^i_I; here j indexes neurons and k indexes spikes, such that w_{ij} is the synaptic weight from neuron j onto neuron i and t_k^j is the k-th spike of the j-th neuron; and h(t) was a one-sided exponential with time constant τ_E = 2 ms. In our simulations we assumed all synapses were excitatory and did not explicitly model effects of inhibition. Synaptic weight values are unitless and specify the postsynaptic conductance change in response to a presynaptic spike, relative to the leak conductance.

Pyramidal cell tuning parameters

In network simulations, we assigned each PC a position (x^i, y^i) corresponding to the peak of its tuning curve, i.e. the (x, y) position eliciting maximal firing. Positions were assigned to approximately tile a grid spanning -1 to 1 m in both dimensions (***Figure 2, 3***) or -2 to 2 m (***Figure 4***), using 3000 or 12000 PCs, respectively. We used the following position-dependent firing-rate equation to calculate the maximal firing rate we expected to be evoked by the trajectory through the simulated environment (the firing rate evoked by the closest point on the trajectory to the neuron's place field center):

$$r^i(x, y) = r_{max} \exp\left[-((x - x^i)^2 + (y - y^i)^2)/(2\lambda_{PL}^2)\right]$$

where r_{max} = 20 Hz is the maximum firing rate (i.e. when the trajectory passes directly through the neuron's peak tuning position) and λ_{PL} = 0.15 m is the length constant determining how close a neuron's peak tuning must be to the trajectory to significantly activate. All neurons had the same r_{max} and λ_{PL}.

Spiking network architecture

We included an excitatory (PC [pyramidal cell]) and inhibitory (INH) population in our spiking network, as well as random spiking inputs from a gating signal G impinging on the PC population. We included excitatory recurrent connections among PCs, connections from PCs to INHs and from INHs to PCs, and only PCs received G inputs. Values for synaptic weights are unitless and specify the postsynaptic conductance change caused by a single presynaptic spike, relative to the leak conductance.

Recurrent excitatory connectivity

Recurrent excitatory synaptic weights were assigned such that neurons with similar position tuning had stronger connectivity, with the connectivity between two neurons i and j with peak tunings separated by d_{ij} given by

$$w_{ij} = w^{PC,PC} \exp\left[-d_{ij}^2/(2\lambda^{PC,PC2})\right].$$

This was motivated by past modeling studies (Káli and Dayan, 2000; Solstad, Yousif, and Sejnowski, 2014) and by analyses of correlated hippocampal activity in the absence of sensory input (Giusti, et al., 2015). Note that in our network architecture, all recurrent connections between position-tuned cells are reciprocal, due to the symmetry of the distance function d_{ij}. All w_{ij} below a minimum value $w_{min}^{PC,PC}$ = 0.1 were set to 0.

Inhibitory connectivity

We included an INH population 1/10 the size of the PC population (300 for ***Figure 2-3***, 1200 for ***Figure 4***). Each PC projected to each INH with probability 0.5 and weight $w^{INH,PC}$ = 0.03. For ***Figure 2-3***, each INH projected to each PC with probability 0.5 and weight $w^{PC,INH}$ = 0.02 (except where $w^{PC,INH}$ was varied in the parameter sweeps in ***Figure 3***). For ***Figure 4*** we decreased $w^{PC,INH}$ to 0.002 to account for the larger inhibitory pool.

Gating input connectivity

Synaptic weights w_i^G on gating inputs G were initially uniform across all position-tuned neurons, and each neuron received an independent instantiation of an upstream G spike train, generated from a Poisson-distributed point process with constant rate r^G. While effective w_i^G varied as a function of each neuron's trajectory-dependent LTP-IE, r^G was identical across all neurons regardless of tuning or LTP-IE status.

Long-term potentiation of intrinsic excitability (LTP-IE) of gating inputs

Although LTP-IE is thought to result from changes in dendritic conductances (Hyun, et al., 2015), we modeled it as an effective synaptic weight change (since conductance changes can be absorbed into connection weights according to the equations underlying conductance-based synaptic dynamics), such that initial weights w_i^G were scaled by a factor σ_i as a function of neuron i's position along the initial trajectory. For computational efficiency we did not model network activity during the initial trajectory, but instead directly calculated the σ_i expected to result from the maximum firing rate r^i_{max} = max(r^i(x, y) | (x, y) ϵ trajectory). From r^i_{max} we computed σ_i according to the following equation (***Figure 1B***):

$$\sigma_i = 1 + \frac{\sigma_{max} - 1}{1 + \exp\left[-\beta_\sigma (r^i_{max} - r_\sigma)\right]}$$

with r_σ = 10 Hz and β_σ = 1. Thus, neurons with position-tuning peaks close to the original trajectory had a σ_i near σ_{max} whereas neurons far away from the trajectory had σ_i near unity.

Sequence replay

We fully simulated the replay epoch only, during which all neurons in the recurrent network received independent stochastic gating input G at constant rate r^G. Replay events arose spontaneously, but in ***Figure 4*** we triggered sequence propagation with a short current pulse to neurons in one of the sensory regions; trigger parameters were unimportant as long as they elicited significant spiking activity in the neurons at the start of the sequence. In ***Figure 2-figure supplement 1F***, replayed positions were decoded during by taking the median peak-tuning of neurons that spiked within 5 ms windows spanning the replay event (decoding was only performed in a window if there were at least 5 spikes in that window).

Replay event detection and speed measurements

To detect spontaneous replay events automatically, we first smoothed the cell-averaged PC firing rate with a Gaussian filter with a 2 ms standard deviation. Next, all timepoints where this smoothed firing rate exceeded 0.5 Hz were marked. Continuous events above 0.5 Hz lasting longer than 30 ms were then labeled as replay events, with two replay events joined into one if the duration of the gap between them was less than 10 ms.

We measured replay speed for an event by first computing the linear position of each PC's place field center along the Z-shaped trajectory shown in ***Figure 2E.*** We then plotted these linear positions of all PCs that spiked during the middle 80% of the replay epoch vs the times the spikes occurred (we removed the first and last 10% to avoid edge effects). Finally, we calculated speed by computing the slope of the best-fit linear regression (minimal least squared error) to the linear-position-vs-spike-time plot.

Spiking network simulations

Single-neuron LTP-IE simulation (***Figure 1***)

LTP-IE-level-dependent voltage distributions (***Figure 1D, E***) were computed using 15-second simulations. In measuring ***Figure 1F-G***, current pulses were 10 ms in duration, and spiking probability was measured within this window. Current-evoked spike probabilities (***Figure 1G***) were computed using 125-second simulations, with pulses presented every 250 ms.

