

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

Rich Pang^{1,2,3} and Adrienne 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, the adaptation of mental processing to changes in task demands, is thought to depend on biological neural networks' ability to rapidly modulate the dynamics governing how they process information. While extensive work has elucidated how network dynamics can be reshaped by slowly occurring structural changes, e.g. the gradual modification of recurrent synaptic patterns, much less is known about how dynamics might be reconfigured over faster timescales of seconds. One compelling example of rapid and selective modulation of network dynamics potentially involved in cognitive flexibility is observed in rodent hippocampus, where short bouts of exploratory behavior cause new activity sequences to preferentially “replay” during subsequent awake rest periods without continued sensory input. Fast mechanisms for selectively biasing sequential activity through networks, however, remain unknown. Using a spiking neural network model, we asked whether a simplified version of sequence replay could arise from three biophysically plausible components: recurrent, spatially organized connectivity; homogeneous, stochastic “gating” inputs; and rapid, activity-dependent scaling of gating input strengths, based on a phenomenon known as long-term potentiation of intrinsic excitability (LTP-IE). Indeed, these enabled both forward and reverse replay of flexible sequences reflecting recent behavior, despite unchanged recurrent weights. Specifically, activation-triggered LTP-IE “tags” neurons in the recurrent network by increasing their spiking probability when gating input is applied, and the sequential ordering of spikes is reconstructed by the existing recurrent connectivity. In a proof-of-concept demonstration, we also show how LTP-IE-based sequences can implement temporary stimulus-response mappings in a straightforward manner. These results elucidate a simple yet previously unexplored combination of biological mechanisms that converge in hippocampus and suffice for fast and flexible reconfiguration of sequential network dynamics, suggesting their potential role in cognitive flexibility over rapid timescales.

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, et al., 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, Hromadka, and Zador, 2009; Mizuseki and Buzsaki, 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, Averbeck, and Chafee, 2010; Harvey, Coen, and Tank, 2012). Functionally, firing sequences are thought to be involved in various cognitive processes, from short-term memory (Davidson, Kloosterman, and Wilson, 2009; Crowe, Averbeck, and Chafee, 2010) to decision-making (Harvey, Coen, and Tank, 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 so-called 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, 2010; Carr, Jadhav, and Frank, 2011; Eagleman and Dragoi, 2012; Jadhav, et al., 2012; Olafsdottir, Bush, and Barry, 2018), but little is known about the mechanisms enabling the underlying sequential activity to become 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 speed, spontaneous activity, and information transmission. Finally, we show how LTP-IE-based sequences can be used to induce temporary stimulus-response mappings in an otherwise untrained recurrent network. We discuss implications for cognitive flexibility and rapid, non-stereotyped 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). Overall neurons received three types of inputs: sensorimotor inputs *S*, carrying tuned “external” information; recurrent input *R* from other cells in the network (developed further, Fig 2); and “gating” inputs *G*, which are assumed to be random, with a homogeneous rate across the network, but independent to each cell. We first provide an overview and intuition of the model that follows.

First, we demonstrate our implementation of LTP-IE in spiking neurons: LTP-IE in a given neuron is activated by that neuron’s firing within a physiologically plausible range (~20 Hz) for 1-2 seconds. LTP-IE specifically augments EPSPs from gating (*G*) inputs (analogous to MEC inputs); consequently, recently activated cells, “tagged” by LTP-IE, will show stronger positive voltage deflections in response to *G* inputs. Thus, when continuous *G* inputs are present, recently active, LTP-IE-tagged cells will, on average, sit closer to spiking threshold. Unless otherwise specified, in all simulations of reactivation dynamics that follow we assume the presence of *G* inputs, as LTP-IE cannot 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 neurons maximally activate when a simulated animal is in different positions in the environment. Consequently, when the simulated animal navigates along a trajectory, LTP-IE tags the cells with place fields along the trajectory, thus storing the trajectory memory as a set of positions, without order or timing information. When *G* inputs are turned on post-navigation (representing an awake resting

state), these recently active LTP-IE-tagged cells sit closer to spiking threshold than cells that were not activated by the trajectory.

Finally, inspired by excitatory recurrence in CA3 ([Lisman, 1999](#)) we assume recurrent excitatory connections in the network, which allows activity to propagate through the network over time. We assume that cell pairs with more similar tuning (nearby “place fields”) have stronger connections ([Lisman, 1999](#); [Kali 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 primarily among LTP-IE-tagged cells, in an order potentially reflecting the original trajectory. We discuss additional model components we excluded for the sake parsimony, such as inhibition, structured G inputs, additional tuning parameters, and LTP-IE extinction, 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 voltage of LTP-IE-tagged cells, which makes them more likely to spike in response to inputs.

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

We first model LTP-IE mathematically and demonstrate its effect on neuronal membrane voltages and spiking. To model LTP-IE, when a naive cell spikes at a sufficiently high rate (~10 Hz) over 1 second, we increase the strength of that cell’s G inputs by a factor σ , in accordance with its spike rate, such that future G inputs 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**).

