

Network Dynamics Underlying the Formation of Sparse, Informative Representations in the Hippocampus

Mattias P. Karlsson and Loren M. Frank

W. M. Keck Center for Integrative Neuroscience and Department of Physiology, University of California, San Francisco, San Francisco, California 94143-0444

During development, activity-dependent processes increase the specificity of neural responses to stimuli, but the role that this type of process plays in adult plasticity is unclear. We examined the dynamics of hippocampal activity as animals learned about new environments to understand how neural selectivity changes with experience. Hippocampal principal neurons fire when the animal is located in a particular subregion of its environment, and in any given environment the hippocampal representation is sparse: less than half of the neurons in areas CA1 and CA3 are active whereas the rest are essentially silent. Here we show that different dynamics govern the evolution of this sparsity in CA1 and upstream area CA3. CA1, but not CA3, produces twice as many spikes in novel compared with familiar environments. This high rate firing continues during sharp wave ripple events in a subsequent rest period. The overall CA1 population rate declines and the number of active cells decreases as the environment becomes familiar and task performance improves, but the decline in rate is not uniform across neurons. Instead, the activity of cells with initial peak spatial rates above ~ 12 Hz is enhanced, whereas the activity of cells with lower initial peak rates is suppressed. The result of these changes is that the active CA1 population comes to consist of a relatively small group of cells with strong spatial tuning. This process is not evident in CA3, indicating that a region-specific and long timescale process operates in CA1 to create a sparse, spatially informative population of neurons.

Key words: hippocampus; learning; CA1; CA3; spatial; dynamics

Introduction

The hippocampus is essential for encoding new memories for space and events (Scoville and Milner, 1957; Squire and Zola-Morgan, 1991; Eichenbaum and Cohen, 2001) and recent work has shown that human hippocampal neurons are active during both memory encoding and retrieval (Gelbard-Sagiv et al., 2008). In the rodent, individual hippocampal neurons fire selectively when the animal is located in a particular subregion of its environment (O'Keefe and Dostrovsky, 1971). These "place cells" encode complex associations available in different locations (Best et al., 2001). Importantly, only $\sim 40\%$ of the principal neurons in hippocampal area CA1 and $\sim 30\%$ of the principal neurons in upstream area CA3 (Wilson and McNaughton, 1993; Lee et al., 2004b; Leutgeb et al., 2004) show clear spatial specificity in a given familiar environment whereas the rest are essentially silent. It is not yet clear how individual neurons become part of a new spatial representation or how this process might differ across hippocampal subregions.

We do know that although both CA3 and CA1 show place

specific activity, these areas differ in their responses to new or altered environments (Guzowski et al., 2004; Lee et al., 2004a,b; Leutgeb et al., 2004; Alvernhe et al., 2008). CA3 place fields shift during the first exposure to a new cue configuration but not during subsequent exposures, whereas the opposite pattern is seen in CA1 (Lee et al., 2004a). This result is consistent with data demonstrating the importance of CA3 for single trial learning (Nakazawa et al., 2003). Furthermore, whereas CA1 cells undergo rapid plasticity in novel environments (Wilson and McNaughton, 1993), CA1 place field plasticity (Austin et al., 1993; Frank et al., 2004) and novelty-specific reactivation (Cheng and Frank, 2008) can extend beyond the first day of exposure to a new place. Together, these results suggest that CA3 and CA1 may respond differently during new experiences.

A parallel set of results has demonstrated that there are increases in activity across multiple brain regions during exposure to novelty (Li et al., 1993; Knight, 1996; Knight and Nakada, 1998; Kirshhoff et al., 2000). These increases have been seen in CA1 (Nitz and McNaughton, 2004; Csicsvari et al., 2007), but it is not yet clear whether these changes are seen in CA3. Furthermore, we do not know whether these increases in activity are relevant to ongoing plasticity in the hippocampal circuit.

We therefore asked whether changes in overall firing are related to changes in place cell coding in CA3 and CA1. We exposed animals to two initially novel environments. The exposures were staggered to allow us to sample CA3 and CA1 activity both during the early days of experience in a new place and also later, once that place was highly familiar. Animals learned to perform an alterna-

Received Sept. 4, 2008; revised Nov. 12, 2008; accepted Nov. 13, 2008.

This work was supported by the Swartz and Sloan Foundations, the John Merck Scholars Program, the McKnight Scholars Program, and National Institutes of Health Grants MH077970 and MH080283. We thank Michael Brainard, Steve Lisberger, Margaret Carr, Sen Cheng, Yuri Dabaghian, Sheri Harris, Caleb Kemere, Steve M. Kim, Ana R. Nathe, and Annabelle Singer for comments on this manuscript.

Correspondence should be addressed to Loren M. Frank, W. M. Keck Center for Integrative Neuroscience and Department of Physiology, University of California, San Francisco, S-859, Box 0444, 513 Parnassus Avenue, San Francisco, CA 94143. E-mail: loren@phy.ucsf.edu.

DOI:10.1523/JNEUROSCI.4261-08.2008

Copyright © 2008 Society for Neuroscience 0270-6474/08/2814271-11\$15.00/0

tion task so that we could determine whether behavioral and neural changes took place on similar timescales. We found fundamentally different patterns of population dynamics in CA3 and CA1, suggestive of specialized roles for these structures in learning.

Materials and Methods

Data collection and preprocessing. Seven male Long–Evans rats (500–600 g) were food deprived to 85–90% of their baseline weight and trained to run on a linear track with one reward site at each end of the track. The pretraining was performed in a different room from the recording experiments. After the animals were accustomed to behaving for liquid reward (sweetened condensed milk), they were implanted with a microdrive array containing 30 independently movable tetrodes according to University of California, San Francisco and National Institutes of Health guidelines (for details, see Frank et al., 2004). After 5–6 d of recovery, animals were once again food deprived to 85% of their baseline weight.

