A model of plasticity-dependent network activity in rodent hippocampus during exploration of novel environments.

Panagiota Theodoni[1,2,*], Bernat Rovira[1,*] and Alex Roxin[1]

[1]Computational Neuroscience Group, Centre de Recerca Matemàtica, Campus de Bellaterra, Edifici C, 08193 Bellaterra, Spain

[2]Institute of Brain and Cognitive Science, NYU Shanghai,

[*]contributed equally

Summary

Place cells of the rodent hippocampus fire action potentials when the animal traverses a particular spatial location in a given environment. Therefore, for any given trajectory one will observe a repeatable sequence of place cell activations as the animal explores. Interestingly, when the animal is quiescent or sleeping, one can observe similar sequences of activation, although at a highly compressed rate, known as "replays". It is hypothesized that this replay underlies the process of memory consolidation whereby memories are "transferred" from hippocampus to cortex. However, it remains unclear how the memory of a particular environment is actually encoded in the place cell activity and what the mechanism for replay is. Here we study how plasticity during spatial exploration shapes the patterns of synaptic connectivity in model networks of place cells. Specifically, we show how plasticity leads to the formation of attracting manifolds: patterns of activity which represent the spatial environment learned. These states become spontaneously active when the animal is quiescent, reproducing the phenomenology of replays. Interestingly, the attractors are formed most rapidly when place cell activity is modulated by an ongoing oscillation. The optimal oscillation frequency can be calculated analytically, is directly related to the plasticity rule, and for experimentally determined values of the plasticity window in rodent slices gives values in the theta range. A major prediction of these model is that the pairwise correlation of place cells which encode for neighboring locations should increase during initial exploration, leading up to the

critical transition. We find such an increase in a population of simultaneously recorded CA1 pyramidal cells from a rat exploring a novel track. Furthermore, in a rat in which hippocampal theta is reduced through inactivition of the medial septum we find no such increase. This is consistent with the model, which predicts orders of magnitude lower learning rates when theta is absent.

Introduction

As an animal explores in any given environment, place cells in the hippocampus fire selectively at particular locations (O'Keefe and Dostrovksy 1971, O'Keefe 1976, Harvey et al. 2009), known as the "place-fields" of the cells. Furthermore, the place fields of an ensemble of place cells remap to completely new positions with respect to one another if the animal enters a distinct environment (Muller and Kubie 1987, Kubie and Muller 1991, Bostock et al 1991). Sequential place cell activation during exploration therefore acts as a unique fingerprint for each environment, providing information needed for navigation and spatial learning. The spontaneous replay of such sequential activation, which occurs within sharp-wave/ripples (SWRs) during quiet wakefulness (Foster and Wilson 2006, Karlsson and Frank 2009, Carr et al. 2011) and sleep (Wilson and McNaughton 1994, Lee and Wilson 2002), suggests that the animal has formed an internal representation of the corresponding environment, presumably during exploration (Wu and Foster 2014). Synaptic plasticity is the clear candidate mechanism for this formation. Nonetheless it remains unclear how changes in the synaptic connectivity are coordinated at the network level in order to generate well-ordered sequences spontaneously.

An additional prominent biomarker of exploratory behavior in the hippocampus is the theta rhythm (4-12Hz). Decreases in theta power due to lesions of the medial septum strongly reduce performance in spatial-memory based tasks (Winson 1978, Mitchell et al. 1982, Dwyer et al. 2007) and other hippocampal dependent tasks (Berry and Thomson 1979, Allen et al. 2002), although they do not eliminate place fields (Brandon et al. 2014, Wang et al. 2015). Despite this, lesioned animals can still reach fixed performance criteria given enough time (Winson 1978, Berry and Thomson 1979). This shows not only that theta is important for learning, but also suggests a potential

mechanism. Namely, it may act as a dial for the learning rate, presumably by modulating the network-wide coordination of synaptic plasticity. How this occurs in unknown.

Here, we study a neural network model to investigate how synaptic plasticity shapes the patterns of recurrent connectivity in a hippocampal circuit as an animal explores a novel environment. We show that spike-timing dependent plasticity (STDP) (Markram et al. 1997, Bi and Poo 1998) during motion-driven sequential place-cell activity leads to the formation of network structure capable up supporting spontaneous bursts. These bursts occur in the absence of place-field input, i.e. during awake-quiescence or sleep. Importantly, the spatio-temporal structure of the bursts undergoes a sharp transition during exploration, exhibiting well-ordered replay only after a critical time of exploration in the novel environment. The underlying rate of this plasticity process, and hence the critical time, is strongly modulated through external oscillatory drive: for very low and high frequency the rate is near zero, while for an intermediate range, set by the time-scale of STDP, it is higher by several orders of magnitude, allowing for learning on realistic time-scales. Our theoretical findings lead us to propose that the theta rhythm accelerates learning by modulating place-cell activity on a time-scale commensurate with the STDP window. This maximizes the growth rate of network-wide patterns of synaptic connectivity which drive spontaneous replay. Finally, we test and verify a main prediction from the model using simultaneous recordings of hippocampal place cells from a rat exploring a novel track. Namely, pairwise correlations between cells with neighboring place fields show a sharp increase over the span of several minutes at the outset of exploration. Furthermore, in a rat in which the medial septum is partially inactivated via muscimol injection, which strongly reduces theta modulation, this increase is not seen.

Results

Plasticity during exploration of a novel environment leads to a transition in the structure of SWRs

We modeled the dynamics of hippocampal place cells as an animal sequentially explored a series of distinct novel ring-like tracks, Fig.1a. The model consisted of recurrently coupled, excitatory

stochastic firing rate neurons, each of which received a place-specific external input on any given track, and inhibition was modeled as a global inhibitory feedback, see e*xperimental procedures* for model details. To model the global remapping of place fields from one track to another, we randomized the position of the peak input, i.e. the place field, of each neuron. In this way, the ordering of cells according to their place field location on one track was random and uncorrelated with their ordering on any other track, see Fig.1b.

We were interested in knowing how the exploration of a novel environment affected the pattern of synaptic connectivity between place cells via a spike-timing dependent plasticity (STDP) rule (Kempter et al. 1999, Song et al. 2000, Pfister and Gerstner 2006), and how this in turn shaped the activity. In order to study "typical" place-cell dynamics during exploration we first exposed the network to a series of distinct tracks until the matrix of recurrent synaptic weights reached a steady state; this occurred by the third track although we simulated a total of ten, see Fig.S1. At this point we exposed the network to another novel track and studied the dynamics in detail; because the exploration process had already become stereotyped, the dynamics on the novel track reflected what would be seen during the exploration of any other novel track in the future.

In our simulations, we modeled the movement of a virtual animal around the track by varying the external inputs to place cells. Specifically, the external input was maximal for cells with place fields at the animal's current position and minimal for cells with place fields at the opposite side of the track, with the input decaying smoothly for intermediate locations like a cosine curve, see e*xperimental procedures* for details. For simplicity we modeled an animal moving around the track at a constant velocity (we will use the metaphor of a virtual animal for conceptual ease, although we really mean that we moved the bump-like external input to our model neurons at a constant rate in time), although a non-constant velocity did not change the results qualitatively, see Fig.S4 and *supplementary materials*. Every three minutes of simulation time we stopped the animal at its

current location for three seconds and removed the place field input. This lead to spontaneous bursting via a synaptic depression-dependent mechanism (Romani and Tsodyks 2015), reminiscent of sharp-waves ripples (SWRs) seen during awake quiescence and sleep, see Fig.2.

As the virtual animal looped around the track the sequential activation of the place cells lead to changes in the recurrent connectivity via the pairwise STDP rule, Fig.2a. Whereas the connectivity between cells was initially unstructured on the novel track, it evolved over the span of minutes to tens of minutes such that cells with nearby place fields were more strongly connected that ones with disparate place fields. Furthermore the resulting connectivity was asymmetric, reflecting the directionality of the animal's motion. This plasticity resulted in changes to the place cell activity, causing place fields to shift backwards, broaden and become negatively skewed as observed in experiment, Fig.2a and Fig.S2. In simulations where the animal explored both directions without bias, the connectivity was more symmetric, see Fig.S4. The connectivity stabilized after about 1hr of simulation and remained stable throughout the simulation, up to 16hrs. Although the changes we observed in the recurrent connectivity and place fields were continuous and smooth during the duration of the simulation, there was a dramatic and sharp transition in the structure of the SWR activity during awake quiescence. We quantified this transition by measuring the mean pairwise cross-correlation of cells with neighboring place fields *on the novel track*. Such cells are "neighbors" only on the novel track and not on any other track by virtue of the global remapping or randomization of place-fields. We call this measure the sequential correlation (SC) because it is high when there is a properly ordered sequence of place-cell activations on any given track. The SC during theta-activity (exploration) was already non-zero at the outset of the simulation by virtue of the external place-field input which generates sequential activity, red squares Fig.2a. On the other hand, the SC was initially near zero during spontaneous SWRs, black circles Fig.2a. Interestingly the SC abruptly increased during SWRs around 40 minutes and remained elevated for the duration of the simulation. Note that the SC for theta-activity also showed a steady increase leading up to

this transition. Space-time plots of place-cell activity show that the abrupt increase in SC coincides with a transition in the spatio-temporal structure of the SWRs: it is initially uncorrelated with the ordering of cells on the novel track, Fig.2c top, and after the transition there is a clear sequential activation reminiscent of so-called "replay" activity, Fig.2c bottom. In fact, before the transition the SWRs are highly structured, but this structure is correlated with a previously explored track and not the novel one, Fig.S3.

The transition in the structure of SWR is strongly dependent on the choice of STDP rule and the frequency of "theta" modulation

We hypothesized that changes in the recurrent connections between place cells in our model during exploration were shaping the spontaneous activity and driving the transition to replay-like activity. The connectivity profile at any point in time could be decomposed into a series of spatial Fourier modes, Fig.3a. We tracked these modes in time, Fig.3b, and discovered that the transition in SWR activity always occurred when the coefficient of the first even mode reached a critical value, Fig.3c. We then asked how the growth of the even mode depended on the details of the STDP rule and the frequency of periodic modulation of the place cell activity. We found that the growth of the even mode depended strongly on the frequency, peaking in the theta range and decreasing by orders of magnitude at low and high frequencies, Fig.3d. In fact, for simulations of up to 1hr, transitions in the SWR activity were observed only when the frequency was in the theta range, Fig.3e. Furthermore, the even mode only grew, and transitions only occurred, when the STDP rule had dominant potentiation at short latencies. For a perfectly anti-symmetric STDP rule, and for a rule with dominant inhibition at short latencies, the even mode did not change and even decreased, respectively, Fig3f. In the latter case SWRs were suppressed entirely (not shown).