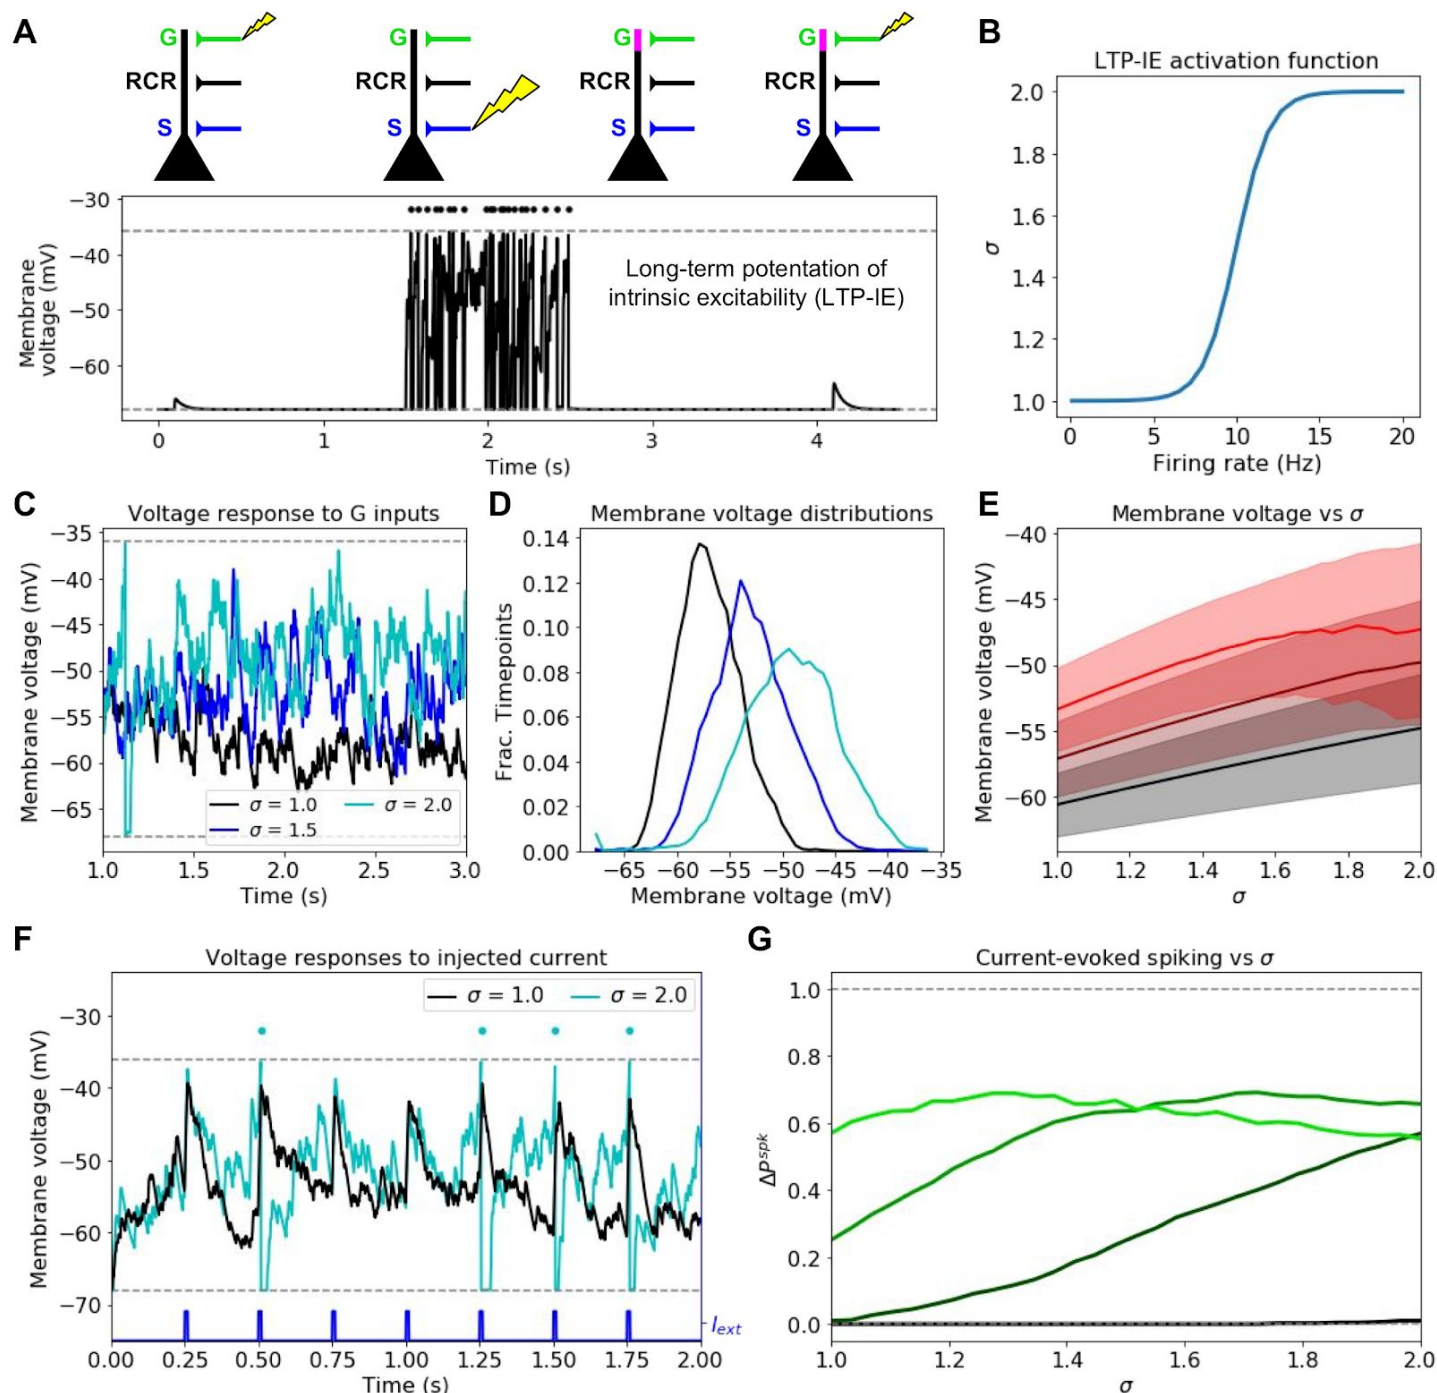


Fig 1: Mechanism and consequences of LTP-IE. A. Demonstration using leaky integrate-and-fire (LIF) neuron model of fast activity-driven LTP-IE, which doubles membrane voltage responses to gate inputs as described in [Hyun, et al., 2013](#), [Hyun, et al., 2015](#). A spike in an upstream gate neuron (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). B. Shifted logistic function for LTP-IE strength (effective weight scaling factor) σ vs. PC firing rate over 1 second. C. Example voltage responses (V_m) of PCs with different σ receiving statistically identical stochastic 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 $\sigma = 2$ case [cyan]). G. Difference between

current-evoked spike probability and spontaneous spike probability as a function of σ for four initial gate input weights (unitless). Color code, in order of increasing lightness: $w^{PG,G} = .004, .008, .012, .016$.

Due to their increased membrane voltage, cells with high LTP-IE also exhibit higher spiking probability, both spontaneously (**Fig 1S1**) and in response to depolarizing current inputs (**Fig 1F,G**), although substantial variability remains, and even with high LTP-IE, current-evoked spiking is not assured (**Fig 1F,G**). Nonetheless, for moderate values of the synaptic strength of initial G inputs, LTP-IE substantially increases the probability of transforming a depolarizing input into an output spike from near zero to approximately 0.5 (**Fig 1G**). (For overly weak initial G inputs, LTP-IE does not substantially facilitate spiking; for overly strong initial G inputs, spontaneous spike rates are already high and increased spiking leads to more frequent voltage resetting [**Fig 1G, 1S1**].) Thus, given moderate initial G input strengths, when a cell undergoes LTP-IE, its chance of spiking in response to future inputs (i.e. its excitability) substantially increases, although variability remains. Since LTP-IE-triggering spike rates are only around 10 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 would shape activity in a recurrent spiking network of 1000 excitatory neurons. Individual neurons were inspired by hippocampal place cells ([O'Keefe, 1979](#); [Moser, Kropff, and Moser, 2008](#)), with firing rates tuned to specific positions ("place fields") within an environment (**Fig 2A-C**). We furthermore assumed preferential connections between neurons with nearby place fields (**Fig 2D** and Methods). Since our primary goal was to evaluate whether LTP-IE alone was sufficient to induce activity sequences we did not include inhibitory neurons, but potential roles for inhibition are described in the Discussion.