Replay in networks of neurons (***Figure 2-3***)

In replay simulations we used 3000 neurons, with place fields distributed on an approximate lattice over a 2 m x 2 m environment. Gating inputs were provided continually, as our simulation represented the awake, quiescent state during which replay is thought to occur. Gating inputs were random homogeneous Poisson point processes, independent to each neuron, with rate r^G that was the same across all neurons.

*Replay dynamics parameter sweep (**Figure 3**)*

For each parameter set we ran ten 10-second simulations starting with different random number generator seeds. We calculated spontaneous event frequency by counting the number of spontaneous replay events arising during each simulation, dividing by the 10-second simulation duration, and averaging over the 10 trials.

To explore the dependence of spontaneous replay event frequency on our network parameters, we first varied the magnitude factor $w^{PC,PC}$ of the recurrent excitatory weight profile and the inhibitory-to-PC connectivity strength $w^{PC,INH}$ (***Figure 3A***), while holding all other parameters fixed. We next varied $w^{PC,PC}$ and $\lambda^{PC,PC}$, the recurrent excitatory connectivity length scale, respectively, while holding all other parameters fixed (***Figure 3B***). Finally, we varied σ_{max}, the maximum LTP-IE level, and r^G, the gate-input rate, while holding all other parameters fixed (***Figure 3C***). Fixed parameter values are given in the table below.

Virtual replay speed measurements and parameter sweep

To explore the dependence of virtual replay speed on network parameters we varied $w^{PC,PC}$ and $\lambda^{PC,PC}$ while measuring replay speed in the following way. For all parameter sets in ***Figure 3A*** that did not exhibit "blowup" behavior (i.e. all parameter sets where replay events were roughly confined to the LTP-IE profile), we ran up to twenty 600 ms trial simulations in which we injected a short current pulse to the initial end of the trajectory to trigger replay. Simulations ended if 20 trials were reached or 10 successfully triggered replay events were observed. To calculate the virtual replay in a given event we first truncated the initial and final 10% of the event to avoid edge effects. We then collected all neurons that spiked during the event and fit a line using least-squares regression to the linear positions of preferred (x, y) of each neuron along the trajectory vs the times those neurons emitted spikes, and the slope of this line was taken to be the virtual replay speed. We only included replay events lasting between 50 and 400 ms for the speed computation, and events exhibiting a

replay speed of more than 10 m/s (events falling outside this range tended not to correspond to unidirectional propagation of the spike sequence along the LTP-IE-defined trajectory). To make the plot shown in ***Figure 3G*** we computed the average virtual replay speeds for all parameter sets exhibiting at least 5 successful replay events.

Power spectral density analysis

Power spectral densities shown in ***Figure 3H*** were computed by running a single simulation for each of two different parameter sets, collecting all replay events, measuring the power spectrum of each event, and averaging over events.

*LTP-IE-based stimulus-response mappings (**Figure 4**)*

For this simulation, in order to construct sufficiently non-interfering LTP-IE-defined paths, we constructed a model network of 12000 neurons with neural place fields ranging from -2 to +2 m in both the x- and y-dimensions and let two clusters of radius of 0.25 m, centered at (-1, 1) and (-1, -1), represent two stimuli S1 and S2, respectively, and two additional clusters, centered at (1, 1) and (1, -1), encode motor outputs M1 and M2, respectively. Readout units summed activity in any of the four clusters at every timestep. Sequence reactivation was triggered by depolarizing current injections into cells contained in the S1 or S2 clusters.

All parameters for spiking network simulations

We used the following values for all simulation parameters unless otherwise noted (e.g. when values were varied in parameter sweeps [***Figure 3***]):

Symbol	**Definition**	**Value**	**Symbol**	**Definition**	**Value**
τ_m^{PC}	PC membrane time constant	50 ms	N^{PC}	Number of PCs in recurrent network	3000 (***Fig 2-3***) 12000 (***Fig 4***)
E^{PC}_{leak}	PC leak potential	-68 mV	N^{INH}	Number of INHs in recurrent network	300 (***Fig 2-3***) 1200 (***Fig 4***)
v^{PC}_{th}	PC spike threshold	-36 mV	$w^{PC,G}$	Gating input weight without LTP-IE	0.8216 (***Fig 2-3***) 0.5 (***Fig 4***)
τ^{PC}_r	PC refractory period	8 ms	$w^{PC,PC}$	Excitatory recurrent weight scale factor	2.6
τ_m^{INH}	INH membrane time constant	5 ms	$w_{min}^{PC,PC}$	Min nonzero recurrent weight	.1
E^{PC}_{leak}	INH leak potential	-60 mV	$w^{PC,INH}$	INH→PC connection weight	0.02 (***Fig 2-3***) 0.002 (***Fig 4***)
v^{PC}_{th}	INH spike threshold	-50 mV	$w^{INH,PC}$	PC→INH connection weight	.03
τ^{PC}_r	INH refractory	2 ms	$p^{PC,INH}$	INH→PC connection	0.5

	period			probability	
E_E	Excitatory synaptic reversal potential	0 mV	$p^{INH,PC}$	PC→INH connection probability	0.5
τ_E	Excitatory synaptic time constant	2 ms	r^G	Gating input firing rate	125 Hz (***Fig 2-3***) 200 Hz (***Fig 4***)
E_I	Inhibitory synaptic reversal potential	-80 mV	$\lambda^{PC,PC}$	Excitatory recurrent connectivity length scale	.053 m
τ_I	Inhibitory synaptic time constant	2 ms	σ_{max}	Max LTP-IE value	2 (***Fig 2-3***) 1.84 (***Fig 4***)
r_{max}	Max position-driven firing rate	20 Hz	r_σ	Threshold firing rate for LTP-IE	10 Hz
λ_{PL}	Position-tuning length constant	0.15 m	β_σ	Scale factor for LTP-IE onset	1
			ΔT	Simulation time step	0.5 ms

Reduced model dynamics

We used a reduced network model and a reduced model of LTP-IE to explore the algorithmic potential of LTP-IE separately from its specific biophysical properties.

The reduced network model operates in discrete timesteps, with each neuron's "voltage" v^i at time t given by the sum of all of its inputs:

$$v^i(t) = \sum_j w_{ij} s_j(t-1) + i_{ext}(t) + i_g \sigma^i$$

where w_{ij} is the synaptic connection weight from neuron j onto neuron i, $s_j(t - 1)$ is 1 if neuron j spiked at t - 1 and 0 otherwise, $i_{ext}(t)$ is the external input current to neuron i at time t, i_g is a constant "gating" current, and σ^i is neuron i's LTP-IE level (ranging from 1 to 2, depending on the neuron). If a neuron's voltage at time t exceeded a threshold voltage v_{th}, we let that neuron spike, with the following caveat: a maximum of n_{max} neurons were allowed to spike at any individual timestep, and if more than n_{max} neurons' voltages exceeded v_{th}, we let only the n_{max} neurons with the highest voltages spike. Neurons that spiked subsequently entered a refractory period that prevented them from spiking for τ_r subsequent timesteps. LTP-IE values σ^i were static throughout each simulation and pre-assigned according to rules that depended on the specific simulation in question.