The analysis of a simplified rate model explains how the interplay between theta-modulation

and the STDP rule govern changes in recurrent connectivity

We sought to understand the mechanism underlying the evolution of the recurrent connectivity seen in simulations by studying a simplified, linear version of our network, which was amenable to analysis. In addition, we made use of the fact that the step-wise increases in synaptic strength due to the STDP rule can be approximated as a smooth process as long as plasticity occurs much more slowly than the firing rate dynamics. When this is the case, the rate of change of the synaptic weight from place cell with place field centered at θ' to one with place field at θ can be written

$$\dot{w}(\theta,\theta') = \int_{-\infty}^{\infty} dT\, STDP(T)\, r(\theta,t)\, r(\theta',t+T) \quad (1)$$

This equation reflects the fact that the total change in the synaptic weight is the sum of all the pairwise contributions from the pre- and post-synaptic cells, with each pair of spikes weighted by the STDP rule with the appropriate latency. The firing rates are then found self-consistently from the firing rate equation. In the end we find that

$$\dot{w}(\theta,\theta') = \dot{w}(\theta-\theta') = \dot{w}_{even}\cos(\theta-\theta') + \dot{w}_{odd}\sin(\theta-\theta') \quad (2),$$

where the growth rates of the even and odd modes $\dot{w}_{even}$ and $\dot{w}_{odd}$ are functions of the STDP rule parameters, the velocity of the animal and the frequency of periodic modulation, see *experimental procedures* for details. It turns out it is possible to understand these dependencies intuitively and comprehensively without having to study the analytical formulas. Specifically, if we wish to isolate the growth rate of the even mode, which is responsible for driving the emergence of replay in the SWR, we can consider place cell pairs where $\theta=\theta'$, i.e. pairs with overlapping place fields. When this is the case we can combine Eqs.(1) and (2) to yield

$$\dot{w}_{even} = \int_{-\infty}^{\infty} dT\, STDP(T)\, r(\theta,t)\, r(\theta,t+T) \quad (3).$$ By appropriately averaging both sides of the equation in time (see *experimental procedures)* the product of the rates in the integral becomes the

autocorrelation (AC) of place-cell activity, and so we can write

$$\dot{w}_{even} = \int_{-\infty}^{\infty} dT\, STDP(T)\, AC(T) \quad (4).$$

Therefore the growth rate of the even mode is found by multiplying the AC of place-cell activity by the STDP window, and integrating. Because the effect of periodic modulation on the AC is straightforward, we can determine graphically how the frequency of modulation interacts with the STDP rule to drive changes in SWR structure.

We first note that if there is no periodic modulation of place-cell activity then the AC will simply reflect the movement of the animal. This will lead to a very broad AC compared to the time-scale of the STDP window. For example, if we assume that the width of the place field is a fraction of the track length (as in our model), then a rat running between 5 and 50cm/sec on a multi-meter track would have an AC which decays on the order of between seconds and tens of seconds. Therefore, the AC is essentially constant compared to a typical STDP window of between tens and hundreds of milliseconds, and the integral in Eq.(3) will give nearly zero. Periodically modulating place-cell activity will increase the growth rate as long as potentiation is dominant at short latencies, Fig.4a. Slow modulation will bias the integral in Eq.(3) toward the potentiation lobe of the STDP (Fig4a left and middle, top), while in an optimal range of frequencies the peaks and troughs of the AC will maximally capture potentiation *and* flip the sign of the depression lobe (Fig4a left and middle, middle). Finally, at higher frequencies the STDP rule undergoes multiple sign flips on a fast time scale, which again gives a near zero value for Eq.(3) (Fig.4a, left and middle, bottom). This means that the maximal growth rate of the even mode occurs for an intermediate range of frequencies: those which modulate place-cell activity on a time scale commensurate with the STDP window, Fig4a right. Note that this has nothing to do with the overall rate of plasticity, which is the same independent of the modulation frequency. That is, even if the AC is flat, recurrent synapses undergo large numbers of potentiations and depressions. Rather, periodic modulation serves to organize the structure of synaptic plasticity at the network level by preferentially strengthening connections

between place-cells with overlapping or nearby place field and weakening others. The optimal range of frequencies for growth of the even mode depends only weakly on the ratio of the width of the potentiation to that of the depression lobe, Fig.4c, but significantly on the total width, Fig.4d. Finally, allowing for triplet interactions (as opposed to just pairwise) in the STDP rule increases the overall growth rate but does not alter the range of optimal frequencies, Fig.4e. On the other hand, the theory predicts that the growth rate of the odd mode is only weakly dependent on the modulation frequency (Fig.4b) as was seen in simulations (Fig.3d), and can be understood by considering Eq.(2) with $\theta - \theta' = \pi/2$. In this case the growth rate depends on the product of the STDP rule with the cross-correlation (CC) of cells with disparate place fields. This CC is an odd function and hence the product with the STDP rule reliably gives a positive growth rate.

Sparse coding allows for the replay of activity from multiple explored tracks during SWRs

The spatio-temporal structure of SWR-like bursts in our model reflected the sequential ordering of place fields on the most recently explored track; correlation with earlier tracks was erased or greatly reduced, Fig.S3. In reality, only a fraction of place cells have well-defined place fields in any given environment, i.e. coding is sparse (Bostock et al. 1991, Fyhn et al. 2007). We incorporated this sparse-coding strategy to our model by providing place-field input to only one half of the total neurons in the network; the other half received constant input. These place cells where chosen randomly from one track to the next and place fields locations where assigned randomly as before. Therefore the overlap in the population of place-cells between any two tracks was also fifty percent. We then allowed the network to evolve by having the virtual animal explore ten distinct tracks, each for 1hr of simulation time, as before. The resulting matrix of synaptic connections was significantly correlated with the ordering of place-cells in all ten explored environments, Fig.5a. Spontaneous activity in the network revealed replay which was simultaneously correlated with several past environments even when the external input was global and entirely non-selective, Fig.5b. When the constant external drive was selective to the subset of place-cells on any given track, replay activity

was robustly correlated with activity on that track only, Fig.5c left. Also, robust replay (SC > 0.5) was found for up to 5-6 distinct tracks, Fig.5c right. Therefore, sparse coding causes the synaptic connectivity matrix to be simultaneously correlated with place field orderings from multiple tracks and allows for robust replay of activity from those tracks given appropriate inputs (cues).

Experimental evidence for a transition to replay

Pairwise reactivations of CA1 place-cells with overlapping place fields during awake SWRs improve during exploration; they are stronger during late exploration than early exploration (O'Neill et al. 2006). This holds true not only for pairwise correlations, but also reactivations of entire neuronal ensembles, at least on linear tracks (Jackson et al. 2006). More recent work has shown that significant replay events during awake SWRs in rats running along a three-arm maze emerge abruptly only after a certain number of runs (Wu and Foster 2014). These results are consistent with our model predictions. We additionally sought to directly test for a time-resolved increase in SC in neuronal data. We looked for this increase in multi-unit recordings of place-cell activity from the hippocampus of rats exploring novel, periodic tracks. We first identified cells with well-defined place fields by extracting the coefficients and phase of the first two spatial Fourier modes of their time-averaged activity as a function of the normalized distance along the track (in radians), see Fig.S7. We kept only those cells for which the ratio of the coefficients exceeded a threshold, indicating strong spatially selective activity, see *experimental procedure*s for details. We then ordered the cells according to their phase (approximately the position of peak firing). The SC of activity over the total duration of the experiment using this ordering was significantly higher than 5,000 randomly reshuffled orderings which on average gave SC = 0, see Fig.6a and b. When the medial septum was inactived via muscimol the SC did not exhibit any dynamics as a funciton of time, Fig.6c. However, once the animal recovered from the muscimol the SC using the proper phase ordering exhibited an initial ramp over the first few minutes of exploration and then remained high (significant difference between first point and all others, t-test with multiple-comparison

Bonferroni correction, average p-value $< 10^{-7}$), see Fig.6d, solid circles. This is consistent with the model result which showed a similar ramping increase when the total activity (and not just SWRs) was considered. On the other hand, the SC computed for shuffled phases showed no dynamics and remained close to zero, see Fig.6c, open squares. This indicates that there are no global changes in neuronal correlations during exploration, which could occur, for example, due to slow changes in theta modulation or neuronal excitability. Rather, there is a sustained increase in pairwise correlations only between those neurons which encode nearby places, and only when strong theta modulation is present.

Finally, we were unable to find a large enough number of significant replay events during putative SWRs in this data to test for a sharp transition, perhaps due to the relatively small number of well-defined place cells.

Discussion

Summary

We have presented a computational model of a hippocampal place-cell network in order to investigate how the exploration of novel environments shapes the patterns of recurrent synaptic connectivity. Because place-fields remap randomly from one environment to the next, the recurrent connectivity, shaped by previous learning, is initially uncorrelated with place-cell activity in a novel environment. Our major finding is that the rate at which spatial correlations in the connectivity emerge during exploration depends almost entirely on the product of the autocorrelation of place-cell activity and the STDP window. This rate is maximum when place-cell activity is periodically modulated on a time-scale commensurate with the STDP rule, which for realistic time constants yields frequencies in the theta range. Furthermore, lower and higher frequencies than theta lead to learning rates which are orders of magnitude slower. This suggests that the role of theta is to accelerate learning. Note that the overall rate of plasticity is not affected by the presence of

oscillations. The number of spike pairs, and hence the number of potentiations and depressions, depends only on the firing rates. Rather, theta oscillations generate repeated pre-post pairings in both directions, which coupled with an STDP rule with dominant potentiation at short latencies bias plasticity toward potentiating events for neurons with neighboring place fields. It is this local potentiation which drives the emergence of replay during spontaneous activity. One signature of this mechanism is a continuous increase in the pairwise cross-correlation in the activity of neighboring place-cells leading up to a critical time. We have found evidence consistent with this by analyzing the activity of simultaneously recorded hippocampal place cells in a rat during the exploration of a novel track.

The assumption of plasticity only at recurrent synapses

In our model we have assumed that plasticity occurs only in the recurrent excitatory synaptic connections, and not in the feed-forward inputs. Therefore we also assume that the place-field input, which peaks at the spatial position of the virtual animal at given moment in time, is itself stable. In fact, consistent with this assumption, it seems most place cells are active from the outset of exploration of a new environment, although see (Hill, A.J. 1978, Frank 2004, Monaco 2014). Furthermore cells tend to exhibit only subtle changes in the size and shape of their place fields over time (Mehta 1997, Mehta 2000), also consistent with our model, see Fig.S2. On the other hand, it has been shown in area CA1 that some place-cells exhibit place-fields only after several minutes of exploration (Frank et al. 2004, Frank et al. 2006). Recent intracellular recordings indicate that appearance of these "hidden" place-fields requires the coincidence of active dendritic propagation and synaptic input via the Schaeffer collaterals (Lee et al. 2012, Bittner et al. 2015). It may be that this mechanism for place-cell "generation" is particularly salient in cells of area CA1 by virtue of being uniquely positioned to compare and integrate both entorhinal and hippocampal inputs. In any case the strongly recurrent nature of the network we study may make it a more relevant model for circuits in area CA3. Nonetheless we would expect that changes in spiking activity arising in CA3 due to plasticity in the recurrent connectivity, as in our model, would be reflected in analogous

changes in the spiking activity of CA1 cells due to the direct inputs via the Schaeffer collaterals. Indeed, the data we have analyzed from place cells of area CA1 show increase sequential correlation as predicted by our recurrent model.