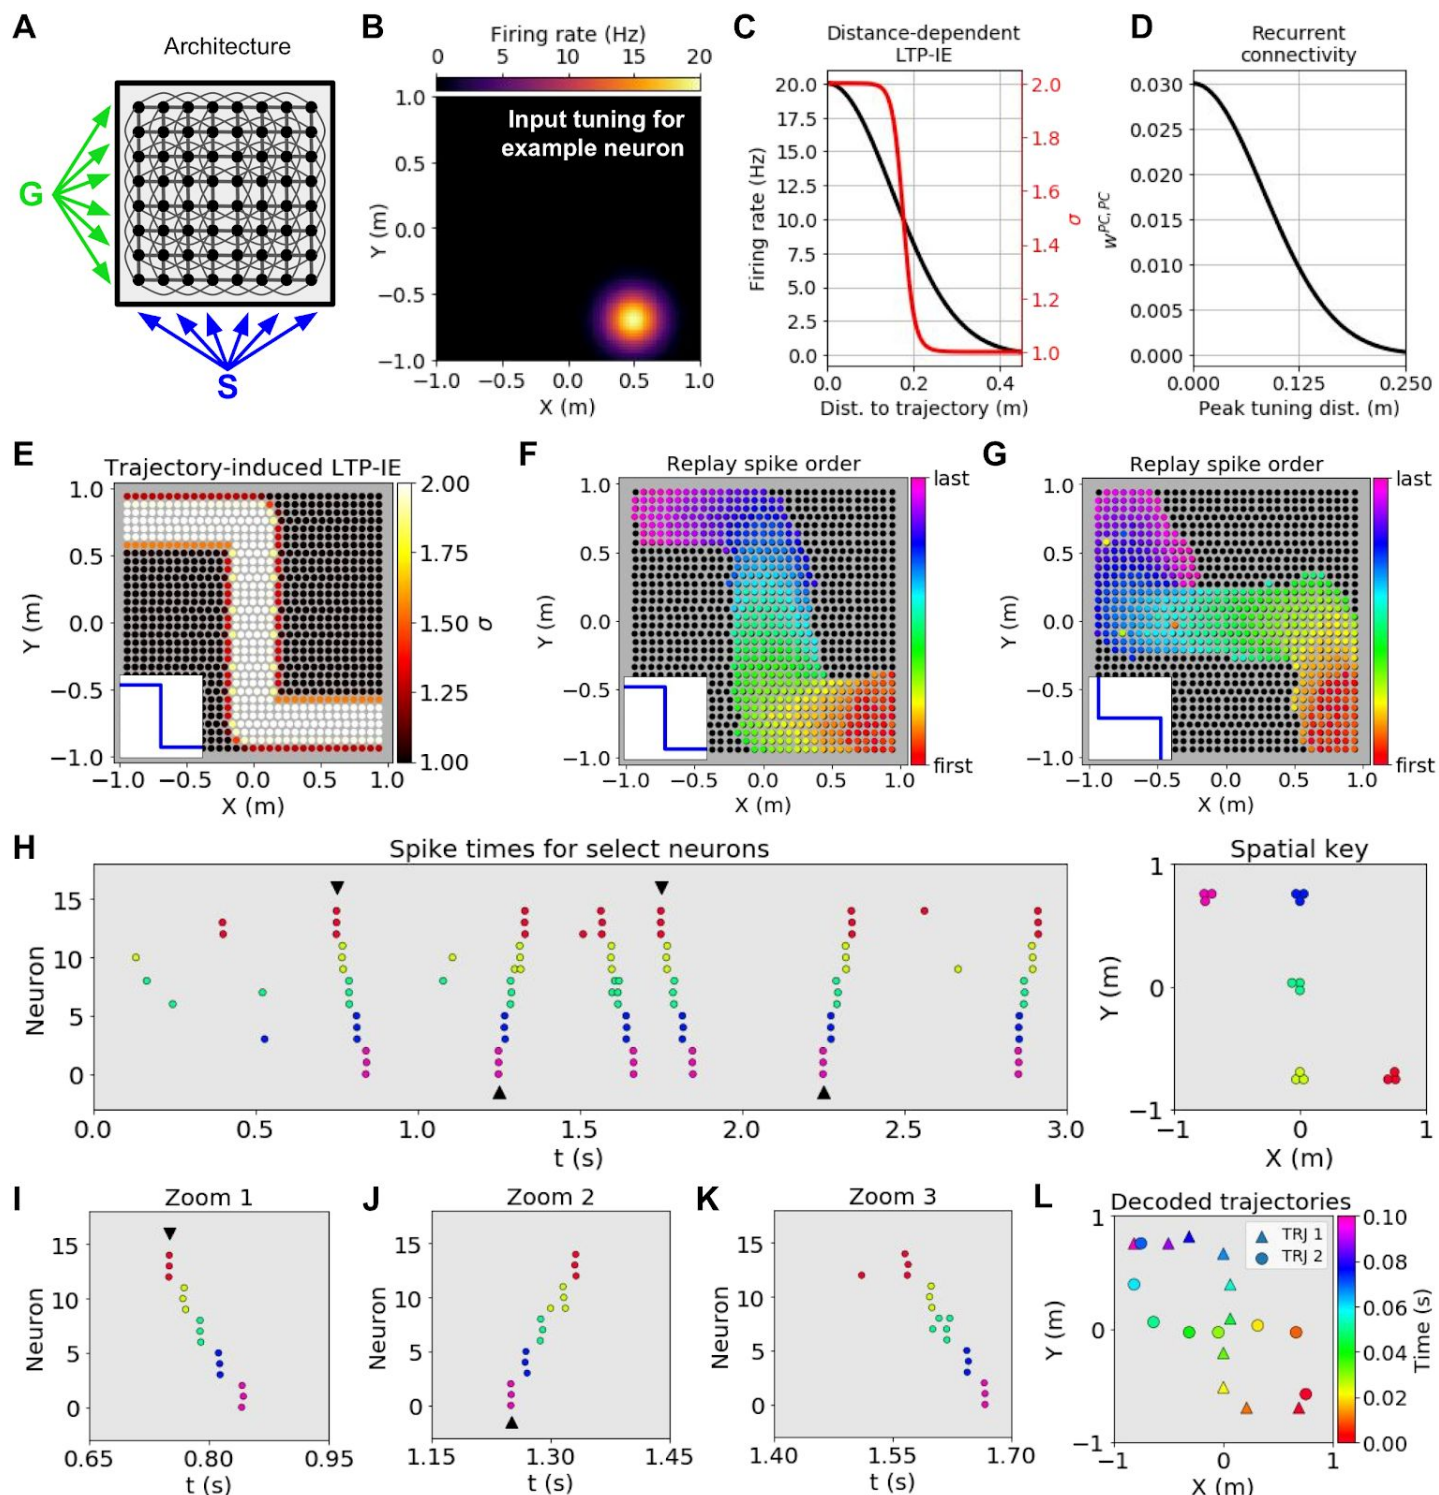


Fig 2: Demonstration of history-dependent sequence replay via LTP-IE. A. Network architecture used in simulations. B. Squared exponential position tuning for an example neuron (20 Hz max. firing rate, 0.15 m length scale). C. Resultant LTP-IE as a function of distance between a cell's maximum tuning and the closest point on the trajectory through the environment (red), computed as a sigmoidal function (Fig 1B) of position-dependent firing rate (black). D. Recurrent weights as squared exponential function of distance between peak tuning positions of two cells. E. LTP-IE profile induced by example trajectory (inset). Cells are positioned according to peak tuning. F. Cells activated during a 100 ms replay epoch, colored by the ordering of the first spikes emitted during replay. Black cells did not activate during replay. G. As in F, but for a different trajectory (inset). H. Raster plot depicting spike times for select cells across several replay epochs. I. Zoom 1. J. Zoom 2. K. Zoom 3. L. Decoded trajectories.

Current-pulse triggers to cells in upper-left or lower-right portion of trajectory are shown by triangles and inverted triangles, respectively. I-K. Zoom in on two triggered replay events (I, J) and one spontaneous event (K). L. Trajectories decoded from replay spikes when LTP-IE was induced by trajectories in F or G.

To investigate network activity depending on recent sensorimotor sequences, we considered the scenario in which a rodent has just completed a short trajectory through its environment, in analogy with experimental setups used to measure neuronal replay in hippocampus ([Foster and Wilson, 2006](#); [Davidson, Kloosterman, and Wilson, 2009](#); [Gupta, et al., 2010](#); [Carr, Jadhav, and Frank, 2011](#)). Since our goal was to understand sequential reactivation during the awake, quiescent period following the trajectory, we only used the trajectory geometry to predict the LTP-IE levels one would expect following its termination. 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 (**Fig 2C**). We then passed the result through the soft-thresholding LTP-IE activation function (**Fig 1B**) to compute the final expected LTP-IE level σ (**Fig 2C**). This allowed us to model the expected LTP-IE profile over the network of neurons as a function of the recent trajectory (**Fig 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? Indeed, when we triggered replay with a depolarizing current pulse into cells at the lower-right corner of the trajectory, an approximately 100 ms sequence of activity cascaded along the path through the network aligned with the LTP-IE profile (**Fig 2F**). Due to the refractory period of ~ 8 ms (see *Methods*), activity propagated in one direction without reversing. When an alternative trajectory was used to induce the LTP-IE profile, triggered replay recapitulated that trajectory instead (**Fig 2G**), indicating the replay of a specific trajectory was not an *a priori* consequence of the recurrent connectivity; indeed, this is to be expected from the spatial uniformity of the recurrence. Additionally, the replay sequence was able to turn corners, suggesting robustness to multidimensional trajectory geometries. Moreover, replay could occur in either direction along the LTP-IE profile depending on which end of the trajectory received the initial trigger (**Fig 2H-L**), in line with empirical observations in which replay often occurs in the reverse direction as the original trajectory ([Foster and Wilson, 2006](#); [Davidson, Kloosterman, and Wilson, 2009](#); [Gupta, et al., 2010](#); [Carr, Jadhav, and Frank, 2011](#)). Replay events also arose spontaneously, in either direction (**Fig 2H, K**), and sometimes—due to the lack of inhibition—arose from the middle of the trajectory, propagating outwards in both directions (not shown). Thus, despite sizable variability in membrane voltages arising from the random gating inputs G , LTP-IE was able to induce 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 through the network? To address this we explored how the network activity evoked by a trigger at the lower-right corner of the Z-shaped LTP-IE profile from **Fig 2E** changed as we varied different pairs of network parameters. For each parameter combination we classified triggered activity into one of four classes: in *fadeout* (**Fig 3A**), activity fails to propagate along the full LTP-IE profile; in *blowup* (**Fig 3B**), activity propagates through whole network instead of being constrained to the LTP-IE profile; in *disordered replay* (**Fig 3C**), activity remains constrained to the LTP-IE profile but does not propagate strictly from one end to the other, as may occur if a spontaneous event arises at the same time as a triggered one. If neither fadeout, blowup, nor disordered replay occurred, we classified the response as *successful replay*. In general, successful replay was sensitive to certain parameter variations but robust to others (**Fig 3D-F**). For instance, whereas a fine balance was required between the recurrent connection strength and the recurrent connectivity length scale (**Fig 3D**), replay was more robust to

variations in the gating input rate r_G and the maximum LTP-IE strength σ_{\max} (**Fig 3E**). (We note, however, that sensitivity to recurrent connection strength and length scale could be improved by increasing the density of cells in the network [**Fig 3S1**].) In general, excessive decreases in any of these four parameters led to fadeout whereas excessive increases led to blowup, with disordered replay occurring when either σ_{\max} or r_G was too large while other parameters retained moderate values (**Fig 3E-F**).