The tetrodes were arranged bilaterally in two 15 tetrode groups centered at AP -3.7 mm, ML ± 3.7 mm. Each group was located inside an oval cannula whose major axis was oriented at a 45° angle to the midline with the more posterior tip of the oval closer to the midline. Tetrodes in the anterior and lateral portion of each group targeted lateral CA3 whereas more posterior and medial tetrodes targeted CA1. Thus, although there was some overlap of CA3 and CA1 tetrodes in individual coronal slices, most CA1 recordings were slightly more posterior than associated CA3 recordings. This arrangement was chosen because the more lateral CA3 cells we targeted tend to project to slightly more posterior/temporal areas of CA1 (Ishizuka et al., 1990), and our goal was to sample neurons from highly interconnected regions of CA3 and CA1.

Each recording day consisted of two or three 15 min run sessions in W-shaped tracks, with rest sessions in a black box before and after each run. We used two geometrically identical but visually distinct tracks that were open to the room but separated from one another by a black barrier (Fig. 1*a*). The tracks had one reward site at the end of each arm, and animals learned to perform a continuous alternation task in which, starting from the center arm of the W-track, they alternated visits to each outer arm for liquid reward (center-left-center-right, and so on) (Frank et al., 2000, 2004). Animal were rewarded only when an outer arm visit was followed by a return to the center arm before the next outbound trajectory was made. Errors were not rewarded, and after an incorrect choice of an outer arm no rewards were given until the animal returned to the center arm. Recording for animals in group 1 ($n = 4$) began on the first day of exposure to track one (T1). These animals experienced T1 for 3 d (two sessions per day), and in three of these animals, track two (T2) was then introduced on day 4. From day 4 onward, these three animals ran at least one session on each track (usually two sessions on T2 and one session on T1). The fourth animal in group 1 ran only on track 1 for two sessions a day for days 1 through 5 and was not exposed to track 2. Animals in group 2 ($n = 3$) were exposed to T1 for 6 d and then ran on both T1 and T2 from day 7 onward. Recordings for group 2 were begun on day 7. The order of track presentations was not the same for every animal, but was kept consistent for each individual animal (supplemental Fig. 1, available at www.jneurosci.org as supplemental material). This design allowed us to compare the activity of single neurons across the two tracks and to obtain recordings across many days as the initially novel environments became familiar.

Each run session was 15 min long, and run sessions were flanked by 20 min rest periods in a high walled black box where neural recording continued. Tetrode positions were adjusted after daily recording sessions for all tetrodes that had poor unit recordings. On rare occasions some tetrodes were moved before recording sessions, but never within 4 h of recording. After data collection electrode locations were identified histologically (Fig. 1*b*) (Frank et al., 2004).

Data were collected using the NSpike data acquisition system (L. M. Frank; J. MacArthur, Harvard University, Cambridge, MA). An infrared diode array with a large and a small group of diodes was attached to the animal's preamplifier array and the animal's position on the track was reconstructed using a semiautomated analysis of digital video of the

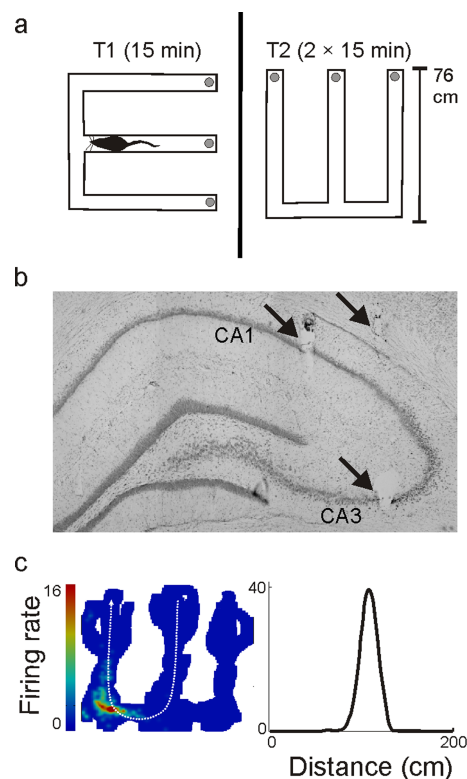


Figure 1. Experimental overview. *a*, Track shapes and orientations in the recording room. The track arms were 76 cm long, and each section of track was 7 cm wide. The gray circles at the end of each arm represent food wells in which rewards were delivered. The black line separating the two W-tracks represents the black barrier that prevented the animal from seeing one track from the other. The diagram of a rat represents the approximate size of an animal relative to the environment. *b*, Histology. Arrows indicate lesions at the ends of three tetrodes. The slice shown was composed from three adjacent photos and contains a lesion from a single tetrode targeting CA3, a tetrode targeting CA1, and a reference tetrode located in the corpus callosum. Overall, most CA1 tetrodes were slightly more posterior than the CA3 tetrodes (see Materials and Methods). *c*, Linearization of firing rates. On the left is a two-dimensional heat map of the occupancy normalized firing rates of one neuron. The animal's location was projected onto one of four linear trajectories. Shown on the right is the linear firing rate for one trajectory from the center to the left arm (shown in white on the heat map). The neuron had a peak spatial rate of 39 Hz and was active on 15% of the track. The linearized rate is higher than the nonlinearized rate because the cell was directional and because two-dimensional smoothing reduces the firing rates.

experiment with custom-written software. Spike data were sampled at 30 KHz, digitally filtered between 600 Hz and 6 KHz (2 pole Bessel for high and low pass) and threshold crossing events were saved to disk. Continuous LFP data from all tetrodes was sampled at 1.5 KHz, digitally filtered between 0.5 and 400 Hz and saved to disk.

After neural data were collected, individual units were identified by clustering spikes using only peak amplitude and spike width as variables. All spike sorting was done using custom software (MatClust, M. Karlsson). It was generally possible to use a single set of cluster bounds defined in amplitude and width space to isolate units across an entire 2–3 h recording session. In the minority of cases in which there was a slight shift in amplitudes across time, units (putative single neurons) were clustered only when that shift was coherent across multiple clusters and where plots of amplitude versus time showed a smooth shift. Only units with at least 100 spikes fired across a day of recording and a clear refractory period were included. Furthermore, only units where the entire cluster was visible were included. Thus, no units in which part of the cluster vanished into the noise or was cutoff by the recording threshold were clustered. Because units were clustered across both run and rest sessions, we were able to identify well isolated units that fired frequently in the rest box but infrequently on one or both tracks. Data from these cells were

included in the average rate analyses and in the calculation of proportion of cells active.