Remapping of place fields for different directions of motion on the same track

Hippocampal place cells actually exhibit global remapping of their place fields depending on the direction of motion of the animal on linear tracks (McNaughton et al. 1983, Muller et al. 1994). This is perhaps not surprising given that the behaviorally relevant information for the animal is not just the position along the track but also the way in which it is facing; e.g. this determines how far away a potential reward is, if located at one or both ends of the track. Studies using periodic tracks have shown no such global remapping, but rather some degree of rate remapping (Schwindel et al. 2016), i.e. the direction of motion affects the shape and amplitude of place fields, but not their position. In the data we have analyzed there is very weak remapping, see Figs.S7c and , and so the assumption of invariance of place field to direction of motion is a good one. In our model we have made this assumption. The consequence of this is that while exclusively clockwise (CW) or counter-clockwise (CCW) motion will lead to highly asymmetric recurrent connectivity Fig.1, exploration of both directions for equal times will give rise to symmetric connectivity, see Fig.S4. In practice any trajectory over a finite amount of time will have a directional bias; in the data we have analyzed the rat spends 54% of the time moving CW and 46% CCW (DOUBLE CHECK THIS), and this will necessarily lead to asymmetries in the connectivity. In linear tracks, due to the global remapping such asymmetries should be even more pronounced.

Forward versus backward replay

The inevitable asymmetry in the recurrent connectivity of our model place-cell network due to plasticity during exploration strongly biases spontaneous activity. On a periodic track this replay would be exclusively CW or CCW depending on the corresponding bias in motion during exploration, while learning on a linear track would always produce forward replay. Previous work has shown that perfectly symmetric connectivity can give rise to both forward and backward replay

in equal measures, due to spontaneous symmetry breaking of activity (Romani and Tsodyks 2015). We would argue that such symmetric connectivity is not robust for the reasons given above, although we cannot rule out the existence of homeostatic mechanisms which would conspire to make it so. Rather, we propose here an alternative mechanism for generating backward replay given asymmetric connectivity based on local sensory input. Specifically, if global input to our model network is constant then replay occurs only in one direction. However, if a localized bump of input is provided to the network, synapses to downstream neurons (in the sense of the asymmetric bias of the connectivity) become rapidly depressed. This prevents the spontaneous activity from propagating forward and forces it to propagate backward, see Fig.S5. In fact, in experiment, when local spatial input is absent, e.g. when the animal is sleeping in a rest box, forward replay is predominant, see (Roumis and Frank 2015). On the other hand, both backward and forward replay are observed when the animal is awake but quiescent on a given track. This is precisely when locally sensory cues are available to the animal, and could potentially shape spontaneous replay events. In fact recent work shows that some neurons in area CA2 fire more strongly during awake quiescence than during exploration (Kay et al. 2016); they may be providing information regarding local sensory cues.

Robustness to changes in the plasticity model and to the presence of spike correlations

Here we have considered a simple phenomenological model of plasticity which depends on the timing of spike pairs. Taking into account spike triplets as opposed to only pairs does not alter our findings qualitatively, see Fig.4e. It remains to be studied how more realistic voltage- or calcium-based plasticity rules interact with the theta-modulation to affect learning in recurrent networks (Clopath et al. 2010, Graupner and Brunel 2012), although at a single-synapse one can find qualitatively similar regimes for an array of plasticity rules in the presence of pre- and post-synaptic oscillations (Albers et al. 2013).

Our results clearly do not depend on the actual spike timing since our model neurons generate spikes as Poisson processes; rather, all lasting changes in the connectivity are due to time-varying

modulations of the firing rates. In fact, recent work with a spiking neuron model suggests that such modulations in the firing rate, as opposed to exact spike timing, are sufficient to explain the effect of plasticity from STDP and more realistic calcium-based plasticity rules in general (Graupner et al. 2016). In any case the contribution of pairwise spike correlations to the evolution of the recurrent connectivity can be formally taken into account in Eq.4, i.e. via its affect on the auto-correlation of place-cell activity.

Other models of place-cell activity

Recurrent network models for place-cell activity provide a parsimonious explanation for electro-physiological phenomena associated with exploratory behavior as well as the generation of sharp-wave bursts during awake quiescence and sleep (Tsodyks et al. 1996, Shen and McNaughton 1996, Cutsuridis and Hasselmo 2011, Romani and Tsodyks 2015, Jahnke et al. 2015). In (Romani and Tsodyks 2015) SWRs were generated spontaneously by virtue of the spatial modulation of the recurrent connectivity, which drives an instability to traveling waves in the absence of place-field input. The presence of short-term depression modulates the amplitude of the waves, leading to narrow bursts. This is the same mechanism we have used here. Alternatively, recent work with a biophysically detailed spiking network model focused on the role of nonlinear dendritic integration on the generation of replay during SWRs (Jahnke et al. 2015). In that work the authors found that the exploration of a virtual linear track in the presence of pairwise STDP lead to highly asymmetric connectivity; this could generate replay activity given a sufficiently synchronous external input which recruited nonlinear dendritic events. In our work, we have sought to explain the replay as an emergent phenomenon which depends only on the network-wide organization of synaptic structure. In doing so we have considered a simple stochastic firing rate model which allowed us to fully characterize how interplay between the STDP rule and the place-cell activity affects learning. Nonetheless, a detailed reproduction of the phenomenology of SWRs certainly requires mechanisms we have not included here.

Spatial learning

It seems reasonable that the learning of tasks which depend on spatial information require the formation of an internal representation of the relevant environment. This is the process we have studied here. While we have not modeled any particular cognitive task, we propose that the network-wide organization of synaptic structure, in order that it be in concordance with the place-field distribution of place cells, should be a necessary step in spatial learning tasks. Our results suggest that this process is dramatically sped up by modulating place-cell activity in the theta range, which is one possible role of this prominent rhythm.

Experimental Procedures

Model description

We simulated a model of n excitatory neurons with global inhibitory feedback. The firing rate of a neuron i evolved according to $\tau \dot{r}_i = -r_i + \phi\left(\frac{1}{n}\sum_{j=1}^{n} \widetilde{w}_{ij} r_j x_j + I_i(t)\right)$ where $\widetilde{w}_{ij} = w_{ij} - w_I$ is the effective connectivity from a cell j to a cell i and consists of a recurrent excitatory synaptic weight and a global inhibitory feedback. The synaptic depression variable x_i obeys $\dot{x}_i = \frac{(1-x_i)}{\tau_x} - U_0 r_i x_i$. The external input is $I_i(t) = I_0 + I_{PF}\cos(\theta_i - \theta(t))(1 + I_{theta}\cos(2\pi f t))$ where θ_i is the place field position of neuron i in radians, $\theta(t)$ is the position of the "virtual animal", I_{PF} was the amplitude of the place-field input. To model theta-modulation we multiplied the place field input by a periodic signal with frequency f and amplitude I_{theta} . This type of multiplicative modulation is seen in intracellular recordings in-vivo, e.g. see Figs.1 and 5 in (Harvey et al. 2009) and Fig.4 in (Lee et al. 2012). A cell i generates a spike in a time interval dt with probability $r_i dt$. Plasticity occurs for every spike pair between cells i and j: $w_{ij} \to w_{ij} + \Delta w_{ij}$ where $\Delta w_{ij} = A_+ \exp(-T/\tau_+)$ if $T = t_i - t_j > 0$ else $\Delta w_{ij} = -A_- \exp(T/\tau_-)$. We implement this as in (Pfister and Gerstner 2006). We furthermore set a minimum value for

synapses at zero and a maximum value w_{max}. We model distinct tracks by spacing place fields uniformly and assigning them randomly to cells. Therefore the ordering of θ_i s from track to track are random and uncorrelated.

Model analysis

How theta modulation and STDP shape recurrent connectivity

We consider a continuum limit of the network neglecting synaptic depression, in which the firing rate equation can be written $\tau \dot{r}(\theta,t) = -r(\theta,t) + \phi\left(\frac{1}{2\pi}\int_{-\pi}^{\pi} d\theta' w(\theta,\theta') r(\theta') + I(\theta,t)\right)$, (5)

$I(\theta,t) = I_0 + I_{PF}\cos(\theta - vt)(1 + I_{theta}\cos(2\pi ft))$,

and assuming the plasticity is slow compared to the rate dynamics we can write

$\dot{w}(\theta,\theta',t) = \int_{-\infty}^{\infty} dT\, STDP(T) r(\theta,t) r(\theta',t+T)$. If we take the transfer function ϕ to be linear we can solve the equations self-consistently for the rates and the weights where

$r(\theta,t) = r_0 + r_{even}\cos(\theta - vt) + r_{odd}\sin(\theta - vt)$ and $\dot{w} = \dot{w}_0 + \dot{w}_{even}\cos(\theta - \theta') + \dot{w}_{odd}\sin(\theta - \theta')$. We choose a balanced STDP to avoid unbounded growth or decay to zero in the mean weights, i.e.

$A_+\tau_+ = A_-\tau_-$. Then $\dot{w}_0 = 0$,

$$\dot{w}_{even} = \alpha I_{PF}^2\left(\frac{1}{1+\tau_+^2 v^2} - \frac{1}{1+\tau_-^2 v^2}\right) + \frac{\alpha}{2} I_{PF}^2 I_{theta}^2\left(\frac{1}{1+\tau_+^2 v_+^2} - \frac{1}{1+\tau_-^2 v_+^2}\right) + \frac{\alpha}{2} I_{PF}^2 I_{theta}^2\left(\frac{1}{1+\tau_+^2 v_-^2} - \frac{1}{1+\tau_-^2 v_-^2}\right), \quad (6)$$

$$\dot{w}_{odd} = \alpha I_{PF}^2\left(\frac{v}{1+\tau_+^2 v^2} - \frac{v}{1+\tau_-^2 v^2}\right) + \frac{\alpha}{2} I_{PF}^2 I_{theta}^2\left(\frac{v_+}{1+\tau_+^2 v_+^2} - \frac{v_+}{1+\tau_-^2 v_+^2}\right) + \frac{\alpha}{2} I_{PF}^2 I_{theta}^2\left(\frac{v_-}{1+\tau_+^2 v_-^2} - \frac{v_-}{1+\tau_-^2 v_-^2}\right), \quad (7)$$

where $v_+ = v + 2\pi f$ and $v_- = v - 2\pi f$ and $\alpha = \frac{A_+\tau_+}{(1 - w_{even}/2)^2 + (w_{odd}/2 - \tau v)^2}$. For modulation frequencies above 1Hz we have $\epsilon = \frac{v}{2\pi f} \ll 1$ and the growth rates simplify to give

$$\dot{w}_{even} = \alpha I_{PF}^2 I_{theta}^2 \left(\frac{1}{1+4\pi^2\tau_+^2 f^2} - \frac{1}{1+4\pi^2\tau_-^2 f^2} \right) \quad (8)$$

$$\dot{w}_{odd} = \alpha I_{PF}^2 \left(\frac{\nu}{1+\tau_+^2 \nu^2} - \frac{\nu}{1+\tau_-^2 \nu^2} \right) \quad (9).$$

This shows explicitly that the growth of the even mode is strongly dependent on the amplitude and frequency f of periodic modulation, while that of the odd mode is independent of these to leading order. See *supplementary information* for a detailed derivation. Interestingly, in a single-synapse model of STDP, maximal potentiation is achieved when the pre- and post-synaptic firing rates are modulated at a frequency commensurate with the STDP window. However, the growth rate of the synapse depends not a difference of terms, as in Eq.8, but rather on a product. Therefore there is no sensitivity to the exact shape of the STDP window, i.e. dominant potentiation or dominant depression at short latencies, or perfectly anti-symmetric.