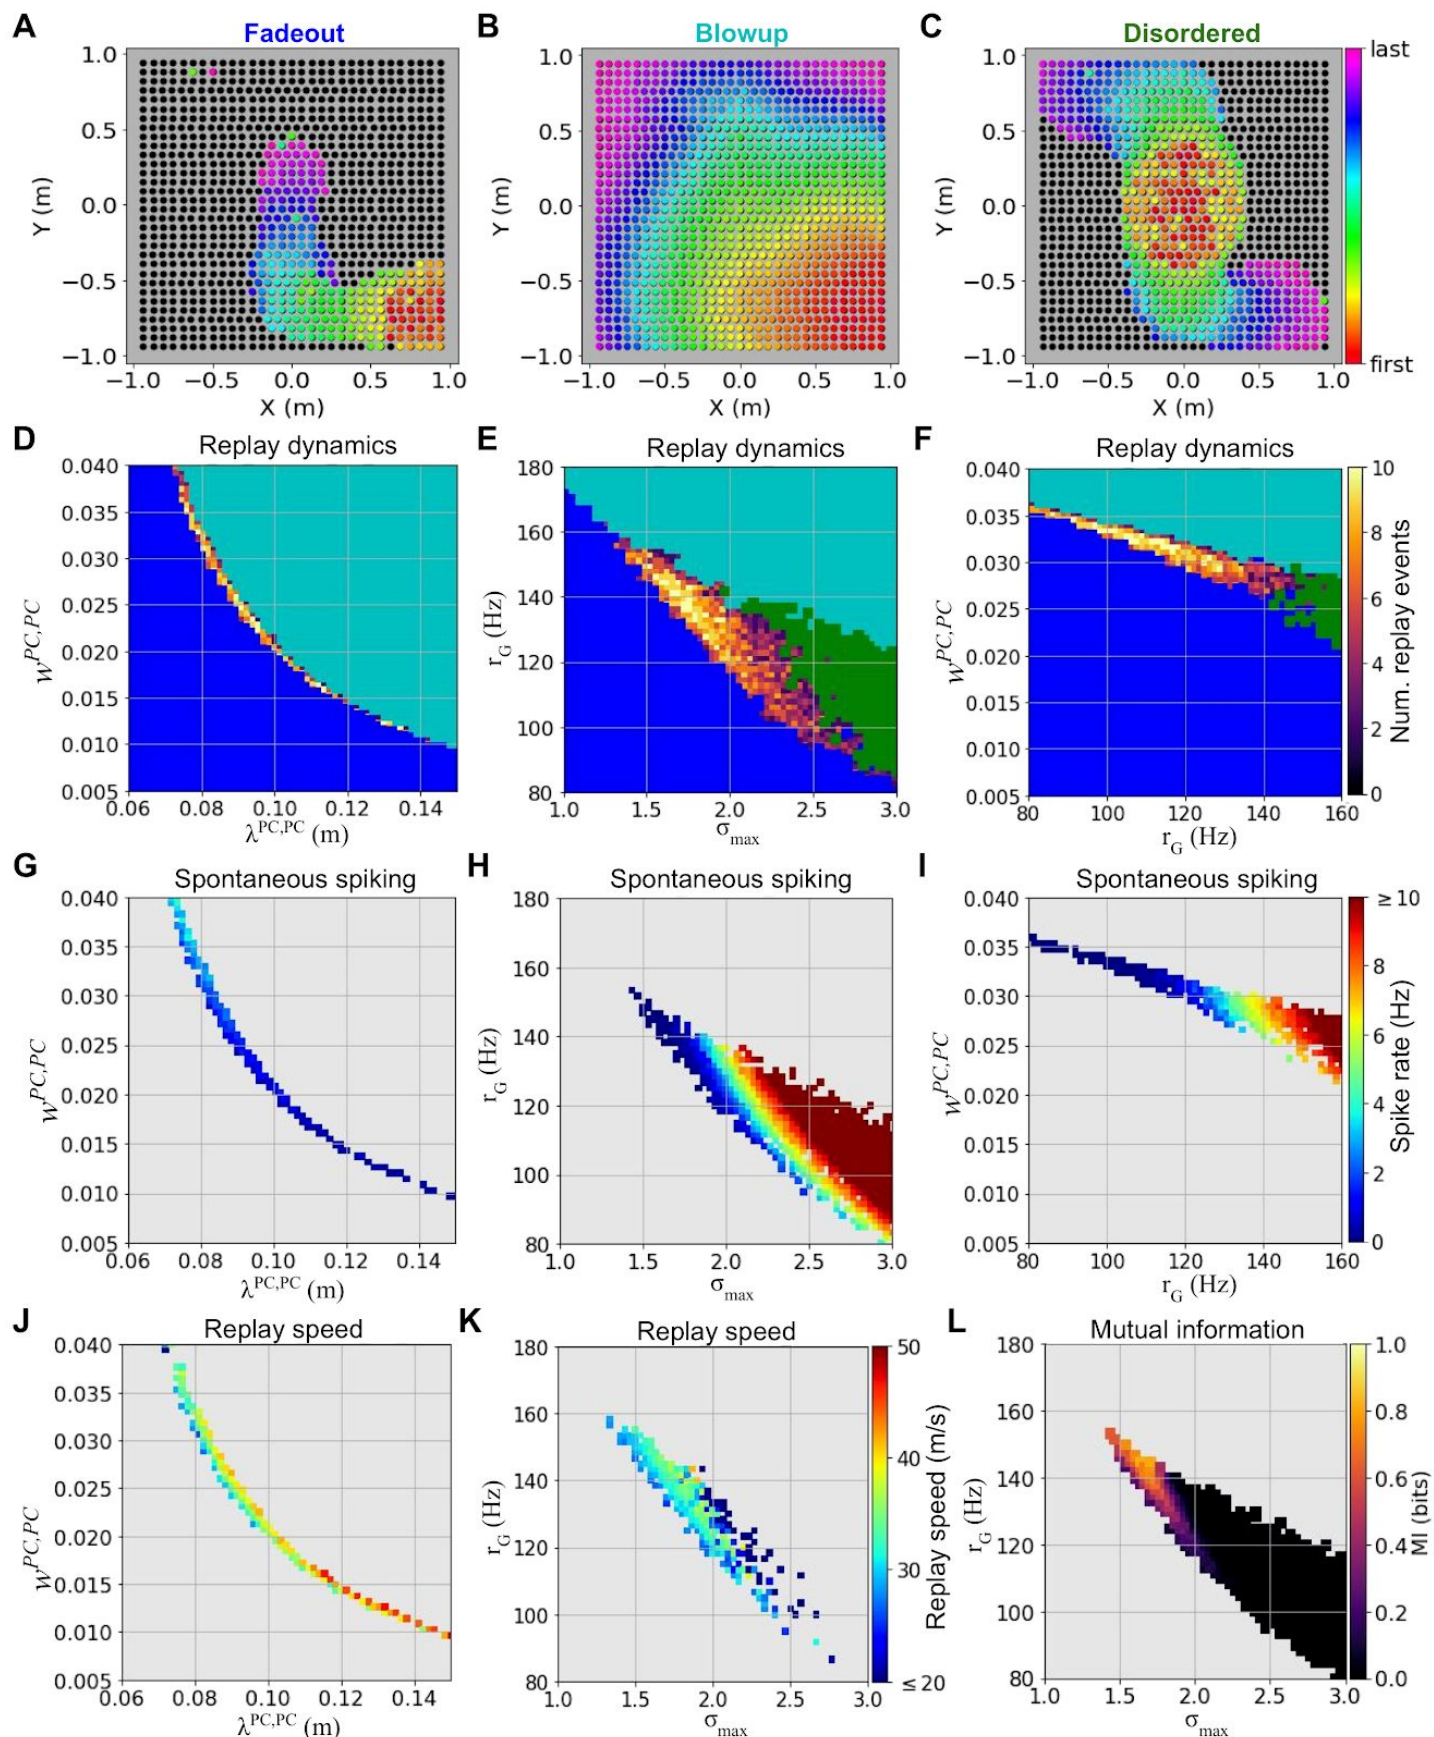


Fig 3: Parameter dependence of LTP-IE-based sequential replay. A. Demonstration of “fadeout”, in which sequential activity initiated at lower right fails to propagate along full trajectory defined by LTP-IE profile. Same format as Fig 2F-G, with colors indicating spiking order (with only first spikes used for neurons spiking multiple

times) during replay epoch. B. Demonstration of “blowup”, in which all neurons activate during replay epoch. C. Demonstration of “disordered” replay, in which LTP-IE-tagged neurons activate during replay epoch but not in either a fully forward or fully reverse order. In this example, activity was triggered by a current injection in the middle of the trajectory, which caused propagation in both directions. D. Profile of replay dynamics regimes as a function of recurrent connectivity length scale $\lambda^{PC,PC}$ and scale factor $w^{PC,PC}$ (see methods). Colors corresponding to Fig 3F colorbar indicate number of successful replay events out of ten trials, for parameter regimes that yielded successful replay. If fewer than 2/10 successful replay events occurred, colors indicate most common dynamics (fadeout: blue; blowup: cyan; disordered: green). E. Same as D but showing dynamics as function of gate input rate r_G and max LTP-IE value σ_{max} . F. Same as E, F but showing dynamics as function of $w^{PC,PC}$ and r_G . G. Spontaneous spike rate of neurons with at least 50% of max LTP-IE value, calculated over 10-second time period absent external stimulation, as function of $\lambda^{PC,PC}$ and $w^{PC,PC}$. Spontaneous spike rates were calculated for all non-fadeout/non-blowup parameter regimes, i.e. where ordered or disordered replay occurred. H. Same as G but as function of r_G and σ_{max} . I. Same as G, H except as function of $w^{PC,PC}$ and r_G . J. “Virtual” replay speed of triggered sequential activity propagation for all parameter regimes in G. K. Same as J but for parameter regimes in H. L. Mutual information between trigger presence and spike detection at upper left between 50-250 ms following trigger, calculated for all non-fadeout and non-blowup parameter regimes. Perfect information transmission is 1 bit.

Within parameter regimes enabling successful replay, precise parameter values modulated specific activity dynamics. For instance, spontaneous spiking (typically dominated by spontaneous replay events) increased with the recurrent connection strength, maximum LTP-IE strength, and gating input rate (**Fig 3G-I**). Furthermore, the “replay speed” (the speed of movement along the trajectory decoded during the replay event) showed the greatest positive relationship with the recurrence length scale, but also exhibited some dependence on other parameters (**Fig 3J-K**). This result corroborates the hypothesis that the speed of replay events emerges from internal network structure, as opposed to the temporal structure of the animal’s original trajectory ([Davidson, Kloosterman, and Wilson, 2009](#)). Note that empirical replay speeds have been