Reduced model simulation of intersecting trajectory replay

Connectivity

We first specified tuning parameters for each neuron with regard to position and orientation in the environment. Each neuron was assigned a position preference in (x, y) space, tiling a (2 m x 2 m) environment; as well as a preferred θ, randomly selected from one of 8 possible angles tiling 2π radians: 0, π/4, -π/4, ..., such that each window of 8 neurons in a set of windows tiling a row of neurons contained all 8 angles. (E.g. the first through eighth neurons with a given preferred y would each have a distinct preferred θ, similarly for the ninth through 16th neurons, and so on.) Neurons were subsequently connected according to their preferred tuning, with the (symmetric) connection weight *w* between two neurons separated by Δx, Δy, and Δθ in their preferred tuning space given by

$$w(\Delta x, \Delta y, \Delta\theta) = \exp\left(-\frac{\Delta x^2 + \Delta y^2}{2\lambda_{xy}^2}\right)\exp\left(-\frac{\Delta\theta^2}{2\lambda_\theta^2}\right)$$

with λ_{xy} and λ_θ specifying connectivity length constants. (Note Δθ is the absolute distance between two angles calculated on the unit circle.)

LTP-IE profile

To assign LTP-IE levels to neurons such that accurate replay would be achieved, we first generated the trajectory through the environment shown in ***Figure 5A, inset***. Neurons with preferred tuning close to the trajectory in (x, y, θ)-space were assigned high LTP-IE levels (σ ~ 2), and neurons far away from the trajectory were assigned low LTP-IE levels (σ ~ 1), according to the following equation:

$$\sigma(x, y, \theta) = 1 + \frac{1}{1 + \exp[-\beta(z - z_0)]}$$

$$z(x, y, \theta) = \exp\left[-.5\left(\frac{\Delta x^2}{\sigma_x^2} + \frac{\Delta y^2}{\sigma_y^2} + \frac{\Delta\theta}{\sigma_\theta^2}\right)\right]$$

where Δx, Δy, and Δθ in this case were the distance between a neuron's preferred x, y, and θ, respectively, and the point on the trajectory that yielded the highest z. σ_x, σ_y, and σ_θ were length constants determining how similar a neurons tuning peak had to be to an (x, y, θ) point on the trajectory to induce LTP-IE.

Trigger

In the reduced model we explicitly triggered replay by injecting an external current pulse with a squared exponential profile focused at either the initial (-1 m, 0.25 m) or final end (-.25 m, 1 m) to evoke forward or reverse replay, respectively. Exact trigger parameters were unimportant as long as they evoked spiking in approximately the correct location.

Parameters

We used the following parameter values for the reduced model simulation of intersecting trajectory replay:

Symbol	Definition	Value	Symbol	Definition	Value
N	Total number of neurons	4096 (64 x 64)	n_{max}	Max number of active neurons at a timestep	10
N^X	Number of neurons in each row of environment	64	λ_{xy}	Connectivity length scale	0.25 m
N^Y	Number of neurons in each column of environment	64	λ_θ	Connectivity orientation-length scale	π/7 radians
N^θ	Number of equally spaced orientations	8	σ_X , σ_Y	LTP-IE length scales	0.2 m
v_{th}	Spike threshold	11	σ_θ	LTP-IE orientation-length scale	π/16 radians
i_G	Gate input	5	β	LTP-IE sigmoid steepness	20
τ_r	Refractory period	6 timesteps	T	Number of timesteps in simulation	50 timesteps

Reduced model simulation of non-spatial mapping

Connectivity

We included five groups of neurons in this simulation: S1, S2, M1, M2, and B. S1 and S2 represented sensory ensembles, each composed of N_S neurons; and M1 and M2 represented motor ensembles, each composed of N_M neurons. B represented a "switchboard" ensembles of neurons, akin to Swan and Wyble's "binding pool" (Swan and Wyble, 2014), comprising N_B neurons. Each motor ensemble, with connection weight w^{MM}.

N^{BS} randomly selected switchboard received connections from S1, and another potentially overlapping N^{BS} randomly selected switchboard neurons received connections from S2. If a switchboard neuron was included in the group receiving connections from S1, n^{BS} neurons from S1 were selected to send connections to that switchboard neuron; similarly for S2. Sensory-to-switchboard connections had synaptic weights w^{BS}.

N^{MB} randomly selected switchboard neurons sent connections to M1, and another potentially overlapping N^{MB} randomly selected switchboard neurons sent connections to M2. If a switchboard neuron was included in the group sending connections to M1, n^{MB} neurons in M1 were selected to receive connections from that switchboard neuron; similarly for M2. Switchboard-to-motor connections had synaptic weight w^{MB}.

There were no recurrent connections within either sensory ensemble or within the switchboard ensemble. These connections were omitted since our goal was to demonstrate a proof-of-concept nonspatial-mapping computation through LTP-IE. Exploring the effects of sensory and switchboard recurrence in a more biophysically accurate model may yield fruitful insights into LTP-IE-based computations, however.

LTP-IE profile

To encode the mapping (S1→M1, S2→M2) we assigned an LTP-IE value of $\sigma = 2$ to all switchboard neurons that either (1) received connections from S1 and sent connections to M1 or (2) received connections from S2 and sent connections to M2. All other neurons were assigned an LTP-IE value of $\sigma = 1$. To encode the mapping (S1→M2, S2→M1) we assigned an LTP-IE value of $\sigma = 2$ to all switchboard neurons that either (1) received connections from S1 and sent connections to M2 or (2) received connections from S2 and sent connections to M1. All other neurons were assigned an LTP-IE value of $\sigma = 1$.

Trigger

We triggered recall of an association encoded in the switchboard LTP-IE profile by injecting an external current pulse at timestep 2 into all the neurons in either S1 or all the neurons in S2.

Parameters

Symbol	**Definition**	**Value**	**Symbol**	**Definition**	**Value**
N	Total number of neurons	2400	N^{BS}	Number of switchboard neurons receiving sensory projections	200
N_S	Number of neurons in each sensory ensemble	100	N^{MB}	Number of switchboard neurons projecting to one motor ensemble	200
N_M	Number of neurons in each motor ensemble	100	n^{BS}	Number of sensory neurons projecting to each switchboard neuron receiving sensory input	10
N_B	Number of neurons in switchboard ensemble	2000	n^{MB}	Number of motor neurons receiving projections from each switchboard neuron projecting to motor ensembles	25
v_{th}	Spike threshold	10	w^{MM}	Recurrent	0.35

				connection strength within each motor ensemble	
i_G	Gate input	3	w^{BS}	Connection strength from sensory ensemble to switchboard neurons	0.88
τ_r	Refractory period	5 timesteps	w^{MB}	Connection strength from switchboard to motor ensemble neurons	2.33
n_{max}	Max number of neurons active at a timestep	50	T	Number of timesteps in simulation	5 timesteps

Code

All code for this work was written in Python 3 and is available at https://github.com/rkp8000/seq_speak.