The emergence of replay

When the virtual animal is first exposed to a novel track, the recurrent connectivity w is initially unorganized with respect to the ordering of place fields, i.e. $w_{even} = w_{odd} = 0$. We have seen in the last section how exploration and theta-activity cause these modes to grow in time. If we remove the place-field input at different points in time during exploration, i.e. $I = I_0$, we can predict when replay should first emerge. To do this we consider the stability of the homogeneous state in firing rates to spatio-temporal perturbations, i.e. $r = r_0 + \delta r_{even} \cos(\theta) e^{\lambda t} + \delta r_{odd} \sin(\theta) e^{\lambda t}$, where δr_{even} and δr_{odd} are small. Plugging this into Eq.(4) yields the characteristic equation for the eigenvalue λ , $\tau^2\lambda^2 + 2\tau\lambda\left(1 - \frac{w_{even}}{2}\phi'\right) + \left(1 - \frac{w_{even}}{2}\phi'\right)^2 + \frac{w_{odd}^2}{4}(\phi')^2 = 0$. From this we can see that there can be no instabilities to stationary bumps $\lambda = 0$ but that traveling waves can emerge ($\lambda = i\omega$) when $w_{even} = \frac{2}{\phi'}$ and they have velocity $\omega = \frac{w_{odd}}{2\tau}\phi'$. Therefore, exploration will

cause the even mode of the connectivity to grow until a critical value is reached, at which point traveling waves correlated with activity on the novel track will emerge spontaneously. This argument still holds when we include the effect of synaptic depression, although the calculation is more involved, see *supplementary information.*

Reduction of the plasticity rule to integral of STDP times AC of place cell activity

To go from Eq.3 to Eq.4 we must expand the product of the firing rates. For the sake of illustration we will take and leave the more general case, with an odd component, for the *supplementary information*. The product

$$r(\theta,t)r(\theta',t+T)=r_0^2+r_0r_1\left(\cos(\theta-vt)+\cos(\theta'-vt-vT)\right)+r_1^2\cos(\theta-vt)\cos(\theta'-vt-vT)$$

$$=r_0^2+\frac{r_1^2}{2}\cos(\theta-\theta'-vT)+f(t)=r_0^2+\frac{r_1^{\,2}}{2}\left(\cos(\theta-\theta')\cos(vT)+\sin(\theta-\theta')\sin(vT)\right)+f(t)\quad,$$

where f(t) is a time-periodic function with period π/v which could range from 100ms to a few seconds depending on whether v is the velocity of the animal (no modulation) or $v_+=v+2\pi f$, i.e. theta-modulation. Note that this time-scale is fast compared to the time-scale of plasticity which occurs on the order of minutes. This separation of time-scales allows us to average both sides of Eq.3; if we furthermore set $\theta=\theta'$ we obtain Eq.4.

Analysis of data from hippocampal recordings

We analyzed spiking data from 124 simultaneously recorded hippocampal neurons of area CA1 in a rat during 35 minutes of exploration of a novel, periodic track, see (Wang et al 2015) for detailed methods on the data collection. In order to identify place cells we first calculated a rate map on the track for each neuron, and then linearized the rate to obtain a one-dimensional place field. We then fit the place field with the function $r=r_0+r_1\cos(x-\psi_1)+r_2\cos(2(x-\psi_2))$ using the built-in ?? function in matlab, thereby extracting the coefficients of the Fourier modes and their phases. We also extracted the coefficients via Fast-Fourier Transform and found identical results. We did not find any cells for which r_2 was significant compared to the first two coefficients and hence limited our analysis to the first two coefficients and the phase, i.e. the place field was centered at

ψ_1 . We excluded all cells for which the ratio r_1/r_0 was below a fixed threshold. Specifically, we set the threshold to select a given percentage of cells with the highest ratio. For Fig.6 we took 15% of cells (both for the data set with and without muscimol). Varying this percentage from 10% to 30% gave qualitatively similar results for all cases, see Fig.S9.

We calculated the average SC over the entire experiment by ordering the neurons according to their place field phases, extracted in the analysis described above (~0.1, see Fig.6). We compared this to 5000 trials in which we randomly reshuffled the phases. This lead to a Gaussian-like distribution with mean near zero ($6x10^{-3}$) and standard deviation 0.016, see fit in Fig.6A. The SC with the proper ordering was hence highly significantly different from other orderings. We also calculated the SC during 3.5 minute epochs of time, yielding 10 time points.

Methods and parameters for figures

Fig.2 Parameters for simulations are:

$$\tau = 10\,ms \qquad A_+ = 0.1, \tau_+ = 20\,ms, A_- = 0.1/3, \tau_- = 60\,ms.$$

$$I_0 = 3\,Hz, I_1 = 22\,Hz, I_{theta} = 1, f = 5\,Hz, v = 1\,grad/sec \qquad U_0 = 0.0008, \tau_R = 800\,ms \qquad w_I = 45$$

The weight matrix W was trained by simulating exploration on 10 distinct linear tracks for one hour each. The value of W was taken as a constant 40 for all synapses at the beginning of exploration of the first environment and the maximum possible weight was set at $w_{max} = 80$. See Fig.S1 for details. For the connectivity profiles, we calculate the mean synaptic weight between pairs of neurons with a difference in place field location $\Delta\theta$ at the given time of the snapshot. There are no autapses, but the curve is made continuous at $\Delta\theta = 0$ by interpolating between adjacent points. The activity profiles are averages over the five seconds preceding the given time of the snapshot, e.g. for 1 min it is the average activity between 55 and 60 seconds. The average is taken with respect to the peak of the place field input, at $\theta = 0$.

To generate the figure in (b) we calculate the SC in 1 sec bins during the entire simulation. Every 180 seconds there is a 3 second period during which the external theta-modulated place field input is removed, in order to model awake quiescence. Activity during this period is spontaneous and is

considered to be SWRs (black circles). During SWR activity we only calculate the SC for the second and third seconds because it takes some time for the place field activity to die away and the SWR activity to emerge, e.g. SC for SWRs is calculated for seconds 181-182 and 182-183 but not for 180-181. For the simulation in (b) when the place field input is removed, the external input is set to a constant value of $I_0=0$ while for (b) and (c) it is set at $I_0=-0.7$. Changing I_0 does not significantly affect the value of SC, only the degree of burstiness of the spontaneous activity, see Fig.S4.

Fig.3 The curves shown in (a) are a cartoon meant to illustrate how the recurrent connectivity can be decomposed into a spatial Fourier series which include even (cosine) and odd (sine) terms. The amplitude of an even term (its coefficient) can lead to a transition in the network dynamics when it reaches a critical value. (b) and (c) The coefficients a1 and b1 are estimated in the following way. The mean recurrent connectivity is calculated as in Fig.2 at a given time during the simulation. The curve $a_0+a_1\cos(\theta)+b_1\sin(\theta)+a_2\cos(2\theta)+b_2\sin(2\theta)$ is then fit to the connectivity using the built-in nonlinear curve fitting function of the software Grace, providing the coefficients. In all cases shown this function provided an excellent fit. Parameters in (d) and (e) are the same as in Fig.2. The parameters in (f) are the same as in Fig.2 with the sole exception of the STDP rule. For the anti-symmetric case (left) the parameters are $A_+=0.1, \tau_+=40\,ms$, $A_-=0.1\,\tau_-=40\,ms$. while for the case on the right they are $A_-=0.1, \tau_+=20\,ms$, $A_+=0.1/3, \tau_-=60\,ms$.

Fig.5 The virtual animal has explored ten distinct environments for one hour each. In each environment, one half of the neurons are modeled as place cells (randomly assigned place field location from uniform distribution around the track) and the other half receive only constant background input. Plasticity occurs via STDP as before between all cell pairs. Place cells in any given environment are chosen randomly with equal probability; hence the overlap in place-cell representation between any two environments is on average one half. The number of neurons is N = 200, so that in any environment there are still 100 place cells, as in previous simulations. (a) The matrix shows the value of the recurrent synaptic weights after allowing them to evolve according to

a pairwise STDP rule for ten hours of simulation time, one hour for each of ten different "environments".

Author Contributions

P.T.. ran numerical simulations. B.R. analyzed data. A.R. developed theory, ran simulations, analyzed data and wrote the paper with input from the other authors.

Acknowledgements

References

Albers, C., Schmiedt, J. T. and Pawelzik, K. R. (2013). Theta-specific susceptibility in a model of adaptive synaptic plasticity. Front. Comp. Neurosci. 7(170) 1-18.

Allen, M.T., Padilla, Y. and Gluck, M.A. (2002). Ibotenic acid lesions of the medial septum retard delay eyeblink conditions in rabbits (oryctolagus cuniculus). Behav. Neurosci. 116:733-738, 2002.

Bittner, K. C., Grienberger, C., Vaidya S. P., Milstein, A. D., Macklin, J. J., Suh, J., Tonegawa, S. and Magee, J. C. (2015). Conjunctive input processing drives feature selectivity in hippocampal CA1 neurons. Nat. Neurosci. 18, 1133-1142.

Bostock, E., Muller, R. U. and Kubie, J. L. (1991). Experience-dependent modifications of hippocampal place cell firing. Hippocampus 1(2): 193-206.

Brandon, M.P., Koenig, J., Leutgeb, J.K. and Leutgeb, S. (2014). New and distinct hippocampal place codes are generated in a new environment during septal inactivation. Neuron, 82: 789-796.

Carr, M.F., Jadav, S.P. and Frank, L.M. (2011). Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval. Nat. Neurosci. 14:147-153.

Clopath, C., Büsing, L., Vasilaki, E. and Gerstner W. (2010). Connectivity reflects coding: a model of voltage-based STDP with homeostasis. Nat. Neurosci. 13(3) 344-352.

Cutsuridis, V. and Hasselmo, M. (2011). Spatial memory sequence encoding and replay during modeled theta and ripple oscillations. Cogn Comput. 3:554-574.

Foster, D.J. and Wilson, M.A. (2006). Reverse replay of behavioral sequences in hippocampal place cells during the awake state. Nature 440:680-3.