After clustering we separated units into putative excitatory and fast spiking inhibitory neurons using standard criteria (Fox and Ranck, 1981; Frank et al., 2001). Putative inhibitory cells were identified based on their high average rate (> 7 Hz) and their narrow spike waveforms (< 0.25 ms measured from peak to trough), as these criteria produced clearly separated groups of neurons (data not shown). Only a relatively small number of putative inhibitory cells were isolated, so we restricted our analyses to putative principal neurons.

Analysis of behavioral data. We measured the performance on the continuous alternation task by analyzing the sequence of center arm to outer arm visits on the W-track. A decision was correct if the animal visited the outer arm that was not visited on the previous outbound trajectory. We calculated the proportion of incorrect decisions made as a function of the amount of experience on the track. For this analysis data were combined across animals as a pool of ones (correct decisions) and zeroes (incorrect decisions) for each day, and using a previously described dynamic state space algorithm (Smith et al., 2004), we estimated the probability of an incorrect decision with confidence limits across all days. From this, we calculated the time constant with confidence limits as the point at which errors decayed to $1/e$ of the total range. We also measured the average amount of time the animals used to travel from one well to another as well as the average velocity for all times included in the analyses of spike data.

Analysis of neural data. All analyses were performed using custom software written in Matlab (Mathworks). Only days that contain data from at least two animals were included, and all analyses for the run periods excluded times when animals were moving along the track at a speed < 3 cm/s to ensure that periods of nonspatial firing during immobility did not contribute to our results. The average activity for each cell was calculated as the total number of spikes fired during running divided by the length of time the animal was running. When comparing cell rates across days, only the first run session was used from each track per day. To measure place field properties we calculated the “linearized” activity of each cell. The behavioral data were separated into four trajectories (center arm to left arm, left to center, center to right, and right to center) and the animal’s linear position was measured as the distance in cm along the track from the reward site on the center arm. We then produced four corresponding occupancy maps and spikes per bin maps using 2 cm spatial bins. Both the occupancy and spikes per bin map were smoothed with a 4 cm SD Gaussian curve with a total extent of 20 cm. The spikes per bin map was divided bin by bin by the occupancy map to produce a smoothed, occupancy normalized firing rate map which captured the spatial specificity seen in the standard 2D place field plot (Fig. 1*c*). Spatial peak rate was defined as the maximum rate across all spatial bins for one 15 min run session. The similarity of spatial coding for single cells across the two environments was computed using place field overlap, defined as two times the sum of the overlapping areas of the linear rate curves divided by the sum of the areas of each curve (Battaglia et al., 2004). This measure is bounded between 0 and 1, where 1 signifies perfect overlap. All trajectories with at least 3 Hz peak rate were averaged together for this measure.

We also examined firing rates during sharp-wave ripple (SWR) events during the rest after the first run of each day. This allowed us to determine whether the changes in rates that we found were specific to the spatial setting and whether these changes persisted after the initial experience in the environment into periods that have been associated with memory consolidation (Buzsáki, 1986; Buzsáki et al., 1992; Sutherland and McNaughton, 2000). SWRs were identified based on peaks in the local field potential (LFP), recorded from one channel from each tetrode in the CA3 and CA1 cell layers. A reference tetrode was positioned the corpus callosum and all neural signals were recorded relative to that reference to eliminate muscle artifacts from the recordings. The raw LFP data were bandpass filtered between 150 and 250 Hz and the SWR envelope was determined using a Hilbert transform. The envelope was smoothed with a Gaussian with a SD of 4 ms and a width of 32 ms. SWRs were defined as contiguous periods when the smoothed SWR envelope stayed above 3 SDs of the mean for at least 15 ms on at least one tetrode.

The full SWR events included times immediately before and after the triggering event during which the envelope exceeded the mean. We calculated the rate in these SWR events by taking the total number of spikes seen during SWRs and dividing it by the total length of all the SWR events during the rest session.

We then examined the relationship between the spatial peak rate and the proportion of the environment over which each neuron was active. The proportion of the environment over which each neuron was active was defined as the length of the region(s) in which the cell fired at above 0.5 Hz linearized firing rate divided by the total length of the track. For this analysis we chose, for each neuron, the first session on the most novel track on each recording day. This restricted the analyses so that the same cell did not provide two data points for different blocks. The results were indistinguishable if this condition was relaxed (data not shown). All statistical tests for differences in proportions used the Z test for proportions, and the associated Z scores are reported in the text.

To visualize the changes that we saw in spatial peak rates over time we constructed smoothed probability distribution functions (PDF) of peak and average rates. To do so we computed the cumulative distribution function (CDF) of the data, smoothed the CDF using a Gaussian curve and computed the PDF as the derivative of the CDF. The Gaussian had a SD of 2.3 Hz for the peak rate analysis and a SD of 0.23 Hz for the average rate analysis. The width of the Gaussian was 10 times the SD. Note that these smoothed distributions were used only for visualization and all statistics were computed on raw data.

We also developed a simulation approach to determine whether a measured set of within day changes in peak spatial rate seen between days 3 and 6 (referred to as the middle block) could account for the changes we observed from the first 2 d (early block) to the final days (days 7 onward, late block). To do so we first constructed a scatter plot with the peak spatial rate of each middle block neuron on the x -axis and the change in that peak spatial rate from the first run of the day to the second run of the day on a track on the y -axis (supplemental Fig. 2, available at www.jneurosci.org as supplemental material). This scatter plot indicated how the peak spatial rate of a neuron related to the change in that peak spatial rate. This relationship allowed us to ask how the peak spatial rates of a group of cells from the early block would evolve if they changed in the manner of cells from the middle block.