Acknowledgments

We would like to acknowledge Stefano Recanatesi, Alison Duffy, Guillaume Lajoie, Ben Lansdell, Yoni Browning, and Kenneth Latimer for helpful discussions regarding this work. We would also like to acknowledge the MONA2: Modeling Neuronal Activity meeting for providing an opportunity for feedback on an earlier version of this work.

Funding

This work was supported by National Institutes of Health (https://www.nih.gov) grants NIH R01DC013693, NIH R01NS104925; The Simons Foundation Collaboration for the Global Brain; The University of Washington Computational Neuroscience Training Grant; and The Washington Research Foundation University of Washington Institute for Neuroengineering. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests

The authors have declared that no competing interests exist.

Supplementary figures

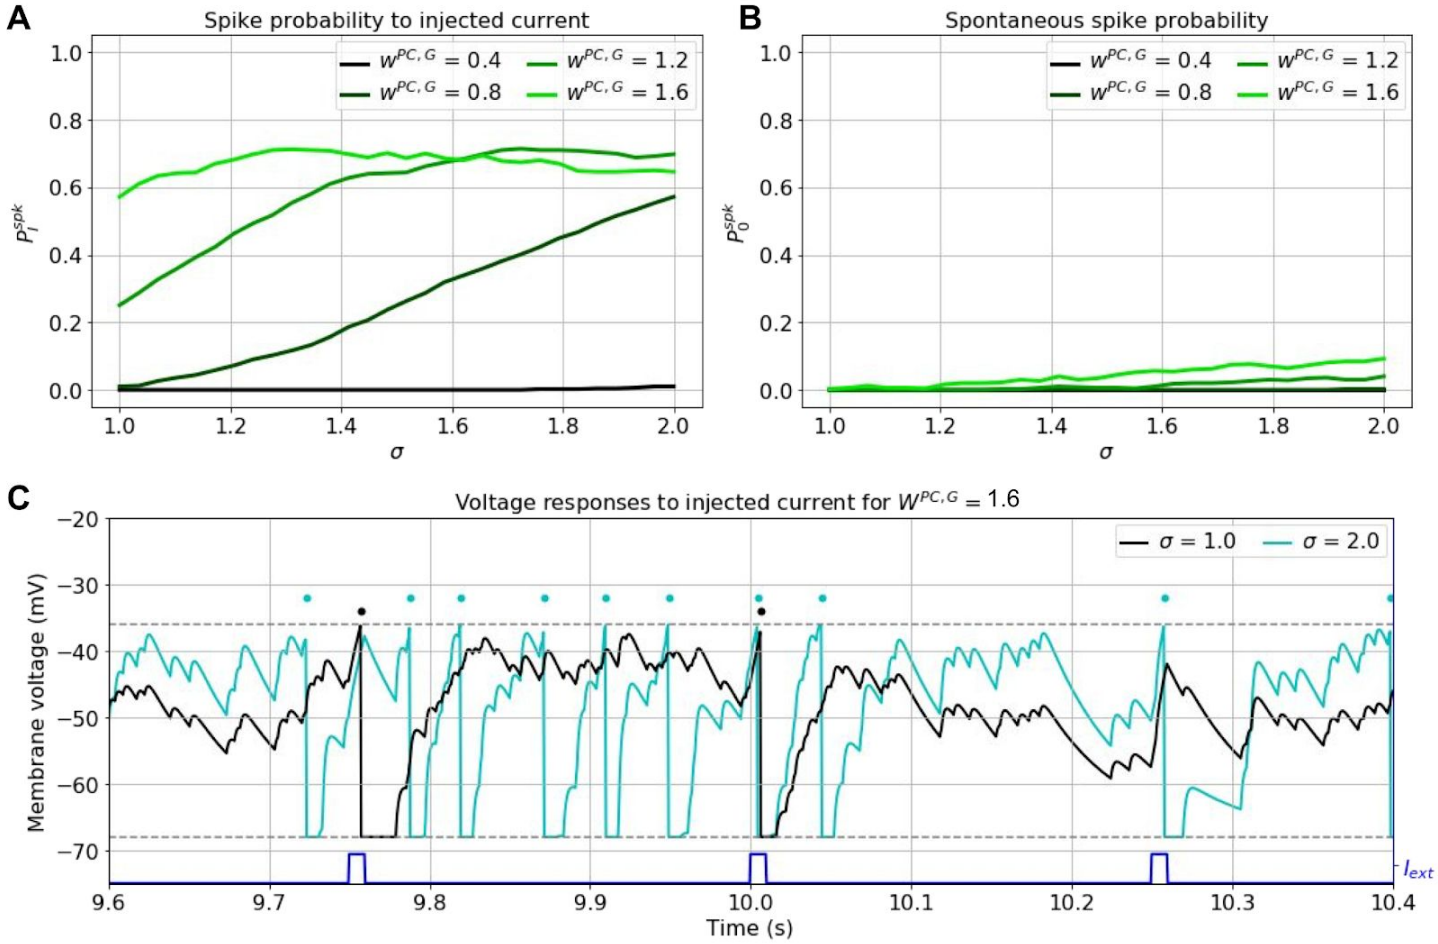

Figure 1-figure supplement 1: Effects of current pulses on spiking under LTP-IE. *A. Probability of current-evoked spiking for different $w^{PC,G}$ and σ. B. Spontaneous spiking probability for different different $w^{PC,G}$ and σ. C. Example voltage traces for neurons with two different σ given large $w^{PC,G}$. Same format as* ***Figure 1F****.*