Frank, L. M., Stanley, G. B. and Brown, E. N. (2004). Hippocampal plasticity across multiple days of exposure to novel environments. J. Neurosci. 24:7681-7689.

Frank, L. M., Brown, E. N. and Stanley, G. B. (2006). Hippocampal and cortical place cell activity: implications for episodic memory. Hippocampus 16:775-784, 2006.

Fyhn, M., Hafting, T., Treves, A., Moser, M.-B. and Moser, E.I. (2007). Hippocampal remapping and grid realignment in entorhinal cortex. Nature, 446, 190-194.

Graupner, M. and Brunel, N. (2012). Calcium-based plasticity model explains sensitivity of synaptic changes to spike pattern, rate, and dendritic location. PNAS 109(10) 3991-3996.

Graupner, M., Wallisch, P. and Ostojic, S. (2016). Natural firing patterns imply low sensitivity of synaptic plasticity to spike-timing compared to firing-rate. J. Neurosci. 36(44): 11238-11258.

Harvey, C. D., Collman, F., Dombeck, D. A. and Tank, D. W. (2009). Intracellular dynamics of hippocampal place cells during virtual navigation. Nature 461, 941-946.

Jackson, J. C., Johnson, A. and Redish, A. D. (2006). Hippocampal sharp waves and reactivation during awake states depend on repeated sequential experience. J. Neurosci. 26(48): 12415-12426.

Jahnke, S., Timme, M. and Memmesheimer, R.-M. (2015). A unified dynamic model for learning, replay and sharp-wave/ripples. J. Neurosci. 35(49): 16236-16258.

Karlsson, M.P. and Frank, L.M. (2009). Awake replay of remote experiences in the hippocampus. Nat. Neurosci. 12(7):913-8.

Kay, K., Sosa, M., Chung, J. E., Karlsson, M., Larking, M. C. and Frank, L. M. (2016). A hippocampal network for spatial coding during immobility and sleep. Nature 531, 185-190.

Kempter, R., Gerstner, W. and van Hemmen, J.L. (1999). Hebbian learning and spiking neurons. Phys. Rev. E 59(4) 4498-4514.

Kubie, J. L. and Muller, R. U., (1991). Multiple representations in the hippocampus. Hippocampus, 1(3), 240-242.

Lee, A.K. and Wilson, M.A. (2002). Memory of sequential experience in the hippocampus during slow wave sleep. Neuron 36: 1183-1194.

Lee, D., Lin, B. and Lee, A. K. (2012). Hippocampal place fields emerge upon single-cell manipulation of excitability during behavior. Science 337(6096): 849-853.

McNaughton, B. L., Barnes, C. A. and O'Keefe, J. (1983). The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. Exp. Brain Res. 52, 41-49.

Muller, R. U. and Kubie, J. L. (1987). The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. J. Neurosci., 7(7), 1951-1968.

Muller, R. U., Bostock, E., Taube, J. S. and Kubie, J. L. (1994). On the directional firing properties of hippocampal place cells. J. Neurosci. 14, 7235-7251.

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

O'Keefe, J. (1976). Place units in the hippocampus of the freely moving rat. Exp. Neurol. 51, 78-109.

O'Neill, J., Senior, T. and Csicsvar, J. (2006). Place-selective firing of CA1 pyramidal cells during sharp wave/ripple network patters in exploratory behavior. Neuron 49, 143-155.

Pfister, J-P and Gerstner, W. (2006). Triplets of spikes in a model of spike timing-dependent plasticity. J. Neurosci. 26, 9673-9682.

Romani, S. and Tsodyks, M. (2015). Short-term plasticity based network model of place cells dynamics. Hippocampus 25:94-105.

Roumis, D. K. and Frank, L. M. (2015). Hippocampal sharp-wave ripples in waking and sleeping states. Curr. Op. in Neurobiol. 35: 6-12.

Schwindel, C. D., Navratilova, Z., Ali, K., Tatsuno, M. and McNaughton, B. L. (2016). Reactivation of rate remapping in CA3. J. Neurosci. 36(36): 9342-9350.

Shen, B. and McNaughton, B. L. (1996). Modeling the spontaneous reactivation of experience-specific hippocampal cell assemblies during sleep. Hippocampus 6:685-692.

Song, S., Miller, K. D. and Abbott, L. F. (2000). Competitive Hebbian learning through spike-

timing-dependent synaptic plasticity. Nat. Neurosci. 3(9) 919-926.

Tsodyks, M. V., Skaggs, W., Sejnowski, T. and McNaughton B. (1996). Population dynamics and theta rhythm phase precession of hippocampal place cell firing: a spiking neuron model. Hippocampus 6, 271-280.

Wang, Y., Romani, S., Lustig, B., Leonardo, A. and Pastalkova, E. (2015). Theta sequences are essential for internally generated hippocampal firing fields. Nat. Neurosci. 18, 282-288.

Wilson, M.A. and McNaughton, B.L. (1994). Reactivation of hippocampal ensemble memories during sleep. Science 265: 676-679.

Wu, X. and Foster, D. J. (2014). Hippocampal replay captures the unique topological structure of a novel environment. J. Neurosci. 34(19): 6459-6469.

Figure 1. A schematic description of the model.

(a) We model the sequential exploration of a number of distinct linear tracks. (b) The network consists of N place cells. The ordering of the place fields is randomly remapped on each track. Therefore, if the cells are properly ordered in any given environment the place field input is represented by a spatially localized bump of activity (upper left and lower right). Sequential activity on a familiar track would look random given an ordering on a novel track, and vice versa (upper right and lower left).

Figure 2. Spatial exploration gives rise to plasticity-dependent emergence of spatio-temporal structure in SWR activity. (a) Top row: Snapshots of the recurrent excitatory connectivity during the simulation of an extended exploration of a novel environment. The recurrent synaptic weights become modulated in the space of the novel track and stabilize. Bottom row: Snapshots of the place cell activity zeroed at the maximum of the place field input. (b) The mean cross-correlation of the activity of place cells with adjacent place fields on a novel track during exploration (called here sequential correlation SC). red: SC during active exploration (theta-activity), i.e. cells are driven by external place field input, black: SC during rest periods, i.e. place field input is removed and intrinsic bursting occurs ("SWRs"). Note the sharp transition in the SC of SWR activity. (c) Top: SWRs at early times do not exhibit sequential replay of activity from the novel track. top: average input to place cells before (blue star), during (black star) and after period of "quiet wakefulness". bottom: space-time plots of the place cell input and firing rate. Note the disordered spatio-temporal structure of the burst activity. Bottom: After a critical transition time SWRs exhibit sequential replay of activity from the novel track. Note the sequential structure of the burst activity.

Figure 3. The transition in the structure of SWR activity occurs when the symmetric component of the recurrent connectivity reaches a critical value. This mode grows only for STDP rules with dominant potentiation, and grows fastest when periodic modulation is in the theta range. (a) The

profile of the recurrent connectivity at any point in time can be decomposed into a series of even (cosine) and odd (sine) modes. Only the first two modes are shown. (b) In the simulations from Fig.2 both even and odd modes grow in time as the recurrent connectivity is shaped by synaptic plasticity. (c) As the feedback inhibition is reduced from $w_I = 65$ (black) to 55 (orange) to 45 (green), the transition shifts to earlier times (40, 20 and 10 mins respectively). In all cases the value of the even mode a_1 reaches precisely the same critical value at the time of the transition. In (b) and (c) the amplitude of the modes is normalized by the mean synaptic weight. (d) The amplitudes of the even (solid circles) and odd (open squares) modes of the recurrent synaptic connectivity after 1hr of exploration, as a function of the frequency of oscillatory modulation. While the odd mode is relatively insensitive to the modulation frequency, the even mode grows maximally over a range of approximately 1-10Hz and it is strongly attenuated at lower and higher frequencies. Note the logarithmic scale. (e) The degree of sequential correlation of spontaneous bursting (SWR) after 1hr of exploration as a function of the frequency of oscillatory modulation. The SWR activity is strongly correlated with the topology of the current environment only in simulations where the frequency lies between 2-8Hz. Inset: The time at which a transition in the SC takes place as a function of frequency, for times up to 1hr. (f) Left: A purely anti-symmetric STDP rule leads to recurrent connectivity which has only odd Fourier modes (red squares). There is no increase or transition in the SC of SWR in these simulations even after 1hr. Right: An asymmetric STDP rule with depression dominating at short latencies leads to recurrent connectivity with a negative amplitude of the even mode, i.e. recurrent excitation is weaker between pairs of place cells with overlapping place fields than those with widely separated place fields. In this case SWR is actually completely suppressed (not shown).

Figure 4. Analysis of a linear firing rate model with STDP reveals the mechanism by which the STDP rule and periodic modulation affect recurrent connectivity. (a) The growth rate of the symmetric mode is proportional to the integral of the convolution of the autocorrelation (AC) of

place cell activity with the STDP window. Left: The autocorrelation (AC) of place cell activity overlaid on the STDP window with dominant potentiation at short latencies. Middle: The convolution (pointwise multiplication) of the STDP window and the AC. Right: The growth rate of the symmetric (cosine) mode of the connectivity as a function of frequency. When the activity is modulated at 1Hz (top), the convolution returns nearly the original STDP window, which being balanced yields near-zero growth rate (i). For 5Hz the potentiating lobe is maintained and some of the depression lobe changes sign and becomes potentiating leading to higher growth rate (ii). For 20Hz the STDP window undergoes sign reversals at a rate faster than the STDP lobes, meaning the integral is again near zero (iii). (b) The growth rate of the even (cosine) and odd (sine) spatial Fourier coefficients as a function of the frequency of external modulation (black curve is the same as in (a) except on a log-log scale). Inset: The growth rate of the even mode normalized by its value for no periodic modulation. Note that modulation in the theta range increases the growth rate by three orders of magnitude. STDP parameters are $A_{-}=0.1, \tau_{+}=20\,ms, A_{+}=0.1/3, \tau_{-}=60\,ms$. (c) Increased growth rate in the theta range does not require fine tuning. (d) The frequency at which growth is maximal depends on the overall width of the STDP window. Broader windows favor slower frequencies. (e) An STDP triplet rule increases the growth rate at all frequencies compared to the pairwise rule, but does not significantly shift the optimal frequency range.

Figure 5. Sparse coding allows for the replay of multiple previously experienced environments. (a) The synaptic weight matrix of the recurrent excitatory connections between neurons in the network. The matrix is shown after all ten environments have been explored; each panel shows the matrix with cells ordered according to the corresponding environment. Note that while the correlation of the weights decreases the older the environment, it does not vanish. The color code is blue (weak) to yellow (strong) and the cell chosen to be place cells in any given environment are ordered in the bottom left quadrant of the matrix. (b) and (c) Spontaneous activity in the network as a function of the strength of external input, i.e. replay activity. (b) Spontaneous activity when all

neurons receive equal, constant external input, called global input. top: The mean firing rate in the network, bottom: The SC of the replay activity with different environments. The SC is highest for the most recently explored environment, but still non-zero for older ones, see the space-time plot of the same replay activity shown with ordering corresponding to environment 10 to 8 (top to bottom). (c) Spontaneous activity when external input selectively targets those neurons which where place cells in any given environment. Left: Firing rate, Right: The SC of the replay with different environments given different selective stimulations. N = 200 neurons in all simulations.