We then took individual cells from the early block, chosen at random, and simulated the evolution of their peak firing rates using the changes measured in the raw peak rate/change in peak rate data from the middle block. For each early block cell we computed its peak spatial firing rate and then found the middle block cell whose peak spatial firing rate during the middle block was closest to that of the early block cell. We then took the change in the peak rate for the selected middle block cell and applied that change to the peak rate of the early block cell. Thus, a new peak rate for the early block cell was obtained by either increasing or decreasing the cell’s peak rate according to the rate change of the matched cell from the middle block. This allowed us to apply the changes in peak rate we measured within a day for the middle block to the early block, thereby evolving the early block population according to the dynamics of the middle block cells.

We repeated the same procedure multiple times, each time using the newly generated rate for the early block cell as illustrated in supplemental Figure 2 (available at www.jneurosci.org as supplemental material). Peak rate changes from the middle novelty block were randomly sampled with replacement to match the size of the actual pool of cells. This meant that single cells could be picked multiple times from the early novelty block without duplicating a particular set of changes in rate. During the evolution process, any cells that dipped < 0 Hz peak rate were set to 0 Hz before the next iteration. The evolution process was repeated with 10,000 random cell picks from the early novelty block, each with 100 iterations through the process.

We also calculated the distributions that would result from random peak rate versus peak rate change relationships. We chose at random a measured peak rate for each neuron and, separately, a measured peak rate change. We then used that population to simulate the dynamics of the early block group. This process was repeated 200 times. These simulations occasionally produced “runaway” cells whose rates went above 70

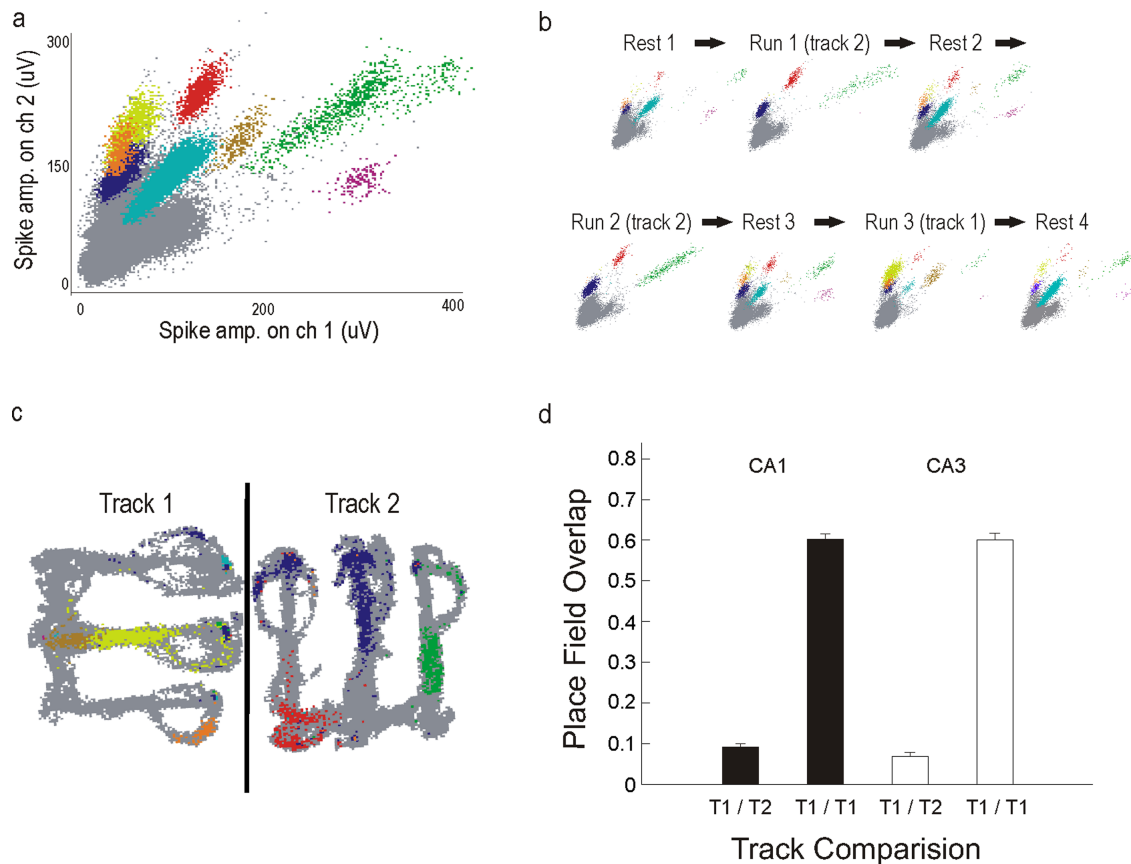


Figure 2. Recording stability and distinct spatial representations for the two tracks. *a*, Spike amplitudes on two of four tetrode channels recorded across 1 d (experimental day 12, animal 6, group 2). Different colors represent different clustered cells; gray points represent spikes that were not clustered. *b*, Spike amplitudes for each rest and run session illustrating the stability of recordings and the presence of some putative single neurons that were well isolated due to their activity during rest sessions but were silent or minimally active on the tracks (e.g., purple cluster). *c*, Spiking as a function of position for one run session on each track where each pixel in which a spike was detected is colored as in *a*. The black line represents the barrier between the tracks. The differences in the spiking patterns between track 1 and track 2 illustrate the distinct spatial codes for these two environments. *d*, Overlap between the linearized firing rates of the two runs from the same day across track 1 and track 2 (T1/T2) and two runs from the same day in track 1 (T1/T1). The low overlap of T1 and T2 indicates that the spatial firing patterns of individual CA1 and CA3 neurons differed across these two environments, whereas the high overlap of T1 with itself indicates that the spatial representation of the first track was generally stable within a day.

Hz, and we excluded these cells from the final distribution. Runaway cells did not result from the simulation based on the actual peak rate–rate change distribution.