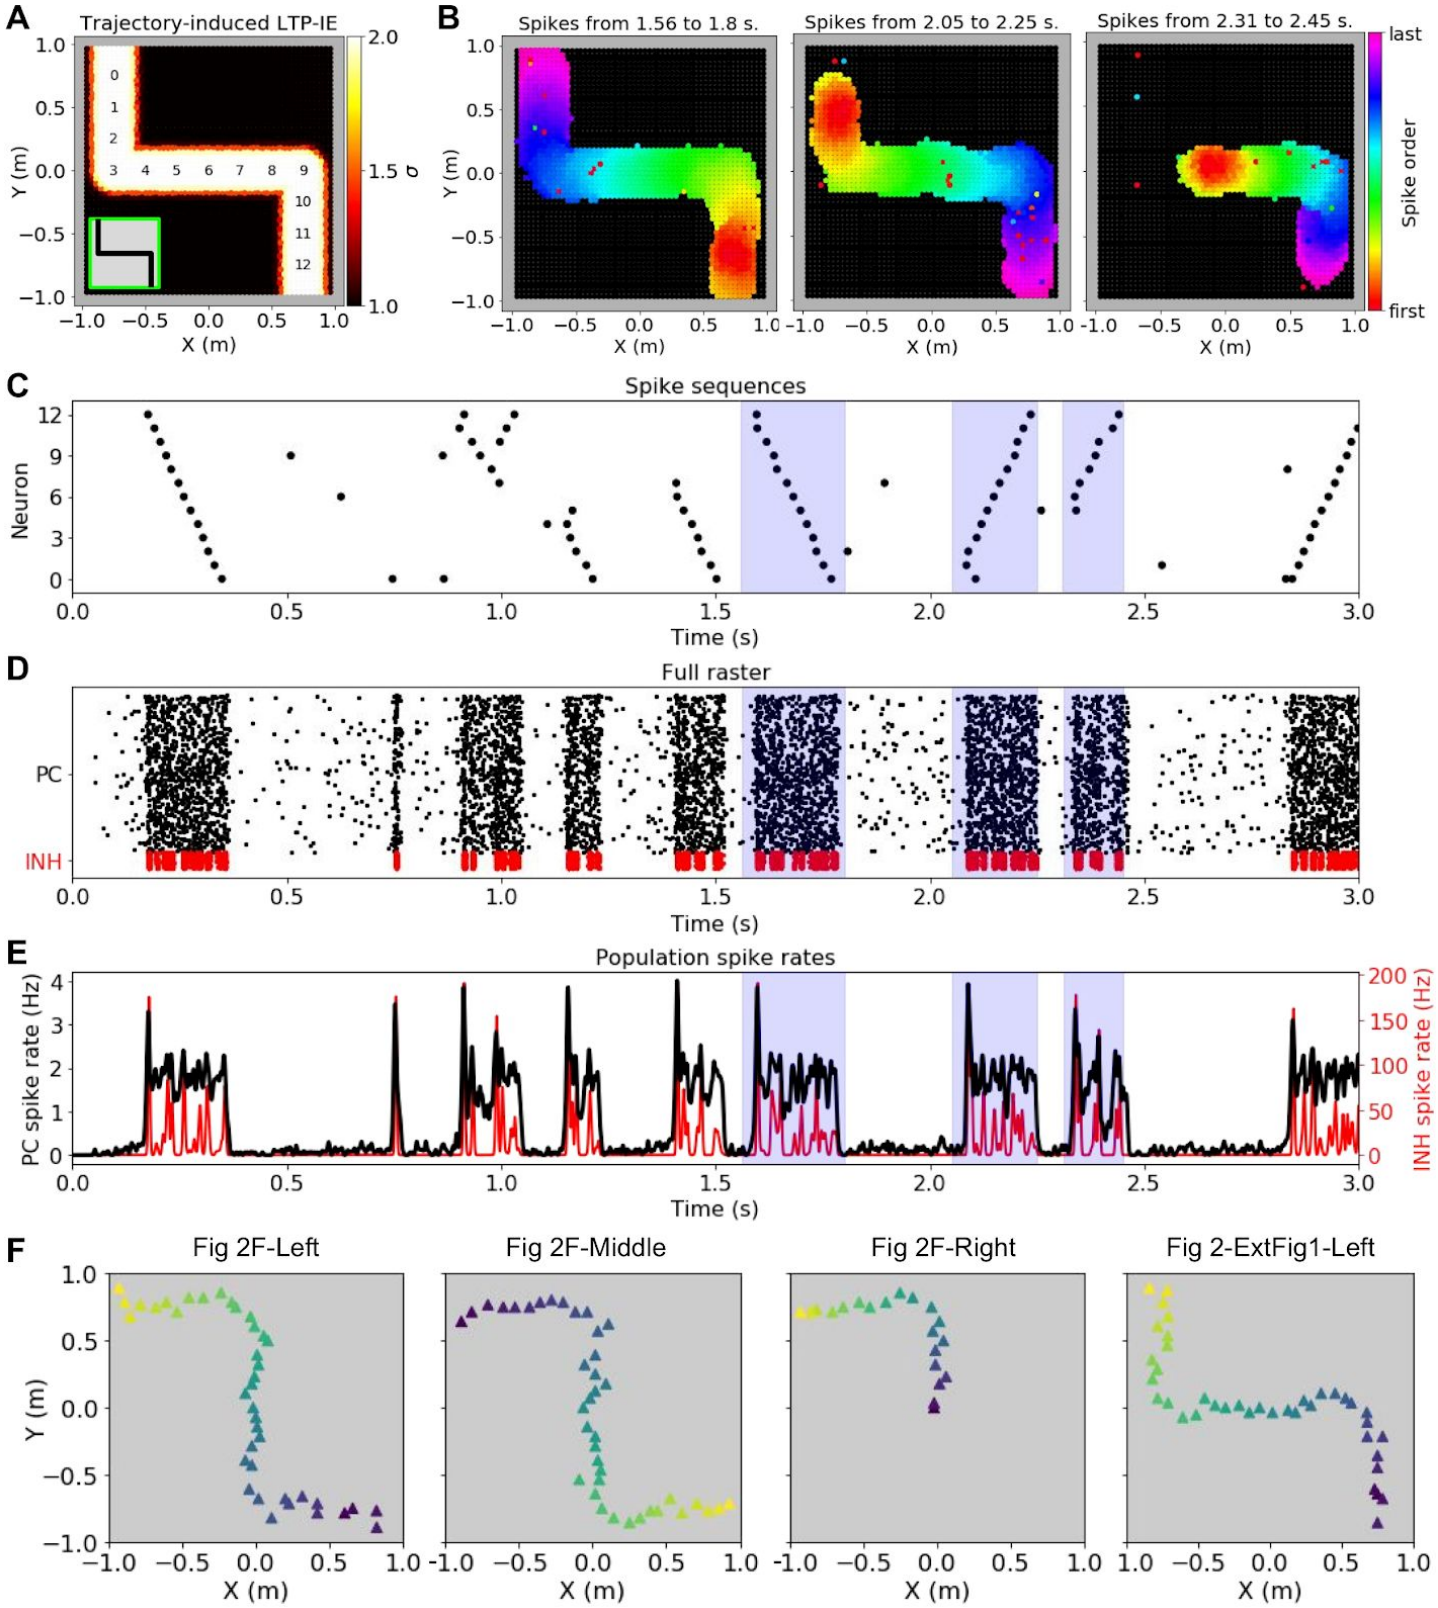

Figure 2-figure supplement 1: *LTP-IE-based replay for an alternative trajectory/sequence to that shown in Figure 2. A-E: As in Figure 2E-H, but for the trajectory shown in A (inset). F: Trajectories decoded from the replay events indicated above each panel (dark blue indicates start and yellow indicates end).*

References

Abeles, Moshe. *Local cortical circuits: An electrophysiological study*. Vol. 6. Springer Science & Business Media, 2012.