measured near ~8 m/s ([Davidson, Kloosterman, and Wilson, 2009](#)) whereas ours were in the range of 20-50 m/s; this range is largely determined by the density of our network, however. When we doubled the density of neurons (2000 neurons with place fields spanning the same environment), allowing us to decrease the recurrence length scale, replay became more reliable, and replay speed decreased substantially across all parameter sets (compare **Fig 3D** with **Fig 3S1A** and **Fig 3J-K** with **Fig 3S1B-C**).

A key component of brain dynamics underlying cognitive flexibility may be the ability to selectively route information from one neural subnetwork to another ([Akam and Kullmann, 2010](#); [Stocco, Lebiere, and Anderson, 2010](#); [Kirst, Timme, and Battaglia, 2016](#)). One substrate for such selective information propagation may be flexible sequential activity patterns that coordinate activity between disparate areas within the network ([Pezzulo, et al., 2014](#); [Jadhav, et al., 2016](#); [Rajan, Harvey, and Tank, 2016](#)), which our results suggest LTP-IE may facilitate. For instance, if activating one end of the LTP-IE profile reliably causes activation at the other, this could be used to quickly carve out a temporary communication channel within the network. To this end, we measured the mutual information (MI) between the presence/absence of a triggering pulse at the lower right end of the trajectory, and subsequent elevated spiking activity at the opposing end in the upper left (where perfect transmission is 1 bit). When we computed MI along the LTP-IE profile as a function of gating input rate r_G and maximum LTP-IE level σ_{max} , which both influence the spontaneous spiking that is most likely to interrupt MI, we found that MI was maximized for a high r_G and moderate σ_{max} , although MI remained below 1 bit for all parameter values, indicating that transmission was never 100% reliable (**Fig 3L**). This suggests that for LTP-IE-induced sequences, information transmission through these sequences may be maximized for moderate values of σ_{max} , so stronger differentiation between LTP-IE-tagged and non-LTP-IE-tagged cells does not necessarily imply improved information transmission. We note further that we could increase MI to 1 by

increasing the density of neurons in the network and shortening their recurrence length scale (**Fig 3S1C**), with information transmission likely improving due to decreased variability in the average signal propagation (since more neurons are involved). Thus, experience-dependent LTP-IE at biologically realistic levels can rapidly and flexibly modulate sequence-based information propagation through a network of spiking neurons.

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 sensorimotor sequences, 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 (**Fig 4A**). We assume such an induction could be generated by appropriately transformed sensorimotor inputs from upstream areas, but we did not model this explicitly, as our goal was to demonstrate the final results.