Results

We isolated large numbers of neurons on each day from both CA1 (total $n = 708$) and upstream area CA3 (total $n = 984$) which provides direct synaptic input to CA1 (Fig. 2*a*). These neurons were approximately equally split between the left and right hemispheres (CA1 left, $n = 290$; CA1 right, $n = 318$; CA3 left, $n = 482$; CA3 right, $n = 502$), and we did not find any differences in firing patterns between the two hemispheres, so all data were collapsed across hemispheres. Recordings were generally stable across the entire 2–3 h session, and we were therefore able to identify putative single neurons that were active on one or both tracks as well as neurons active only during the rest sessions (Fig. 2*b*). As expected (Frank et al., 2004; Leutgeb et al., 2004), CA3 and CA1 place fields were evident during the first exposure to the track. Furthermore, there were distinct spatial representations of the two tracks in both areas (Fig. 2*c,d*). Only a small number of putative inhibitory neurons were recorded, so we did not include data from these neurons in this study.

CA1-specific rate dynamics

We first sought to characterize the timeline of the previously observed increase in CA1 activity during novelty. We compared

absolute firing rates across the two tracks for animals in group 1 and found that individual CA1 neurons were on average more active on the more novel T2 (paired $t_{(145)} = 2.04$, $p < 0.05$), a result consistent with previous studies. Using recordings across multiple sequential days, we were able to quantify the decay of the rate change. For every recording session, we calculated the average firing rates of all recorded CA1 principal neurons, including those that fired robustly in the rest box but were largely silent on the tracks, and combined these rates into an estimate of the population average rate. We included all clustered principal neurons rather than imposing a cutoff because we wanted to minimize the selection bias for the measure of population activity. Rates slowly declined across multiple days of experience (T1, $R^2 = 0.0259$; slope \pm 95% bounds = -0.0896 ± 0.0620 Hz/d, $p < 0.005$; T2, $R^2 = 0.0441$; slope = -0.249 ± 0.191 Hz/d, $p < 0.05$) indicating a long timescale process rather than a transient response to novelty (Fig. 3*a*, left).

The same dynamics were observed in recordings from group 2 (Fig. 3*a*, right). Again, CA1 activity was initially high in the more novel T2 and slowly decreased over multiple days of exposure ($R^2 = 0.0553$, slope = -0.088 ± 0.033 Hz/d, $p < 10^{-4}$) whereas activity on the more familiar T1 was lower on all recording days (paired $t_{(457)} = 4.96$, $p < 10^{-6}$) and did not decline appreciably with further experience ($R^2 = 0.0034$; slope = -0.022 ± 0.0373 , N.S.). We should note that although we used a linear regression to




Figure 3. Hippocampal rate dynamics. *a*, Population average rates in CA1 as a function of experiment day. The set of days associated with each group is indicated at the top. Rates were high in each novel environment and then decayed over the course of a few days. Track 1, Group 2 illustrates the final firing rate level reached when animals had 6 d of experience on the track before recording began. *b*, Population average rates in CA3 as a function of experiment day. There were no significant changes in CA3 rates across days. *c*, Rates pooled from all animals and both tracks, plotted as a function of the number of days spent on the track. Top, Percentage of CA1 and CA3 cells active for the three novelty blocks, where “active” is defined as having a mean rate >0.1 Hz. Error bars represent SEM, calculated by repeated resampling of the data with replacement (bootstrap).

quantify the decline in slopes, the rates from group 1, track 1 and group 2, track 2 appear somewhat sigmoidal, whereas the rates from group 1, track 2 are more linear or exponential in character. The large variability in individual neuron mean rates makes it difficult to determine the exact shape of these curves, but it is nevertheless clear that there is a substantial decrease in rate over time. Furthermore, because the rate difference is observed both within single recording sessions and across days, the effect is un-

likely to be due to within-day confounds such as environment sequence or cell isolation instability, or across-day confounds such as quantity of recorded cells or tetrode placement. Finally, none of the potential behavioral confounds we examined, including differences in velocity, head direction, or ripple-related activity in the local field potential, could explain these results (supplemental Fig. 3, available at www.jneurosci.org as supplemental material). In particular, animals moved, on average, more slowly in the novel track, and thus the higher rates seen in novel environments were associated with lower movement speeds. Given the previously established positive correlation between speed and CA1 firing rates (McNaughton et al., 1983), we would have expected lower rather than higher rates if velocity differences caused the firing rates differences.

The increase in rate was not present one synapse upstream in CA3 (Fig. 3*b*). CA3 neurons had on average lower activity than CA1 for most recording days, and unlike neurons in CA1, did not show a consistent decline in activity as the environments became more familiar (linear regression, both groups, N.S.). We compared the activity of single CA3 cells across the two tracks and found no significant difference in activity for either group ($t_{(185)} = 0.1$, $t_{(612)} = 1.54$, N.S.). Thus, unlike in CA1, no significant population rate change was observed in CA3 either within or across days. These results indicate that the change in CA1 was not due to inheritance of higher rate activity from CA3.

We pooled data from both groups of animals and found that the ensemble of neurons in CA1 emits more than twice as many spikes per second in a novel compared with a more familiar environment (Fig. 3*c*). Here, day 1' represents the combined data across all animals from the first day of exposure to each track. CA1 population activity during initial exposures was ~ 0.9 Hz per neuron compared with ~ 0.4 Hz observed during more familiar exposures, and the majority of the rate decline occurred during the first 5 d of exposure to an environment. This trend could also be seen when we examined histograms of average firing rates across days (supplemental Fig. 4, available at www.jneurosci.org as supplemental material). When we pooled CA3 neurons across all animals, we found that population rates were flat, averaging ~ 0.3 Hz across all exposure days (linear regression, N.S.).