Abeles, M., et al. "Synchronization in neuronal transmission and its importance for information processing." *Progress in brain research*. Vol. 102. Elsevier, 1994. 395-404.

Bakker, Arnold, et al. "Pattern separation in the human hippocampal CA3 and dentate gyrus." *Science* 319.5870 (2008): 1640-1642.

Beggs, John M., and Dietmar Plenz. "Neuronal avalanches in neocortical circuits." *Journal of neuroscience* 23.35 (2003): 11167-11177.

Beggs, John M., and Dietmar Plenz. "Neuronal avalanches are diverse and precise activity patterns that are stable for many hours in cortical slice cultures." *Journal of neuroscience* 24.22 (2004): 5216-5229.

Bi, Guo-qiang, and Mu-ming Poo. "Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type." *Journal of Neuroscience* 18.24 (1998): 10464-10472.

Botvinick, Matthew, and Takamitsu Watanabe. "From numerosity to ordinal rank: a gain-field model of serial order representation in cortical working memory." *Journal of Neuroscience* 27.32 (2007): 8636-8642.

Brunel, Nicolas. "Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons." *Journal of Computational Neuroscience* 8.3 (2000): 183-208.

Buzsáki, György. "Hippocampal sharp waves: their origin and significance." *Brain research* 398.2 (1986): 242-252.

Campanac, Emilie, and Dominique Debanne. "Spike timing-dependent plasticity: a learning rule for dendritic integration in rat CA1 pyramidal neurons." *The Journal of physiology* 586.3 (2008): 779-793.

Caporale, Natalia, and Yang Dan. "Spike timing–dependent plasticity: a Hebbian learning rule." *Annu. Rev. Neurosci.* 31 (2008): 25-46.

Carr, Margaret F., Shantanu P. Jadhav, and Loren M. Frank. "Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval." *Nature Neuroscience* 14.2 (2011): 147.

Carrillo-Reid, Luis, et al. "Imprinting and recalling cortical ensembles." *Science* 353.6300 (2016): 691-694.

Chenkov, Nikolay, Henning Sprekeler, and Richard Kempter. "Memory replay in balanced recurrent networks." *PLoS Computational Biology* 13.1 (2017): e1005359.

Crowe, David A., Bruno B. Averbeck, and Matthew V. Chafee. "Rapid sequences of population activity patterns dynamically encode task-critical spatial information in parietal cortex." *Journal of Neuroscience* 30.35 (2010): 11640-11653.

Csicsvari, Jozsef, et al. "Fast network oscillations in the hippocampal CA1 region of the behaving rat." *J Neurosci* 19.RC20 (1999): 1-4.

Csicsvari, Jozsef, et al. "Ensemble patterns of hippocampal CA3-CA1 neurons during sharp wave–associated population events." *Neuron* 28.2 (2000): 585-594.

Curto, Carina, Anda Degeratu, and Vladimir Itskov. "Flexible memory networks." *Bulletin of Mathematical Biology* 74.3 (2012): 590-614.

Dan, Yang, and Mu-ming Poo. "Spike timing-dependent plasticity of neural circuits." *Neuron* 44.1 (2004): 23-30.

Daoudal, Gaël, Yasuhiro Hanada, and Dominique Debanne. "Bidirectional plasticity of excitatory postsynaptic potential (EPSP)-spike coupling in CA1 hippocampal pyramidal neurons." *Proceedings of the National Academy of Sciences* 99.22 (2002): 14512-14517.

Daoudal, Gaël, and Dominique Debanne. "Long-term plasticity of intrinsic excitability: learning rules and mechanisms." *Learning & memory* 10.6 (2003): 456-465.

Davidson, Thomas J., Fabian Kloosterman, and Matthew A. Wilson. "Hippocampal replay of extended experience." *Neuron* 63.4 (2009): 497-507.

Debanne, Dominique, and Mu-Ming Poo. "Spike-timing dependent plasticity beyond synapse-pre-and post-synaptic plasticity of intrinsic neuronal excitability." *Frontiers in synaptic neuroscience* 2 (2010): 21.

Diehl, Peter U., and Matthew Cook. "Unsupervised learning of digit recognition using spike-timing-dependent plasticity." *Frontiers in Computational Neuroscience* 9 (2015): 99.

Eagleman, Sarah L., and Valentin Dragoi. "Image sequence reactivation in awake V4 networks." *Proceedings of the National Academy of Sciences* 109.47 (2012): 19450-19455.

Ego-Stengel, Valérie, and Matthew A. Wilson. "Disruption of ripple‐associated hippocampal activity during rest impairs spatial learning in the rat." *Hippocampus* 20.1 (2010): 1-10.

Egorov, Alexei V., et al. "Graded persistent activity in entorhinal cortex neurons." *Nature* 420.6912 (2002): 173.

Fiete, Ila R., et al. "Spike-time-dependent plasticity and heterosynaptic competition organize networks to produce long scale-free sequences of neural activity." *Neuron* 65.4 (2010): 563-576.

Foster, David J., and Matthew A. Wilson. "Reverse replay of behavioural sequences in hippocampal place cells during the awake state." *Nature* 440.7084 (2006): 680.

Foster, David J., and James J. Knierim. "Sequence learning and the role of the hippocampus in rodent navigation." *Current opinion in neurobiology* 22.2 (2012): 294-300.

Fuentemilla, Lluís, et al. "Theta-coupled periodic replay in working memory." *Current Biology* 20.7 (2010): 606-612.

Gerstner, Wulfram, et al. *Neuronal dynamics: From single neurons to networks and models of cognition*. Cambridge University Press, 2014.

Gilson, Matthieu, Anthony Burkitt, and Leo J. Van Hemmen. "STDP in recurrent neuronal networks." *Frontiers in Computational Neuroscience* 4 (2010): 23.

Giusti, Chad, et al. "Clique topology reveals intrinsic geometric structure in neural correlations." *Proceedings of the National Academy of Sciences* 112.44 (2015): 13455-13460.

Griffith, J. S., and G. Horn. "An analysis of spontaneous impulse activity of units in the striate cortex of unrestrained cats." *The Journal of Physiology* 186.3 (1966): 516-534.

Gupta, Anoopum S., et al. "Hippocampal replay is not a simple function of experience." *Neuron* 65.5 (2010): 695-705.

Haga, Tatsuya, and Tomoki Fukai. "Recurrent network model for learning goal-directed sequences through reverse replay." *eLife* 7 (2018): e34171.

Hahnloser, Richard HR, Alexay A. Kozhevnikov, and Michale S. Fee. "An ultra-sparse code underliesthe generation of neural sequences in a songbird." *Nature* 419.6902 (2002): 65.