Figure 6. The SC of place cells in rat hippocampus during the exploration of a novel track shows an initial increase and plateau. (a) Firing rate profiles from the best 15% of place cells (n=19) with an illustration of how cells can be ordered according to their phases. (b) The SC calculated over the entirety of the experiment (35 min) given the correct ordering (red dot) and for 5000 reshuffled ordering. (c) The time course of SC given the proper ordering (red symbols) shows an initial ramp and then a plateau. Error bars are S.E.M. The time course of SC for the reshuffled orderings shows no dynamics. Error bars are the standard deviation over the 5000 shuffled trials. (d) Same as in (c) for the rat with muscimol injection in the medial septum. There is no increase in SC during the experiment. Total number of cells are n = 124 with no muscimol and 128 with muscimol.

Figure S1. The recurrent connectivity reaches a stationary state after exploration of 10 environments. (a) The mean synaptic weight of recurrent excitatory connections as a function of time for 10 consecutive simulations of 1hr each. The initial condition was all the weights identical and set to a value of 40. After several environments the temporal evolution of the recurrent weights becomes stereotyped, i.e. the learning process is identical from one environment to the next. (b) the distribution of synaptic weights at the end of the 10hrs of simulation. About 10% of synaptic weights are maximally depressed to zero, 40% are maximally potentiated (threshold set to 80) and 50% take on intermediate values. As new environments are explored the distribution remains fixed

but individual synapses change in time.

Figure S2. Place fields shift and become skewed during the exploration of a novel environment. (a) Place field shift as a function of time. (b) Place field skewness as a function of time. (c) Place field width as a function of time. The place field shift is calculated as the difference between the position of the maximum of the place field input and the maximum of the place cell activity. Therefore, if the place cell activity leads the place field input, this indicates that the place field has effectively shifted *backwards*, i.e. the place cell spikes most strongly before the center of its place field (as defined at an earlier time) is reached. The skewness is calculated as (mean-mode)/width of the average place cell activity as defined in (a). This is known as Pearson's first skewness coefficient. The width is calculated at half maximum amplitude. The place-cell activity is taken from the simulation from Fig.2.

Figure S3. The transition in SC corresponds to a switch in the spatio-temporal structure of spontaneous traveling wave (TW) activity. (a) SC during SWR events during the simulation shown in Fig.3. Circles: SC of activity with ordering of neurons according to their place fields in the currently explored environment. Squares: SC using ordering from previously explored environment (10th of the 10 environments explored for 1hr each). (b) Space-time plots of the spontaneous activity at 3min (before transition) and at 48min (after transition). Note the switch in the spatio-temporal structure of the TW from reflecting the topology of the familiar to that of the novel environment.

Figure S4. Transitions in the SWR activity are insensitive to the exact trajectory of the animal. (a) left: The position of the virtual rat as a function of time over one minute. right: The velocity as a function of time. The position was modeled as a random walk process with a 10sec correlation time constant. (b) The SC during theta activity (red squares) and SWRs (black circles) as a function of

time. The right-hand panel shows three snapshots of the recurrent connectivity at 1m, 30m and 1hr (top to bottom). All parameters are the same as in Fig.2 with the exception of the motion of the virtual animal (place-cell input). (c) Same as in (b) but with a perfectly anti-symmetric STDP rule with parameters $A_+ = 0.1, \tau_+ = 40\,ms$, $A_- = 0.1\,\tau_- = 40\,ms$. The position of the animal ϕ in radians is found by solving $\dot{\phi} = v$ where the velocity v $\tau_{OU}\dot{v} = -v + \sigma_{OU}\xi$ follows an Ohrnstein-Uhlenbeck process with time constant $\tau_{ou} = 10\,sec$ and amplitude $\sigma_{ou} = 0.002\,rad/s$.

Figure S5. Forward replay occurs spontaneously, but backward replay requires location-specific input. (a) The firing rate r (top), short-term synaptic depression variable x (middle) and external input to place cells (bottom) during a period of "sleep" or "quiet wakefulness", i.e. in the absence of theta-modulated place field input. Spontaneously occurring bursts (SWRs) always travel forward when the input is globally homogeneous, reflecting the asymmetry in the underlying recurrent connectivity. However, a strong location-specific input (just after 500ms and 5 sec) can transiently depress the synapses in downstream neurons, facilitating the propagation of activity backwards. Right: blow-up of activity showing forward and backward replay. (b) Raster plot generated from the same simulation as in (a). The external input is $I_i = -0.8\,Hz$ except for the two spatially-modulated inputs which are presented for 50ms each and have the form $I_i = 40\left(1 + \cos\left(\theta_i\right)\right) Hz$.

Figure S6. The degree of "burstiness" of spontaneous activity can be modulated by a global external input. (a) The temporal maximum of the mean firing rate (averaged over the network) as a function of the external input. The input is a constant, i.e. there is no place-specific input. Simulations are run for 100sec and the first 10sec are discarded to avoid transients. (b) "LFP"s of the spontaneous activity, actually the total input to the network (argument of the nonlinear transfer function) for I = -0.8 (left), -0.5 and 1 respectively. (c) Space-time plots of the spontaneous activity shown with neurons ordered according to their place fields in the last three environments explored. Parameter values are identical to those in Fig.2 and 3.

Figure S7. Place cell statistics. (a) Three sample cells. Spike maps (top), rate maps (middle) and linearized rate along the track with muscimol and post-muscimol (bottom). Each linearized rate profile is fit by a curve $r = r_0 + r_1 \cos(\theta - \phi_1) + r_2 \cos(2(\theta - \phi_2))$ and the coefficients extracted. (b) Top left: The values of the coefficients of the first cosine mode vs. the spatially homogeneous mode for all cells. Top right: The histogram of the ratio r_1/r_0 for all cells. Bottom left: Histogram of firing rates. (c) There is only weak rate-remapping for clockwise versus counter-clockwise motion. (Left to right) The first two Fourier coefficients and phase of the fit to the linearized rate profiles in the clockwise versus the counter-clockwise direction, and a sample cell.

Figure S9. The SC with and without muscimol calculated with different percentages of neurons (with the highest ratio r1/r0).