Changes in CA1 activity were associated with the initial recruitment of a large fraction of CA1 neurons and the subsequent reduction in the number of active cells. We divided our data into three blocks: early (days 1–2), middle (days 3–6), and late (days 7–14), to provide accurate estimates of the initial, intermediate, and stable periods of the rate curve (Fig. 3*c*, top). We then applied a 0.1 Hz mean rate threshold to designate active neurons and counted the number of active neurons out of the total number of isolated neurons, once again including cells that were active on the track and cells that were inactive on the track but could be isolated in the rest sessions before and after the behavior. This threshold was chosen to approximately match the proportion of active CA1 neurons in familiar environments reported in previous studies (Wilson and McNaughton, 1993; Lee et al., 2004b; Leutgeb et al., 2004), but our results were consistent across a range of thresholds. We found that 151/275 (55%) of cells were active in the early novelty block and 203/406 (50%) in the middle block whereas only 239/600 (40%) were active in the late block (early vs late, $Z = 4.17$, $p < 10^{-4}$; note: some cells provided data for both tracks 1 and 2, which in turn provided data for two novelty blocks). There was no evidence for similar changes in CA3, in which the proportion of active cells was constant across

the three novelty blocks, with 108/373 (29%), 185/591 (31%), and 234/776 (30%) active (early vs late, $Z = 0.42$, N.S.).

This decrease in the number of active cells in CA1 prompted us to ask whether the decrease in rate was present for both active and minimally active neurons. We therefore computed the population firing rate change for active neurons (average rate >0.1 Hz) across recording days. We found that the population mean rates for active CA1 neurons also declined over time, but that, as expected, there were no changes in the population mean rates for CA3 neurons (supplemental Fig. 5, available at www.jneurosci.org as supplemental material).

We then asked whether the changes in CA1 activity were related to changes in behavioral performance on the task. We found that there were broad similarities in the timescale of the improvement in performance (Fig. 4*a,b*) and the decline in the CA1 population average rate. To quantify the overall similarity we calculated an exponential fit to the behavioral data as a function of environment novelty, collapsing across tracks, and found that the decline in errors had a time constant of 3.55 ± 0.7 d (mean \pm 95% confidence bound) which was comparable to the exponential time constant for the decay of CA1 rates of 3.12 ± 2.21 d. In addition, when we collapsed data across groups we found that the first exposure to the novel track 2 was associated with higher error rates compared with the same day's exposure to track 1 even though the required alternation task was identical for both tracks (z test for proportions, $p < 0.03$), although neither within group comparison reached significance. Animals also moved more slowly in the more novel tracks, taking longer times to travel between reward sites (track 1 vs track 2, days 1–3, $t_{(1023)} = 3.32$, $p < 0.001$). This change was observable across days with times converging to baseline levels after ~ 4 d (supplemental Fig. 3*c,d*, available at www.jneurosci.org as supplemental material). Overall, the similarity in the time course of behavioral and firing rate change both across days and within a day suggests that firing rate dynamics are related to an ongoing learning process. At the same time, the high variability of behavioral performance across animals (supplemental Fig. 6, available at www.jneurosci.org as supplemental material) made it impossible for us to relate the rates for individual animals to their behavioral performance, so although behavioral learning and CA1 rate changes appear to be correlated, we could not determine whether rate changes predicts behavioral performance.

We then asked whether the high novel environment rates seen

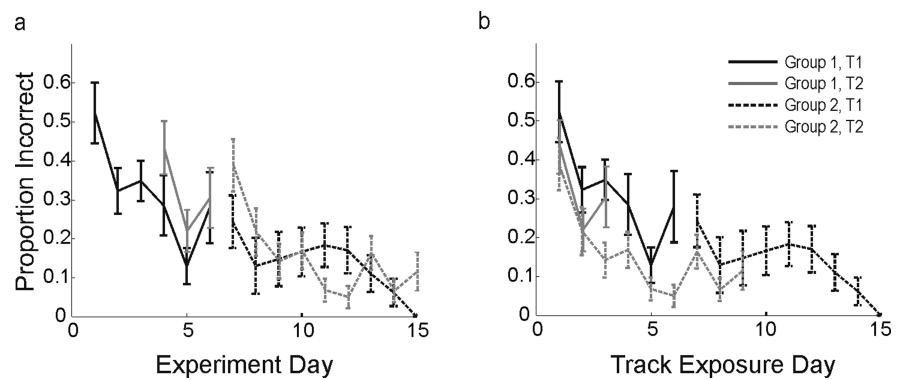


Figure 4. Proportion of outbound error trials. *a*, Proportion of outbound error trials for each group and track plotted as a function of experiment day. Each point represents the proportion of error trials as a function of the total number of trials on a day of experience on a track, collapsed across animals in each group. Error bars represent SEM based on the binomial distribution applied to the number of error trials out of the total number of trials. Animals performed worse in the more novel track 2 compared with the more familiar track 1 within a day when data from both tracks were collapsed (Z test for proportions, $p < 0.03$), but the comparison was not significant for either track taken alone. *b*, Proportion of outbound error trials for each group and track plotted as a function of the day of exposure to each track. There is an overall trend of improvement in each track with a time course similar to that of the decline in CA1 firing rates.