Hartley, Tom, et al. "Space in the brain: how the hippocampal formation supports spatial cognition." *Philosophical Transactions of the Royal Society B: Biological Sciences* 369.1635 (2014): 20120510.

Harvey, Christopher D., Philip Coen, and David W. Tank. "Choice-specific sequences in parietal cortex during a virtual-navigation decision task." *Nature* 484.7392 (2012): 62.

Hattan, David, et al. "Tyrosine phosphorylation of Kv1. 2 modulates its interaction with the actin-binding protein cortactin." *Journal of Biological Chemistry* 277.41 (2002): 38596-38606.

Herrmann, M., J. A. Hertz, and A. Prügel-Bennett. "Analysis of synfire chains." *Network: computation in neural systems* 6.3 (1995): 403-414.

Hyun, Jung Ho, et al. "Activity-dependent downregulation of D-type K+ channel subunit Kv1. 2 in rat hippocampal CA3 pyramidal neurons." *The Journal of Physiology* 591.22 (2013): 5525-5540.

Hyun, Jung Ho, et al. "Kv1. 2 mediates heterosynaptic modulation of direct cortical synaptic inputs in CA3 pyramidal cells." *The Journal of Physiology* 593.16 (2015): 3617-3643.

Ikegaya, Yuji, et al. "Synfire chains and cortical songs: temporal modules of cortical activity." *Science* 304.5670 (2004): 559-564.

Isaac, John TR, et al. "Hippocampal place cell firing patterns can induce long-term synaptic plasticity in vitro." *Journal of Neuroscience* 29.21 (2009): 6840-6850.

Jadhav, Shantanu P., et al. "Awake hippocampal sharp-wave ripples support spatial memory." *Science* 336.6087 (2012): 1454-1458.

Káli, Szabolcs, and Peter Dayan. "The involvement of recurrent connections in area CA3 in establishing the properties of place fields: a model." *Journal of Neuroscience* 20.19 (2000): 7463-7477.

Karlsson, Mattias P., and Loren M. Frank. "Awake replay of remote experiences in the hippocampus." *Nature Neuroscience* 12.7 (2009): 913.

Klampfl, Stefan, and Wolfgang Maass. "Emergence of dynamic memory traces in cortical microcircuit models through STDP." *Journal of Neuroscience* 33.28 (2013): 11515-11529.

Ko, Ho, et al. "Functional specificity of local synaptic connections in neocortical networks." *Nature* 473.7345 (2011): 87.

Koulakov, Alexei A., Tomáš Hromádka, and Anthony M. Zador. "Correlated connectivity and the distribution of firing rates in the neocortex." *Journal of Neuroscience* 29.12 (2009): 3685-3694.

Laje, Rodrigo, and Dean V. Buonomano. "Robust timing and motor patterns by taming chaos in recurrent neural networks." *Nature Neuroscience* 16.7 (2013): 925.

Lee, Tyler P., and Dean V. Buonomano. "Unsupervised formation of vocalization-sensitive neurons: a cortical model based on short-term and homeostatic plasticity." *Neural Computation* 24.10 (2012): 2579-2603.

Leutgeb, S., K. E. Ragozzino, and S. J. Y. Mizumori. "Convergence of head direction and place information in the CA1 region of hippocampus." *Neuroscience* 100.1 (2000): 11-19.

Leutgeb, Jill K., et al. "Pattern separation in the dentate gyrus and CA3 of the hippocampus." *Science* 315.5814 (2007): 961-966.

Lisman, John E., and Marco A. Idiart. "Storage of 7+/-2 short-term memories in oscillatory subcycles." *Science* 267.5203 (1995): 1512-1515.

Lisman, John E. "Relating hippocampal circuitry to function: recall of memory sequences by reciprocal dentate–CA3 interactions." *Neuron* 22.2 (1999): 233-242.

Lisman, John. "Working memory: the importance of theta and gamma oscillations." *Current Biology* 20.11 (2010): R490-R492.

Luczak, Artur, et al. "Sequential structure of neocortical spontaneous activity in vivo." *Proceedings of the National Academy of Sciences* 104.1 (2007): 347-352.

Mamad, Omar, et al. "Medial septum regulates the hippocampal spatial representation." *Frontiers in Behavioral Neuroscience* 9 (2015): 166.

Markram, Henry, et al. "Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs." *Science* 275.5297 (1997): 213-215.

Masquelier, Timothée, Rudy Guyonneau, and Simon J. Thorpe. "Spike timing dependent plasticity finds the start of repeating patterns in continuous spike trains." *PloS One* 3.1 (2008): e1377.

McCloskey, Michael, and Neal J. Cohen. "Catastrophic interference in connectionist networks: The sequential learning problem." *Psychology of Learning and Motivation*. Vol. 24. Academic Press, 1989. 109-165.

Mizuseki, Kenji, et al. "Activity dynamics and behavioral correlates of CA3 and CA1 hippocampal pyramidal neurons." *Hippocampus* 22.8 (2012): 1659-1680.

Mizuseki, Kenji, and György Buzsáki. "Preconfigured, skewed distribution of firing rates in the hippocampus and entorhinal cortex." *Cell Reports* 4.5 (2013): 1010-1021.

Molter, Colin, Naoyuki Sato, and Yoko Yamaguchi. "Reactivation of behavioral activity during sharp waves: a computational model for two stage hippocampal dynamics." *Hippocampus* 17.3 (2007): 201-209.

Mongillo, Gianluigi, Omri Barak, and Misha Tsodyks. "Synaptic theory of working memory." *Science* 319.5869 (2008): 1543-1546.

Moser, Edvard I., Emilio Kropff, and May-Britt Moser. "Place cells, grid cells, and the brain's spatial representation system." *Annu. Rev. Neurosci.* 31 (2008): 69-89.

Nicola, Wilten, and Claudia Clopath. "Supervised learning in spiking neural networks with FORCE training." *Nature Communications* 8.1 (2017): 2208.

Ohtsuki, Gen, and Christian Hansel. "Synaptic potential and plasticity of an SK2 channel gate regulate spike burst activity in cerebellar Purkinje cells." *iScience* 1 (2018): 49-54.

O'Keefe, John. "A review of the hippocampal place cells." *Progress in Neurobiology* 13.4 (1979): 419-439.

O'Keefe, John, and Neil Burgess. "Dual phase and rate coding in hippocampal place cells: theoretical significance and relationship to entorhinal grid cells." *Hippocampus* 15.7 (2005): 853-866.

Ólafsdóttir, H. Freyja, Francis Carpenter, and Caswell Barry. "Task demands predict a dynamic switch in the content of awake hippocampal replay." *Neuron* 96.4 (2017): 925-935.