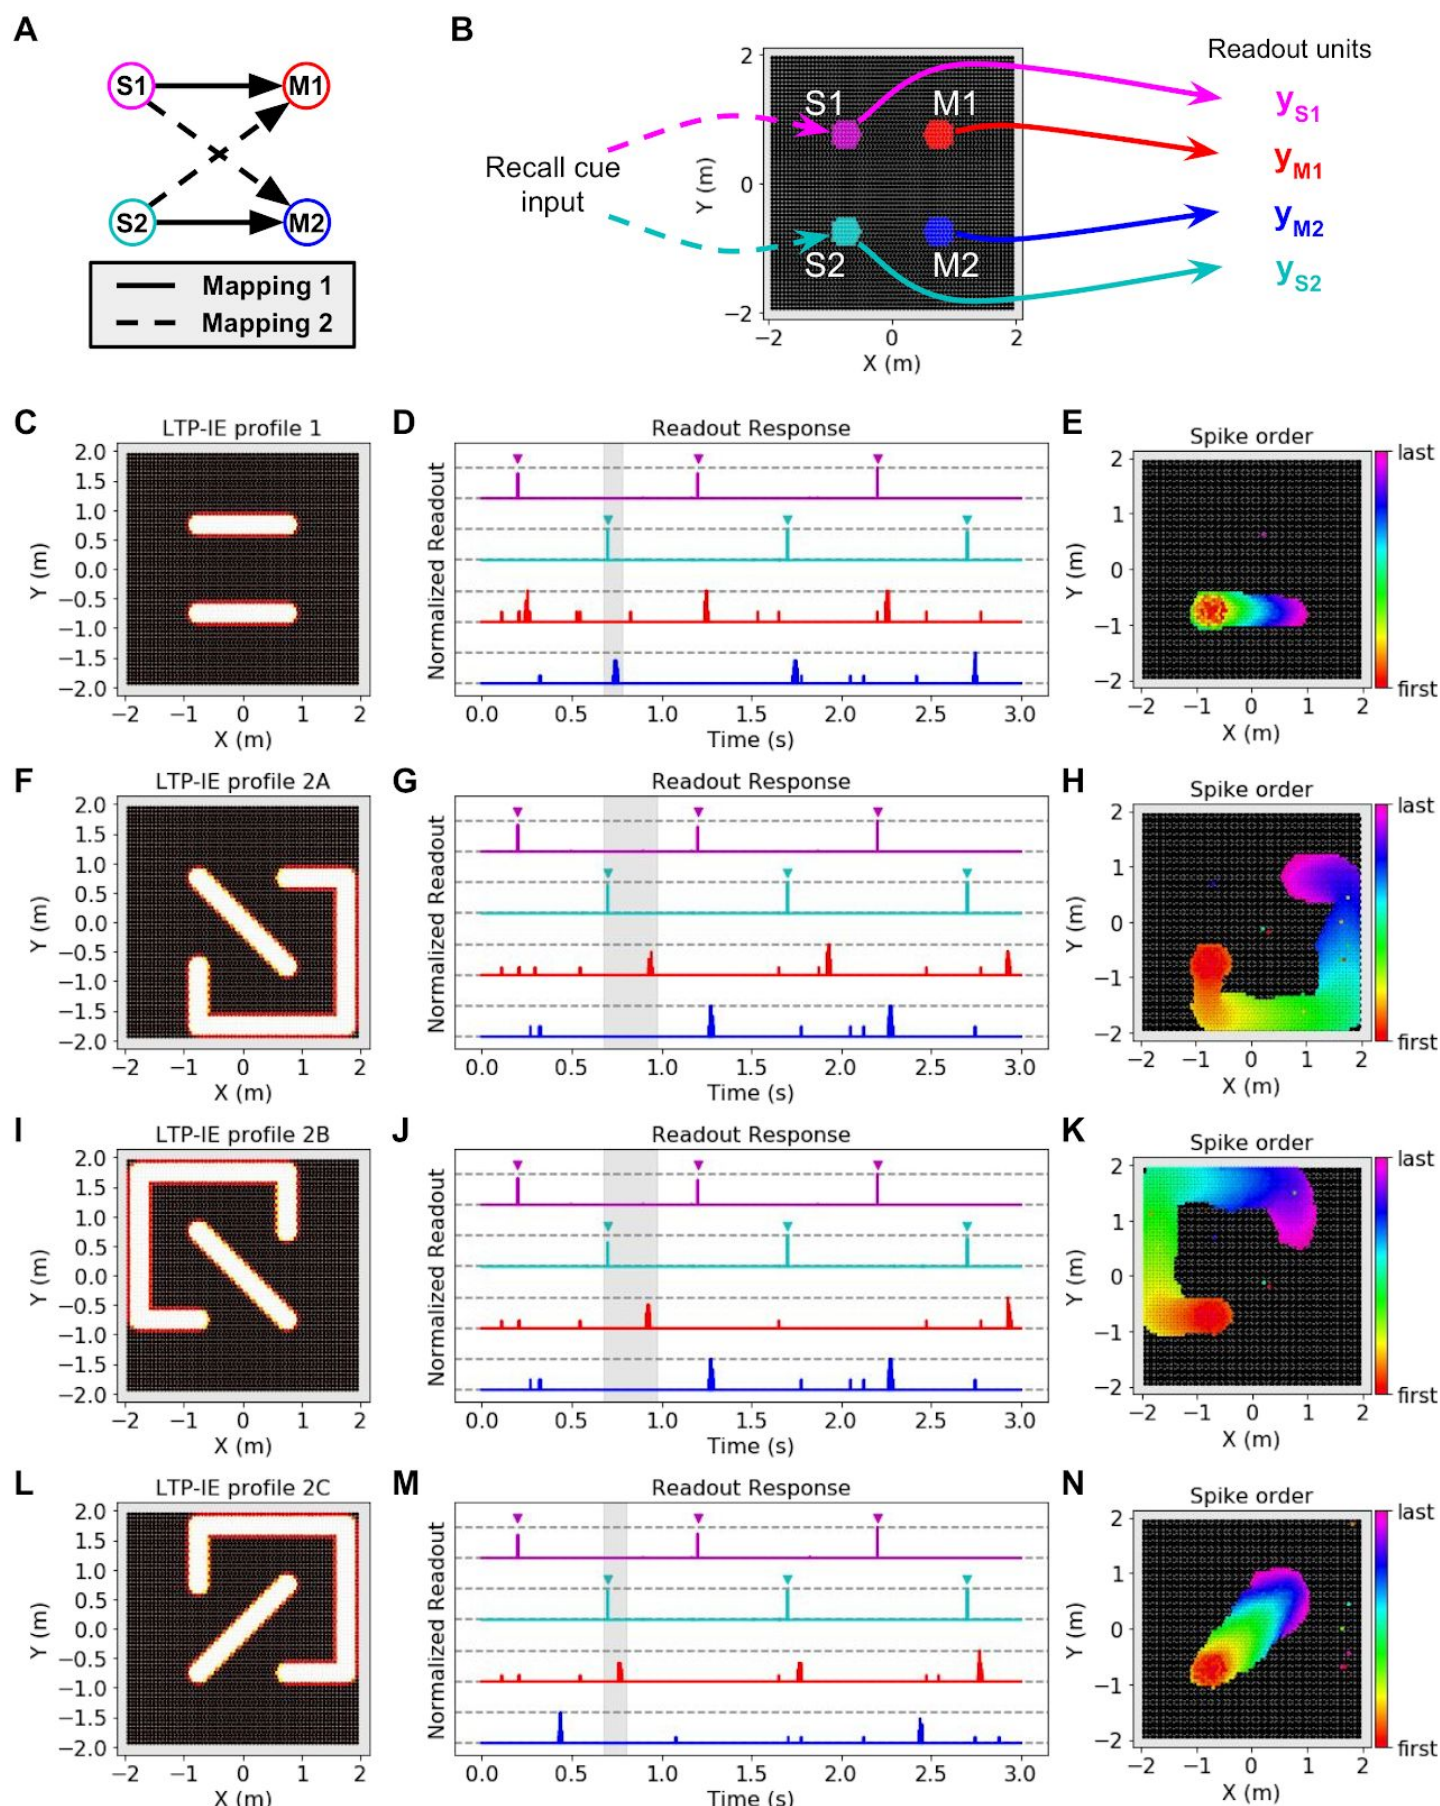


Fig 4: Using LTP-IE-based sequential activity propagation to maintain and decode pairwise associations. A. Example of two different mappings between a pair of stimuli (S1, S2) and a pair of responses

(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. Depiction of simulation structure. Either S1 or S2 is activated with a recall cue (injected current input into colored cells). Readout units average all units from either S1, S2, M1, or M2, indicated by colors. C. Example LTP-IE-based encoding of Mapping 1 with neurons laid out in 2-D coordinate space. D. Time-dependent readout responses (normalized to maximum readout response over the 3 s), 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; 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. E. Order of first spikes of all neurons that spiked during shaded time period in D. F. Example LTP-IE-based encoding of Mapping 2 with neurons laid out in same 2-D coordinate space as in C. G. Readout responses for same recall cue pattern as D. H. First spike order of all neurons that spiked during time-period shaded in G. I-K. Same as F-H but for an LTP-IE profile encoding Mapping 2 that is distinct from that in F. L-N. Same as F-H, I-K but for a third alternative LTP-IE profile encoding Mapping 2.