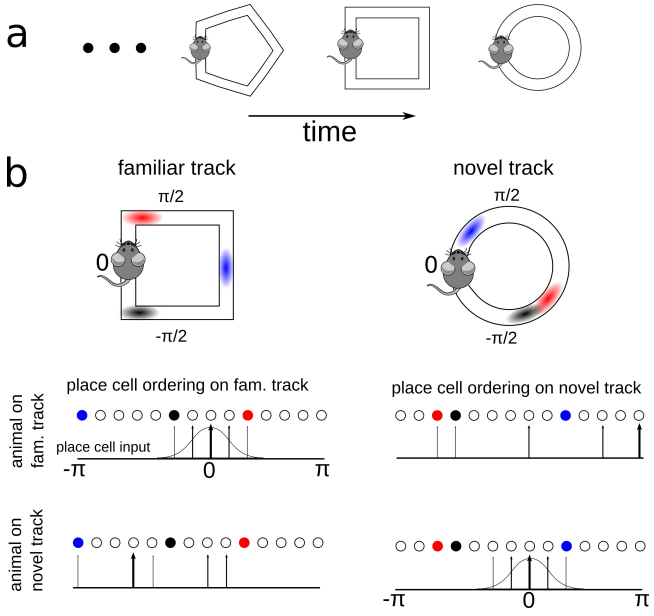

FIGURE 1

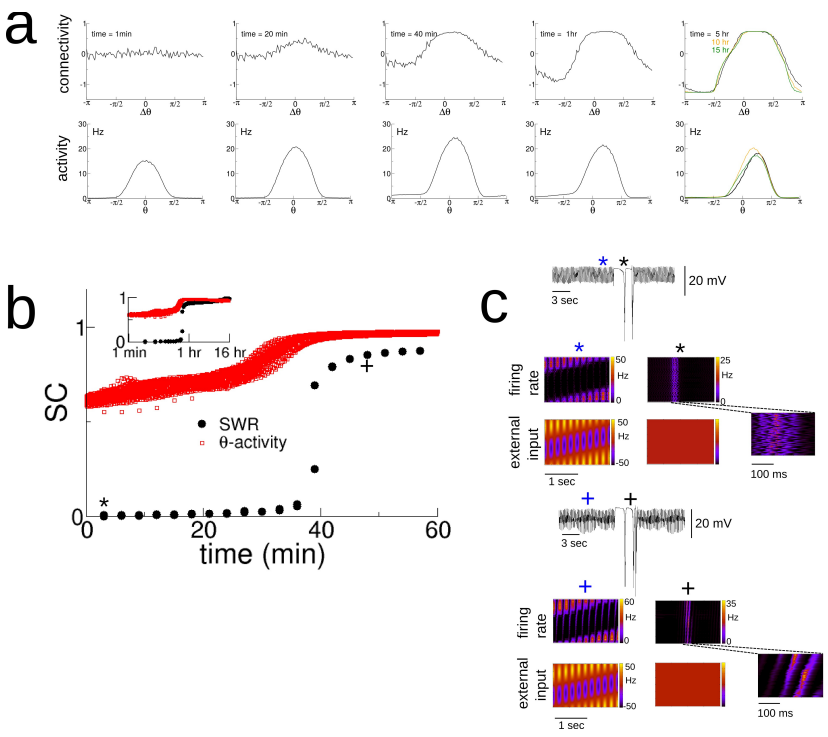

FIGURE 2

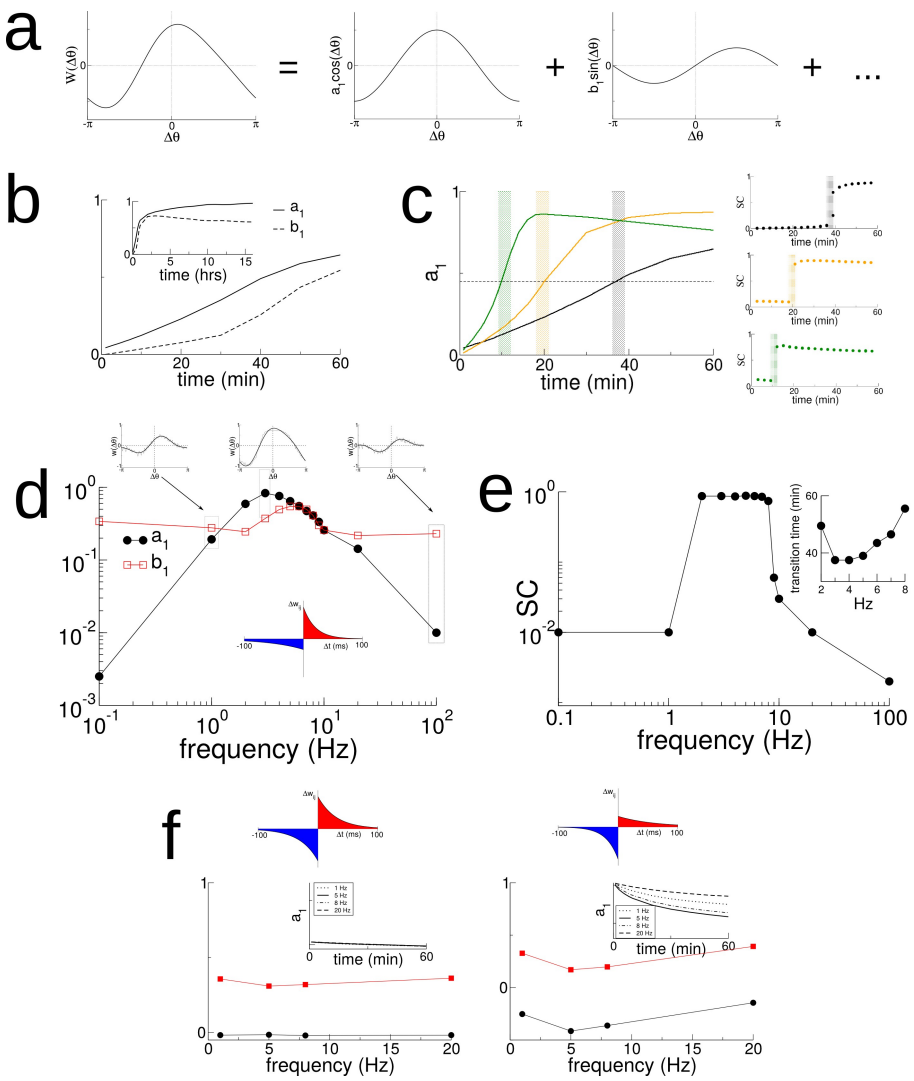

FIGURE 3

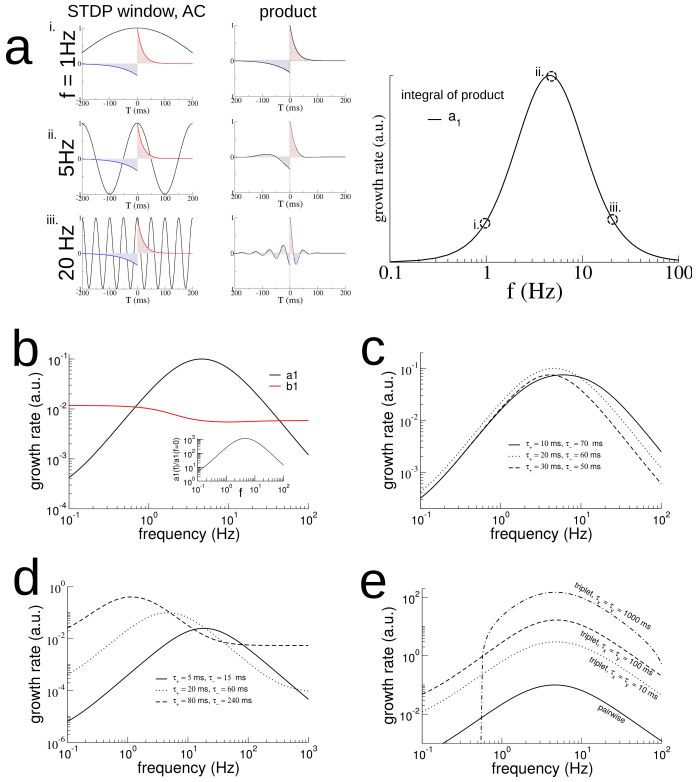

FIGURE 4

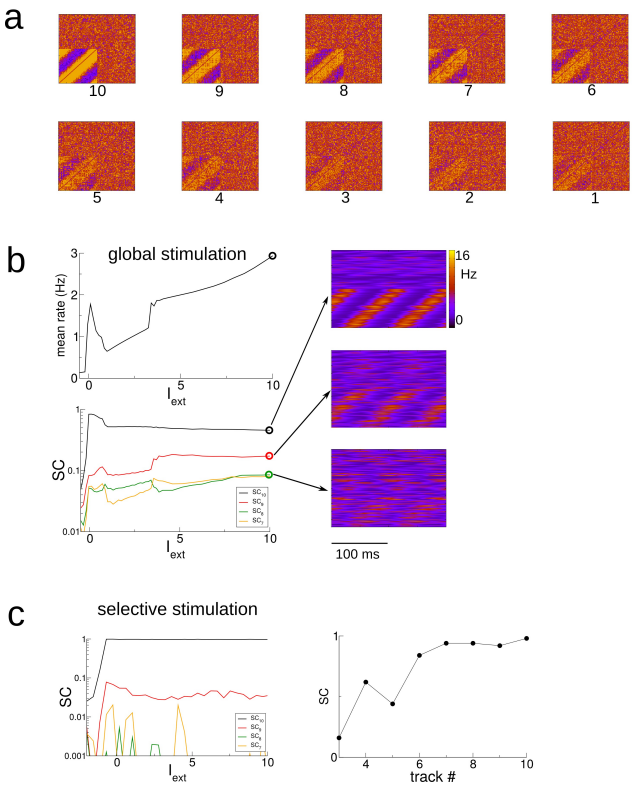

FIGURE 5

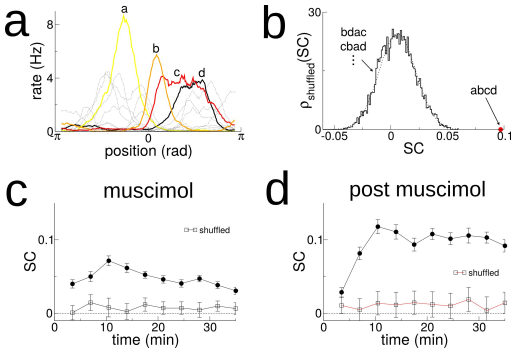

FIGURE 6

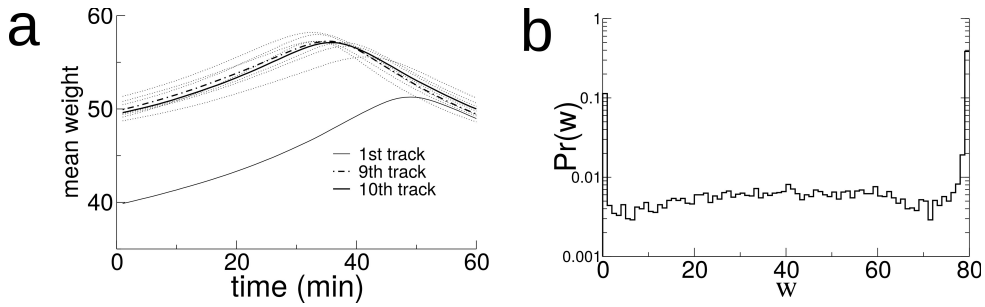

FIGURE S1