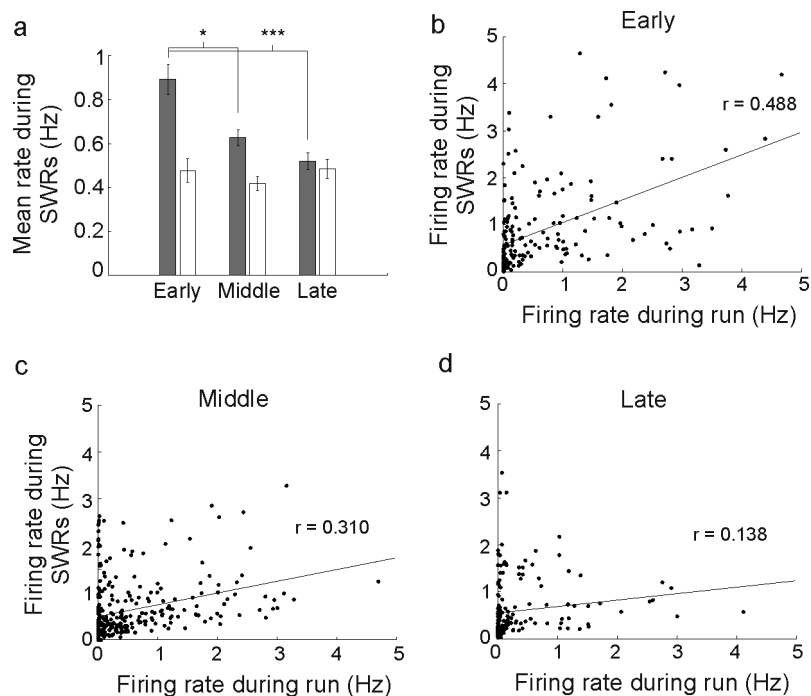


Figure 5. Firing rate changes during SWR events in the rest box mirror changes seen during the run on the track. *a*, Mean rate of CA1 (gray bars) and CA3 (white bars) neurons from SWRs detected during the rest after the first run of each day of recording. CA1 SWR rates declined from the early to the middle and late novelty blocks, and the rates during SWRs were similar to the rates during the run session. CA3 SWR rates, in contrast, did not change across novelty blocks. *b*, Scatter plot of individual neuron's firing rates during running versus rates during rest SWRs. The linear fit (black line) was associated with a significant positive correlation between run and rest rates ($r = 0.488$, $p < 10^{-11}$), indicating that rate during the run was predictive of later activation during the rest. This fit remained significant when the two apparent outlier points at which the rate during running was >4 Hz were removed ($r = 0.420$, $p < 10^{-8}$). *c, d*, Scatter plot of rates for the middle (*c*) and late (*d*) block groups (middle, $r = 0.310$, $p < 10^{-6}$; late block, $r = 0.138$, $p < 0.05$). Error bars represent SEM. * $p < 0.02$, *** $p < 10^{-4}$.

in CA1 were restricted to the run session. Somewhat surprisingly, we found that these higher CA1 rates persisted into the subsequent rest period. Previous studies have documented preferential reactivation of cells associated with new experiences (Cheng and Frank, 2008; O'Neill et al., 2008), but it is not yet clear whether this reactivation interacts with the rate changes we observed. We calculated the average firing rate of CA1 and CA3 neurons during

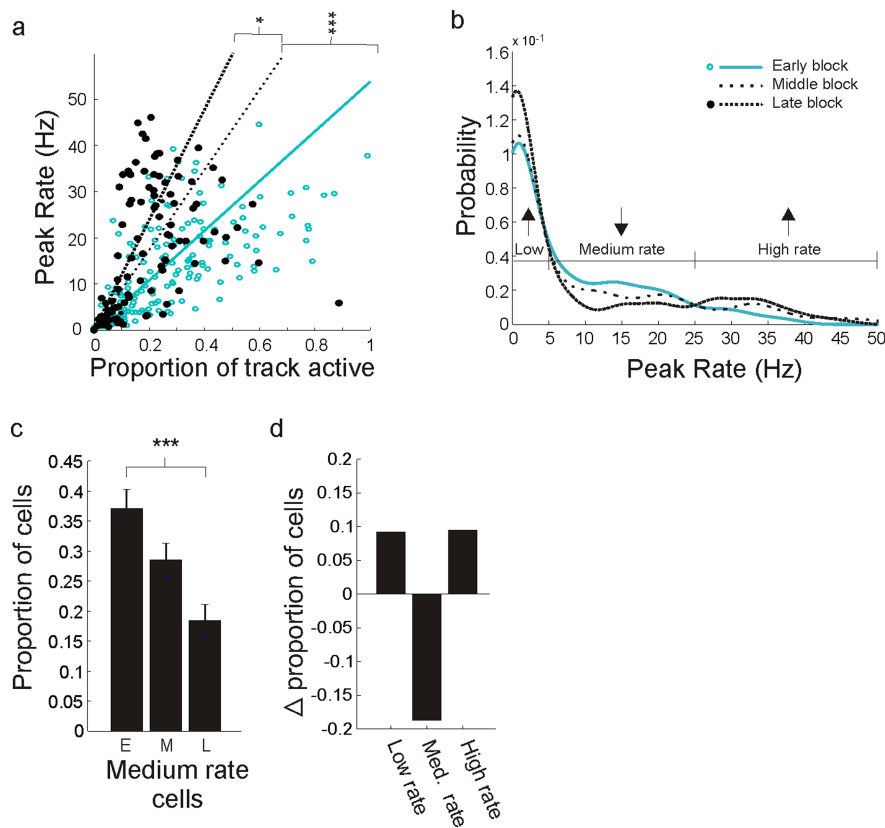


Figure 6. Differential change as a function of peak spatial rate in CA1. *a*, Long time scale evolution of spatial specificity in CA1. The scatter plot shows CA1 peak spatial rates versus proportion active, defined as the proportion of the track at which the cell had >0.5 Hz linearized occupancy-normalized rate. Regression lines were constrained with a zero y -intercept. Points from middle block were omitted for clarity. There was a clear increase in specificity illustrated by the increase in the slope of the regression line: high rate cells from the late block tended to cover smaller portions of the environment than high rate cells from the early block. *b*, Peak rate distributions for the early, middle, and late novelty blocks. The distributions were smoothed for visual clarity, which reduces the sharp increase in proportion near zero, but all statistics were performed on the raw, unsmoothed data. The distributions were significantly different: early versus middle, KS test, $p < 0.05$; early versus late, $p < 0.001$. *c*, Proportion of cells in the intermediate 5–25 Hz peak rate groups as a function of novelty. There was a highly significant decrease in the number of medium rate cells. *d*, Change in proportions for the low, intermediate, and high peak spatial rate groups. The decrease in the intermediate group resulting in equal magnitude increases in the low and high groups. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

SWR events that occurred during the rest period after the first run of each day. The results mirrored our findings for the rates during running (Fig. 5*a*), in that the CA1 SWR population mean rates were ~ 0.9 Hz in the early block but dropped to ~ 0.6 Hz for the middle block (early vs middle, rank-sum, $p < 0.02$) and were ~ 0.5 Hz in the late novelty block (early vs late, rank-sum, $p < 10^{-4}$). Furthermore, when we examined the relationship between average rates during the run and rates during SWRs after the run for individual CA1 neurons, we found a strong, significant correlation for the early and the middle block and a weak correlation for the late block (Fig. 5*b–d*) (early block: $r = 0.488$, $p < 10^{-10}$; middle block $r = 0.310$, $p < 10^{-6}$; late block $r = 0.138$, $p < 0.05$). There was also a significant decrease in the correlation from the early to the late plot ($p < 0.001$) indicating that run and SWR rates decouple over time.