LTP-IE was able to induce multiple stimulus-response mappings in the network, each in multiple ways. As an example we let two regions in the network represent two stimuli S1 and S2, and two regions represent motor outputs M1 and M2, with the average activity in each region read out by a corresponding downstream readout unit (**Fig 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 a trigger in each stimulus region to subsequently activate the correct motor output (**Fig 4C-N**). (We discuss intersections in the Discussion.) We first verified our idea with a simple mapping in which each stimulus was connected via a one-dimensional LTP-IE profile to its appropriate motor output (**Fig 4C-E**). Indeed, triggering each stimulus region with a depolarizing current input led to subsequent activation of its corresponding motor output after a short delay (**Fig 4D-E**), although there were occasionally low levels of spontaneous activity. We next explored a mapping requiring at least one LTP-IE path to take a roundabout course through the network (**Fig 4F-H**). 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 (**Fig 4G-H**). Responses, however, appeared less reliable than when shorter LTP-IE paths were used, likely due to the increased opportunities for propagation to fail due to variability in the baseline G inputs. Notably, the same stimulus-response mapping could be implemented via LTP-IE in several different ways (**Fig 4F, I, L**), 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 single stereotyped 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, yet highly flexible substrate for inducing temporary stimulus-response mappings in a recurrent spiking network, a potential key property enabling rapid and flexible induction of temporary information-processing patterns in brain networks underlying cognitive flexibility.

Discussion

Structured sequential spiking activity is a key feature of recurrent neural network dynamics, potentially reflecting information flow and computations within the network. One compelling example of spike sequences observed is the so-called “replay” of neural sequences representing recent sensorimotor sequences through the environment. This 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 navigational planning ([Foster 2012, Curr. Opin. Neurobiol.](#); [Pfeiffer 2013, Nature](#)) and memory consolidation ([Carr, Jadhav, and Frank, 2011](#); [Jadhav, et al., 2012](#); [Olafsdottir, Carpenter, and Barry, 2017](#); [Zielinski, Tang, and Jadhav, 2017](#)). Such replay serves as a compelling entrypoint for uncovering the biophysical mechanisms that support rapid modulation of information processing. Accordingly, here we take the perspective that sequential spike patterns reflect information transmission between distinct subnetworks. While substantial work has shown how modifications to the structure of recurrent network connections can induce sequences into a network's dynamical repertoire ([Sussillo and Abbott, 2009](#); [Klampfl and Maass, 2013](#); [Laje and Buonomano, 2013](#); [Rajan, Harvey, and Tank, 2016](#)), little was known about alternative mechanisms acting on faster timescales. Here we have demonstrated the sufficiency of an empirically observed, 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](#); [Rebola, Carta, and Mulle, 2017](#)), applied to a spatially organized recurrent spiking network, to coerce a network into generating selective sequences. These are capable of recapitulating recent sensorimotor sequences, implicating their potential role in memory storage, as well as encoding simple stimulus-response mappings, suggesting their suitability for organizing flexible computations.

Additionally, as observed *in vivo* and as predicted by our model, the sequences that replay during awake quiescence are significantly compressed relative to the initial behavioral timescales, suggesting replayed spikes may occur within the appropriate time windows for long-term, spike-timing-dependent plasticity mechanisms (STDP) ([Markram, et al. 1997](#); [Bi and Poo, 1998](#); [Dan and Poo, 2004](#)). For example, it was shown in CA1 place cells 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](#)). Our model reveals a mechanism by which spatial sequences could be transiently stored in the first place, in order to be later embedded into long-term memories.

Finally, although our work was inspired primarily by hippocampus and the notion of LTP-IE-based short-term sequences, the core concept of excitability-based network dynamics might naturally extend to cortical areas, and to more diverse timescales ([Zhang and Linden, 2003](#); [Tittley, 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 taken as evidence of Hebbian (associative) synaptic plasticity, our work suggests the same result might also arise simply via increasing the excitability of the neurons in the ensemble. Distinguishing these two mechanisms in cortical circuits and understanding how they interact presents an exciting avenue for further investigation.

Comparison to existing sequence-generation 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 (5-10 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 the network; 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 model, however, memories are directly entered into the network by physiological spiking patterns (**Fig 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 (**Fig 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](#)) assume that the sequences produced by the network are initially induced in the network via modification of recurrent network weights, for example 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 when implemented computationally, 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 in nature (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 (**Fig 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 successfully 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. Overall, it is likely that memory and computational flexibility over short timescales relies on a host of plasticity and dynamics mechanisms, but we propose that LTP-IE may play a significant role.

Model limitations

As a proof-of-concept account demonstrating the sufficiency of a small number of mechanisms to quickly bias selective sequence production in a network, we did not explicitly model several features in order to keep our model parsimonious. First, we did not include inhibition, as we discovered that LTP-IE-based sequences could arise through excitatory connectivity alone. Among a number of possible roles, inhibition in hippocampal and cortical networks supports network oscillations ([Wang and Buzsaki, 1996](#); [Jadhav, et al., 2012](#)) and balances excitatory activity to stabilize dynamics ([Brunel, 2000](#); [Chenkov, Sprekeler, and Kempter, 2017](#)). Indeed, empirical hippocampal sequence replay co-occurs with brief oscillatory events called sharp-wave ripples ([Carr, Jadhav, and Frank, 2011](#)), suggesting that hippocampal inhibitory circuitry may contribute to oscillatory replay patterns, perhaps serving as a clock-like signal for controlling signal propagation. An additional role for inhibition in sequence replay may be to induce competition among reactivating representations, for example so that multiple sequences would not replay simultaneously, or so activity traveling along a branching LTP-IE-path would follow only one branch at a time. Inhibitory balance might also serve to decrease the parameter sensitivity of our model to small variations in $w^{PC,PC}$ or $\lambda^{PC,PC}$ (**Fig 3D**) by providing a counterbalance to over-excitation, but precisely how the structure and dynamics of inhibitory interactions with our network would modulate LTP-IE-based sequences is beyond the scope of this work.

A further issue with LTP-IE-based sequences surrounds trajectories that intersect themselves. At an intersection point, our current model would typically lead to sequential reactivation that spread in multiple directions, thus failing to recapitulate the original trajectory. This, however, could be resolved by the

embedding of trajectories in a higher-dimensional space beyond just the x-y plane. Indeed, neuronal firing is typically modulated by multiple stimulus features; for example, hippocampal neurons are also tuned to head-direction (θ) ([Leutgeb, Ragozzino, and Mizumori, 2000](#)) and running velocity ([McNaughton, Barnes, and O'Keefe, 1983](#)). Thus, if excitatory connectivity between neurons with similar x- and y-tuning were only present if neurons had similar θ -tuning also, the trajectory representation would expand into a third θ dimension, such that it might no longer intersect with itself and replay would reflect the actual trajectory travelled. Consequently, extending our assumption of preferential connectivity among similarly tuned neurons to additional input dimensions could increase the capacity of our network to store non-interfering sequences.

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. Our model demonstrates that replay can occur absent additional mnemonic information in such, 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.

Finally, our model begs the question of how information would be erased from the network or prevented from entering it in the first place. 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. For example, the neurotransmitter acetylcholine is known to interact with the voltage-gated potassium channel Kv1.2 ([Hattar, 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., 2015](#)). Thus, sufficient machinery for erasing or selectively modulating the overall effects of LTP-IE, say, as a function of arousal or motivation state, may converge in the brain region that has served as the inspiration for our model.

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 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 virtual 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 to the animal.

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](#); [Gupta, et al., 2010](#); [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 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-dimensional trajectories if the recurrent architecture reflected a 3-dimensional space, as might occur within the three-dimensional hippocampal place-cell network of free-flying 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 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. Whether this could be done via biological mechanisms remains an open question, but our present work suggests its 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](#)). 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

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

$$\tau_m \frac{dv^i}{dt} = -(v^i - E_{leak}) - g_E^i(t)(v^i - E_E) + I_{ext}^i(t)$$

where v^i is the membrane voltage of neuron i , $\tau_m = 50$ ms is the membrane time constant, $E_{leak} = -68$ mV is the leak/resting potential, g_E is the time-varying conductance of excitatory synaptic inputs, $E_E = 0$ mV is the excitatory synaptic reversal potential and I_{ext} is externally injected currents (expressed as voltage deflections). If v exceeded a threshold $v_{th} = -36$ mV at any time step, a spike was emitted and v was reset to E_{leak} for a refractory period $\tau_r = 8$ ms. This outsize 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 ([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

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

where 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 $\tau_E = 2$ ms. In our simulations we assumed all synapses were excitatory and did not explicitly model effects of inhibition. Weights were unitless since they acted only to scale conductance changes.