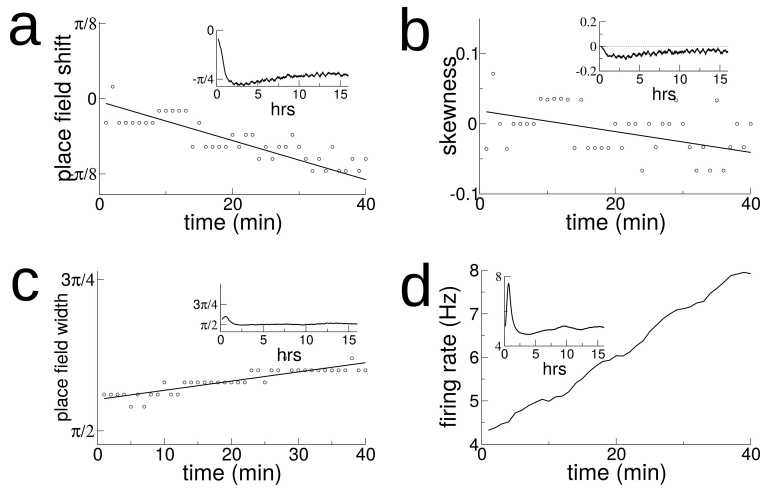

FIGURE S2

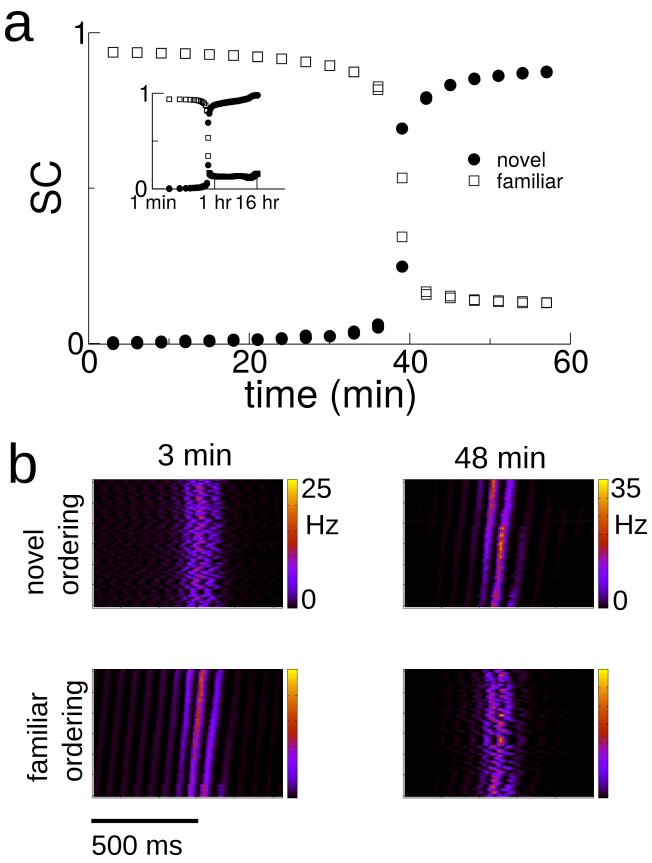

FIGURE S3

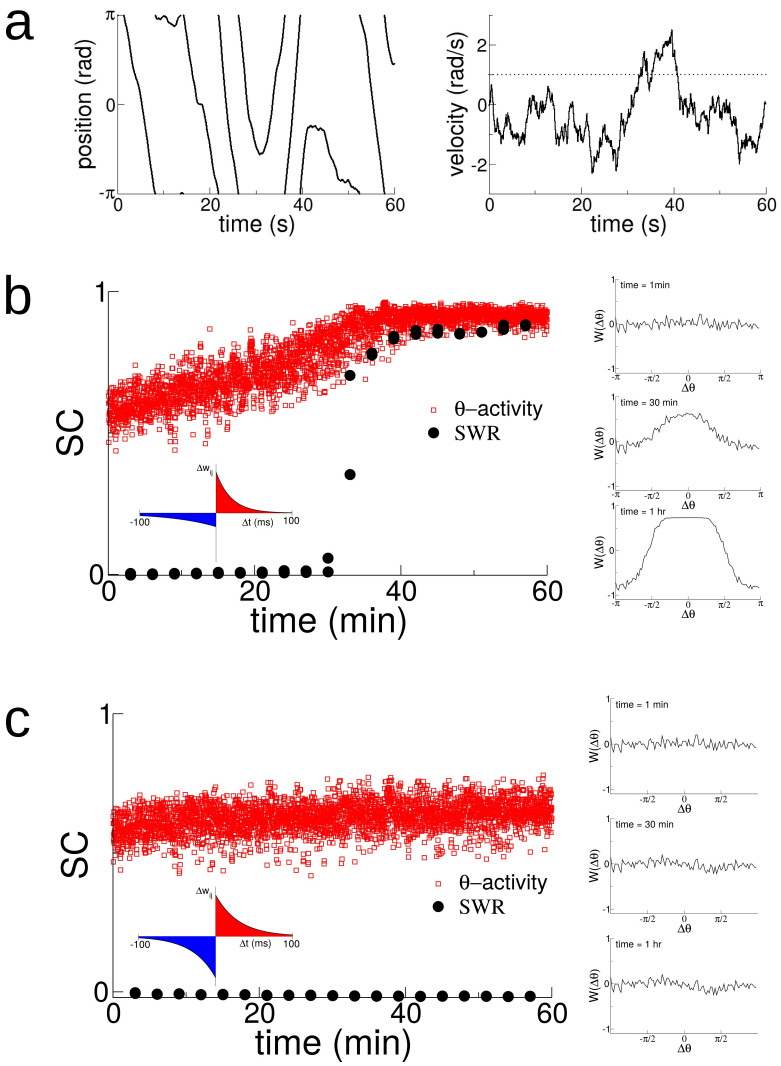

FIGURE S4

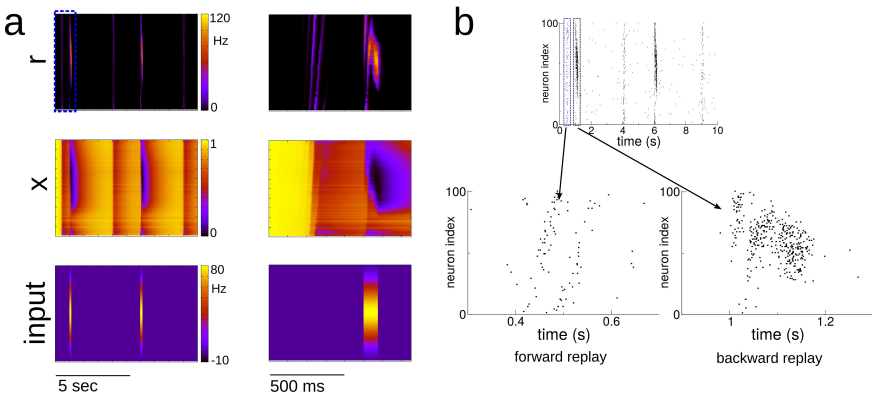

FIGURE S5

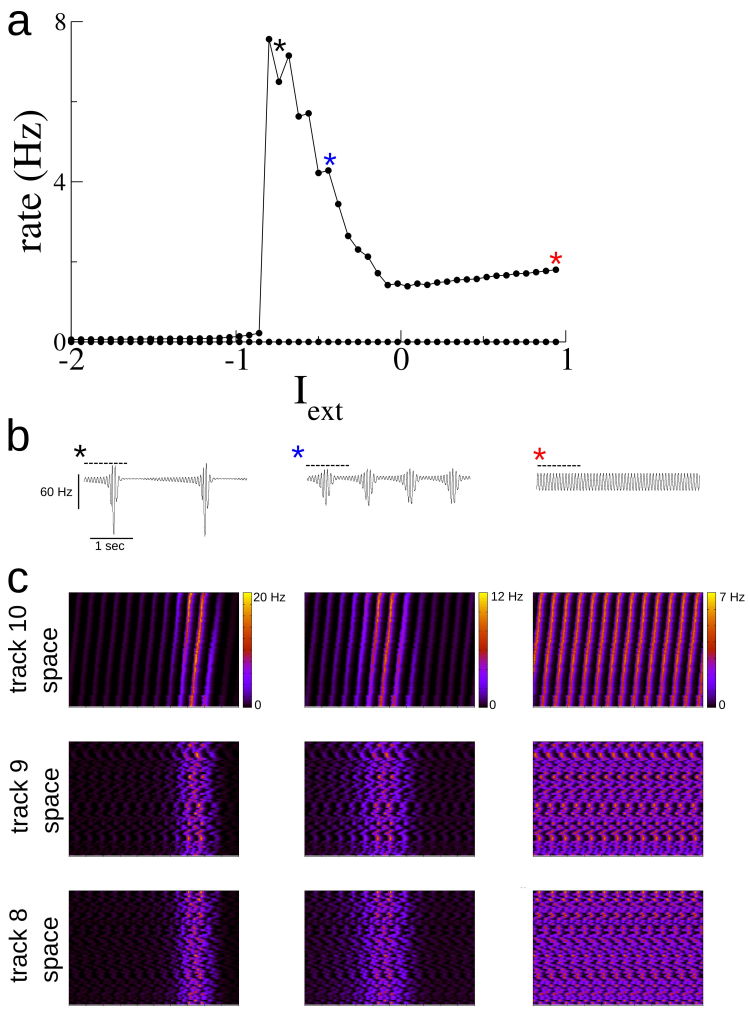

FIGURE S6

FIGURE

S7

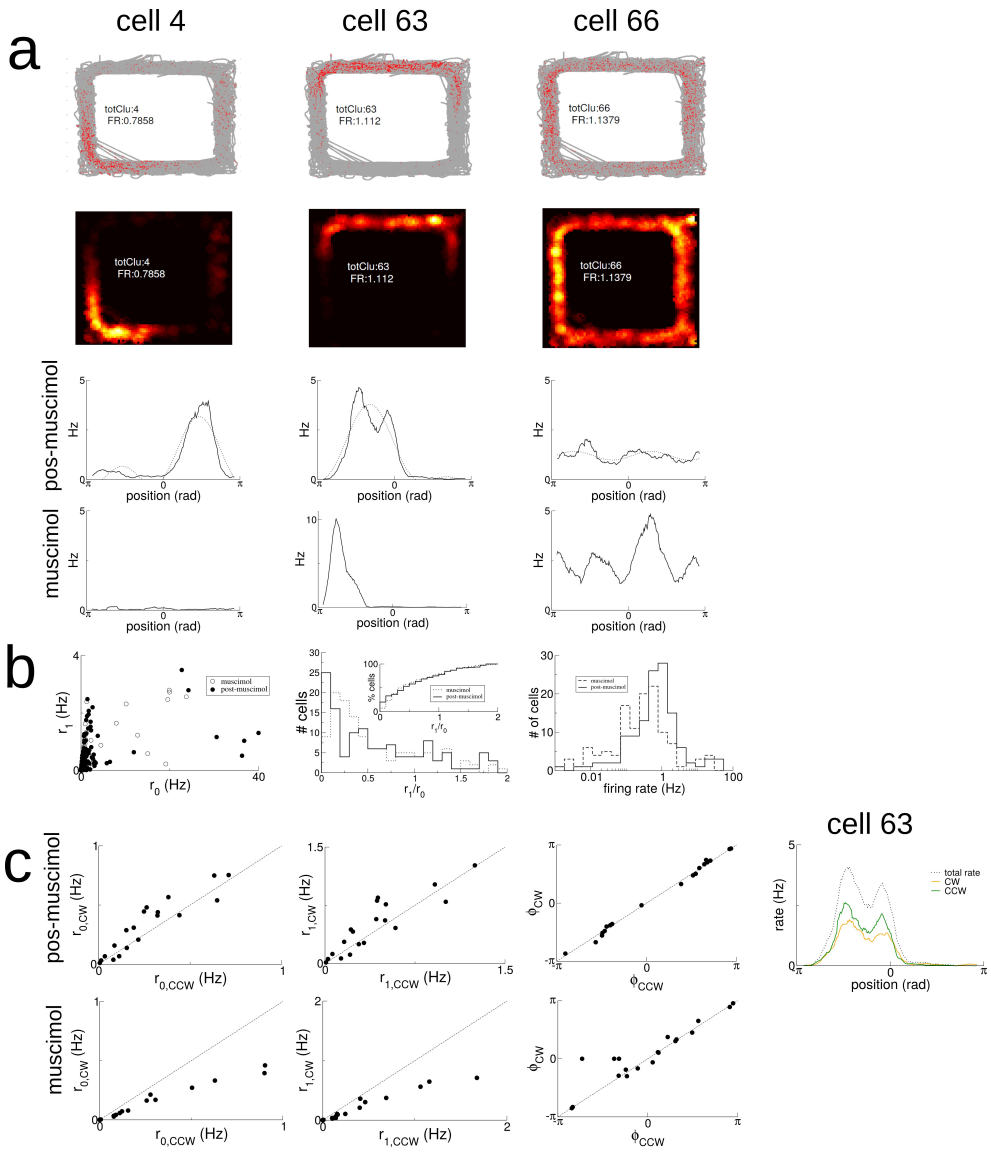

FIGURE S8

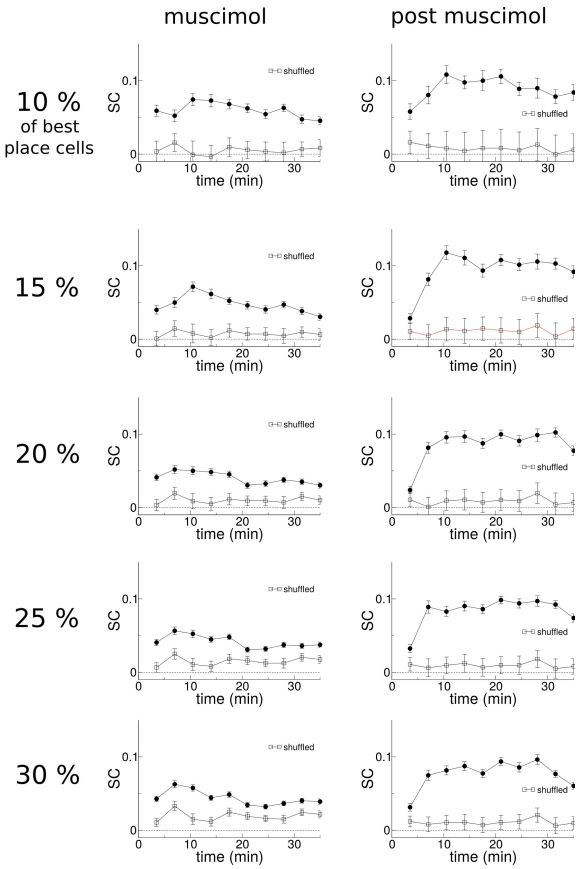

rat A992, square track