Ólafsdóttir, H. Freyja, Daniel Bush, and Caswell Barry. "The role of hippocampal replay in memory and planning." *Current Biology* 28.1 (2018): R37-R50.

Pastalkova, Eva, et al. "Internally generated cell assembly sequences in the rat hippocampus." *Science* 321.5894 (2008): 1322-1327.

Paz, Jeanne T., et al. "Multiple forms of activity-dependent intrinsic plasticity in layer V cortical neurones in vivo." The Journal of physiology 587.13 (2009): 3189-3205.

Pfeiffer, Brad E., and David J. Foster. "Hippocampal place-cell sequences depict future paths to remembered goals." *Nature* 497.7447 (2013): 74.

Pignatelli, Michele, et al. "Engram cell excitability state determines the efficacy of memory retrieval." *Neuron* 101.2 (2019): 274-284.

Plenz, Dietmar, and Tara C. Thiagarajan. "The organizing principles of neuronal avalanches: cell assemblies in the cortex?." *Trends in neurosciences* 30.3 (2007): 101-110.

Raastad, Morten, and Gordon MG Shepherd. "Single-axon action potentials in the rat hippocampal cortex." *The Journal of Physiology* 548.3 (2003): 745-752.

Raffone, Antonino, and Gezinus Wolters. "A cortical mechanism for binding in visual working memory." *Journal of Cognitive Neuroscience* 13.6 (2001): 766-785.

Rajan, Kanaka, Christopher D. Harvey, and David W. Tank. "Recurrent network models of sequence generation and memory." *Neuron* 90.1 (2016): 128-142.

Ramirez-Villegas, Juan F., et al. "Dissecting the synapse-and frequency-dependent network mechanisms of in vivo hippocampal sharp wave-ripples." *Neuron* 100.5 (2018): 1224-1240.

Rebola, Nelson, Mario Carta, and Christophe Mulle. "Operation and plasticity of hippocampal CA3 circuits: implications for memory encoding." *Nature Reviews Neuroscience* 18.4 (2017): 208.

Remington, Evan D., et al. "Flexible sensorimotor computations through rapid reconfiguration of cortical dynamics." *Neuron* 98.5 (2018): 1005-1019.

Renart, Alfonso, et al. "The asynchronous state in cortical circuits." *Science* 327.5965 (2010): 587-590.

Rezende, Danilo J., Daan Wierstra, and Wulfram Gerstner. "Variational learning for recurrent spiking networks." *Advances in Neural Information Processing Systems*. 2011.

Salinas, Emilio. "Fast remapping of sensory stimuli onto motor actions on the basis of contextual modulation." *Journal of Neuroscience* 24.5 (2004): 1113-1118.

Salinas, Emilio. "Context-dependent selection of visuomotor maps." *BMC neuroscience* 5.1 (2004): 47. [2004a]

Scarpetta, Silvia, and Antonio de Candia. "Neural avalanches at the critical point between replay and non-replay of spatiotemporal patterns." *PLoS One* 8.6 (2013): e64162.

Schiller, Daniela, et al. "Memory and space: towards an understanding of the cognitive map." *Journal of Neuroscience* 35.41 (2015): 13904-13911.

Schrader, Sven, et al. "Detecting synfire chain activity using massively parallel spike train recording." *Journal of neurophysiology* 100.4 (2008): 2165-2176.

Shew, Woodrow L., et al. "Neuronal avalanches imply maximum dynamic range in cortical networks at criticality." *Journal of neuroscience* 29.49 (2009): 15595-15600.

Sjöström, Per Jesper, Gina G. Turrigiano, and Sacha B. Nelson. "Rate, timing, and cooperativity jointly determine cortical synaptic plasticity." *Neuron* 32.6 (2001): 1149-1164.

Solstad, Trygve, Hosam N. Yousif, and Terrence J. Sejnowski. "Place cell rate remapping by CA3 recurrent collaterals." *PLoS Computational Biology* 10.6 (2014): e1003648.

Song, Sen, and Larry F. Abbott. "Cortical development and remapping through spike timing-dependent plasticity." *Neuron* 32.2 (2001): 339-350.

Sourdet, Valérie, et al. "Long-term enhancement of neuronal excitability and temporal fidelity mediated by metabotropic glutamate receptor subtype 5." *Journal of Neuroscience* 23.32 (2003): 10238-10248.

Sreenivasan, Kartik K., Clayton E. Curtis, and Mark D'Esposito. "Revisiting the role of persistent neural activity during working memory." *Trends in Cognitive Sciences* 18.2 (2014): 82-89.

Sussillo, David, and Larry F. Abbott. "Generating coherent patterns of activity from chaotic neural networks." *Neuron* 63.4 (2009): 544-557.

Swan, Garrett, and Brad Wyble. "The binding pool: A model of shared neural resources for distinct items in visual working memory." *Attention, Perception, & Psychophysics* 76.7 (2014): 2136-2157.

Titley, Heather K., Nicolas Brunel, and Christian Hansel. "Toward a neurocentric view of learning." *Neuron* 95.1 (2017): 19-32.

Vallar, Giuseppe, and Tim Shallice, eds. *Neuropsychological impairments of short-term memory*. Cambridge University Press, 2007.

Veliz-Cuba, Alan, et al. "Networks that learn the precise timing of event sequences." *Journal of Computational Neuroscience* 39.3 (2015): 235-254.

Wang, Huai-Xing, et al. "Coactivation and timing-dependent integration of synaptic potentiation and depression." *Nature Neuroscience* 8.2 (2005): 187.

Williams, Ronald J., and David Zipser. "A learning algorithm for continually running fully recurrent neural networks." *Neural Computation* 1.2 (1989): 270-280.

Yamamoto, Jun, and Susumu Tonegawa. "Direct medial entorhinal cortex input to hippocampal CA1 is crucial for extended quiet awake replay." *Neuron* 96.1 (2017): 217-227.

Yartsev, Michael M., and Nachum Ulanovsky. "Representation of three-dimensional space in the hippocampus of flying bats." *Science* 340.6130 (2013): 367-372.

Yger, Pierre, Marcel Stimberg, and Romain Brette. "Fast learning with weak synaptic plasticity." *Journal of Neuroscience* 35.39 (2015): 13351-13362.

Yoder, Ryan M., and Kevin CH Pang. "Involvement of GABAergic and cholinergic medial septal neurons in hippocampal theta rhythm." *Hippocampus* 15.3 (2005): 381-392.

Zhang, Wei, and David J. Linden. "The other side of the engram: experience-driven changes in neuronal intrinsic excitability." *Nature Reviews Neuroscience* 4.11 (2003): 885.

Zielinski, Mark C., Wenbo Tang, and Shantanu P. Jadhav. "The role of replay and theta sequences in mediating hippocampal‐prefrontal interactions for memory and cognition." *Hippocampus* (2017).