Neuron tuning and network structure

In network simulations, we assigned each neuron a position (x^i, y^i) corresponding to the peak of its tuning curve, the (x, y) position eliciting maximal firing. Positions were assigned so as to approximately tile a grid spanning -1 to 1 m in both dimensions (**Fig 2, 3**) or -2 to 2 m (**Fig 4**), using 1000 or 4000 neurons, 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 $\lambda_{PL} = 0.15$ 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} .

We explicitly modeled two sources of synaptic input to the neurons in our recurrent network: recurrent inputs and gating inputs. Recurrent 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

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

This was motivated by past modeling studies ([Kali 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 $w_{min}^{PC,PC} = 0.001$ were set to 0.

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 $r_{max}^i = \max(r(x, y) \mid (x, y) \in \text{trajectory})$. From r_{max}^i we computed σ_i according to the following equation (**Fig 1B**):

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

with $r_{\sigma} = 10$ Hz and $\beta_{\sigma} = 1$. 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, during which all neurons in the recurrent network received independent stochastic gating input G at constant rate r^G . Replay was triggered by applying a short depolarizing current pulse to a set of neurons with peak position-tuning within 0.4 m (**Fig 2-3**) or 0.3 m (**Fig 4**) of a central trigger location (e.g. $(x = 1 \text{ m}, y = -0.75 \text{ m})$ in **Fig 2F**) for 3 ms, of amplitude 8 mV (**Fig 2-3**) or 7 mV (**Fig 4**); exact trigger parameters were unimportant, as long as they elicited significant spiking activity in the neurons at the start of the sequence. In **Fig 2L**, replayed positions were decoded during by taking the median peak-tuning of neurons that spiked within 10 ms windows spanning the replay event.

Simulations

Single-neuron LTP-IE simulation (Fig 1)

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

Replay in networks of neurons (Fig 2-3)

In replay simulations, we used 1000 neurons, with place fields distributed on an approximate lattice over a 2 m x 2 m environment. Replay was triggered by a short, depolarizing current pulse into all cells with place field centers within a 0.4 m radius of either (-1 m, 0.75 m) or (1 m, -0.75 m), although our results were not sensitive to these precise numbers, so long as sufficiently many LTP-IE-tagged cells at one end of the trajectory were activated by the trigger.

Replay dynamics parameter sweep (Fig 3)

For each parameter set we ran ten 800 ms simulations starting with different random number generator seeds (which controlled the instantiation of gating inputs, which is the only stochasticity in our simulation). Each simulation contained one replay trigger at 500 ms, and we counted how many of the ten triggers resulted in either fadeout, blowup, disordered replay, or successful replay. If successful replay occurred 2 or more times out of the 10 simulations, we colored the point in parameter space by how many successful replay events occurred; otherwise we colored it by the most common unsuccessful event type (fadeout: blue; blowup: cyan; disordered: green).

We first varied the magnitude factor $w^{PC,PC}$ and the length scale $\lambda^{PC,PC}$ of the recurrent excitatory weight profile (**Fig 2D**), while holding $\sigma_{\max} = 2$ and $r_G = 125$ Hz. We next varied σ_{\max} and r_G while holding $w^{PC,PC} = 0.03$ and $\lambda^{PC,PC} = 0.083$ m, (**Fig 3E**). Finally, we varied $w^{PC,PC}$ and r_G while fixing $\lambda^{PC,PC} = 0.083$ m and $\sigma_{\max} = 2$ (**Fig 3F**).

Spontaneous activity was determined by computing spike rates over 10 seconds without no additional current inputs. Replay speed was calculated as the slope of a line fit to the neuronal positions along the trajectory profile and the times at which they activated during replay. Note that for high levels of spontaneous activity but without blowup, two replay events sometimes occurred in this window, distorting the speed measure (dark blue dots in **Fig 3K**).

To measure mutual information between a replay trigger at one end of the trajectory and elevated spiking activity at the opposing end, we defined the trigger as an injected current into cells with place fields within 0.4 m of (1 m, -0.75 m) and measured elevated spiking activity of neurons within 0.4 m of (-1 m, 0.75 m) between 50 and 250 ms following the trigger. The binary trigger variable corresponded to the presence or absence of the trigger, and the binary opposing-end-spiking variable was positive if > 90% of the opposing-end neurons spiked and otherwise negative.

LTP-IE-based stimulus-response mappings (Fig 4)

For this simulation, in order to construct sufficiently non-interfering LTP-IE-defined paths, we constructed a model network of 4000 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.

Other simulation parameters

Unless otherwise noted we used the following values for all simulation parameters:

Symbol	Definition	Value	Symbol	Definition	Value
τ_m	Membrane time constant	50 ms	N	Number of neurons in recurrent network	1000 (Fig 2-3) 4000 (Fig 4)
E_{leak}	Leak potential	-68 mV	w^G	Initial gating input weight	0.008216
V_{th}	Spike threshold	-36 mV	$w^{PC,PC}$	Max recurrent weight	.03 (Fig 2-3) .034 (Fig 4)
τ_r	Refractory period	8 ms	$w_{min}^{PC,PC}$	Min nonzero recurrent weight	.001
E_E	Excitatory synaptic reversal potential	0 mV	$\lambda^{PC,PC}$	Recurrent connectivity length scale	.083 m (Fig 2-3) .0835 m (Fig 4)
τ_E	Excitatory synaptic time constant	2 ms	σ_{max}	Max LTP-IE value	2
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
r^G	Gating input firing rate	125 Hz (Fig 2-3) 105 Hz (Fig 4)	ΔT	Simulation time step	0.5 ms

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.

References

- Akam, Thomas, and Dimitri M. Kullmann. "Oscillations and filtering networks support flexible routing of information." *Neuron* 67.2 (2010): 308-320.
- Bakker, Arnold, et al. "Pattern separation in the human hippocampal CA3 and dentate gyrus." *Science* 319.5870 (2008): 1640-1642.
- 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.
- 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 Kempster. "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.
- 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.
- Davidson, Thomas J., Fabian Kloosterman, and Matthew A. Wilson. "Hippocampal replay of extended experience." *Neuron* 63.4 (2009): 497-507.
- 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.

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.

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 underlies the generation of neural sequences in a songbird." *Nature* 419.6902 (2002): 65.

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.

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.
- Jadhav, Shantanu P., et al. "Coordinated excitation and inhibition of prefrontal ensembles during awake hippocampal sharp-wave ripple events." *Neuron* 90.1 (2016): 113-127.
- 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.
- Kirst, Christoph, Marc Timme, and Demian Battaglia. "Dynamic information routing in complex networks." *Nature Communications* 7 (2016): 11061.
- Klampfl, Stefan, and Wolfgang Maass. "Emergence of dynamic memory traces in cortical microcircuit models through STDP." *Journal of Neuroscience* 33.28 (2013): 11515-11529.
- 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.
- 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.
- McNaughton, B. L., Carol A. Barnes, and J. O'keefe. "The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats." *Experimental Brain Research* 52.1 (1983): 41-49.
- 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.
- 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.
- 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.
- Pezzulo, Giovanni, et al. "Internally generated sequences in learning and executing goal-directed behavior." *Trends in Cognitive Sciences* 18.12 (2014)
- 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.
- 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.
- Schiller, Daniela, et al. "Memory and space: towards an understanding of the cognitive map." *Journal of Neuroscience* 35.41 (2015): 13904-13911.
- 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.
- 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.
- Stocco, Andrea, Christian Lebiere, and John R. Anderson. "Conditional routing of information to the cortex: A model of the basal ganglia's role in cognitive coordination." *Psychological Review* 117.2 (2010): 541.
- 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.
- 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.

- Wang, Xiao-Jing, and György Buzsáki. "Gamma oscillation by synaptic inhibition in a hippocampal interneuronal network model." *Journal of Neuroscience* 16.20 (1996): 6402-6413.
- 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).