Network selection of cells with strong spatial tuning

We then asked whether the specific tuning properties of individual CA1 place cells influence whether or not they are silenced. We examined the proportion of the track on which each neuron was active and found that this proportion tended to decrease as environments became familiar (rank-sum, median early block proportion = 0.08; median late block proportion = 0.03, $p < 10^{-3}$),

indicating that spatial activity became more focused with experience. This trend persisted when we considered only cells with substantial spatial activity (peak spatial rate of at least 5 Hz, rank-sum, median early block proportion = 0.30; median late block proportion = 0.22, $p < 0.005$).

One might expect that the overall decrease in both total firing and in the spatial extent of activity was a result of a suppression of activity in CA1. In actuality, however, some cells showed an enhancement of spatial firing over time. We examined the relationship between the peak activity within each cell's place field and proportion of the environment on which each neuron was active (Fig. 6*a*; supplemental Fig. 7, available at www.jneurosci.org as supplemental material). In the early block (open blue circles; blue line) many CA1 cells were active across large fractions of the environment, yet tended to have low peak spatial rates. As the environment became familiar (middle block, dotted line; late block, filled black circles and dashed black line) cells tended to be active across smaller proportions of the environment, and there appeared to be some segregation into a higher and a lower rate group. Note that middle block points were omitted for clarity. We first quantified these changes using a regression and found that the slopes of regression fits to the data increased with experience (early vs middle block, $p < 0.0001$; middle vs late block, $p < 0.02$).

To determine whether the apparent segregation into higher and lower peak rate groups was real we plotted the distribution of peak rates for each of the three novelty blocks. We found clear differences across novelty blocks [early vs middle, Komolgorov–Smirnov (KS) test, $p < 0.05$; early vs late, $p < 0.001$] (Fig. 6*b*). The three novelty block distributions intersected at two central points, providing a natural grouping of peak spatial rates. Using this grouping we compared the early and late blocks and found that the proportions of cells with intermediate peak rates decreased by approximately half (5–25 Hz, $Z = -4.56$, $p < 10^{-6}$) (Fig. 6*c*). This decrease was associated with a redistribution of neurons into both the high peak rate and low peak rate groups (Fig. 6*d*), resulting in more than a doubling of the proportion of neurons (from 7 to 17%; $Z = 3.1$, $p < 0.01$) in the high peak rate group and equal magnitude increase (from 55 to 65%, $Z = 2.0$, $p < 0.05$) in the low peak rate group. This segregation into two rate groups suggests that the activity of some neurons with intermediate peak rates is enhanced, whereas the activity of others is suppressed.

We then asked whether the enhancement of intermediate peak rate neurons could have been a side effect of a process that selected neurons based on mean rate. We examined the distribution of average rates for the cells in the early, middle and late novelty block and found that, other than the expected shift toward lower rates, there were no other systematic changes in the distributions (Fig. 7). This shows that the segregation process is

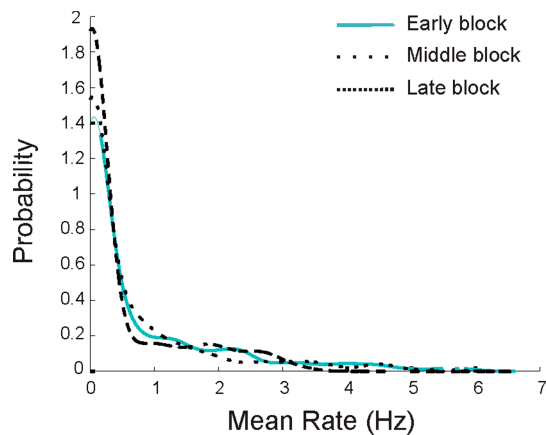


Figure 7. Cell selection process is not dependent on average rates. Shown are the mean rate distributions for the early, middle, and late novelty blocks. Although the average rate distributions did change somewhat, there was no evidence for selection based on average rate. Thus, our data indicate that the population selection rule did not depend on overall activity but was instead related to spatial specificity.

not simply based on the total activity of the cells, but instead favors cells with strong, spatially localized activity. Furthermore, the changes across the novelty blocks differed in CA3. The linear regression on the scatter plots of peak rate versus proportion active showed a increase in slope from the early to the middle block ($p < 0.01$) but no subsequent change between the middle and late blocks (Fig. 8*a*). An inspection of this plot makes it clear that there were only relatively modest changes, in that cells remained active across large fractions of the track in both the early and the late blocks. Furthermore, there was no obvious increase in the number of high peak rate neurons. Instead, the peak rate distributions of CA3 neurons changed only slightly, with a decrease in the proportion of peaks near zero and a slight increase in the proportion of high rate cells (KS test, early vs middle, N.S, early vs late, $p < 0.05$) (Fig. 8*b*).

What governs whether a CA1 cell's peak activity is enhanced or suppressed? If it depends on the current peak rate of the cell, one might expect to see a transition point within the 5–25 Hz peak rate range, where the peak spatial rate of cells with rates below the transition point will tend to decrease, whereas the peak spatial rate of cells with rates above it tend to increase. Ideally, one would record from the same cell across many days to accurately describe this relationship, but as it is very difficult to conclusively demonstrate that the same neurons were recorded across days, we instead analyzed the dynamics of neurons within single recording sessions. We reasoned that if the place field changes are due to a single continuous process, the dynamics of single neurons within a day should be sufficient to predict the changes in the population across days. All animals were exposed to either T1 or T2 twice a day, so we were able to measure differences in activity of single neurons across consecutive runs in an environment. For each cell, we compared the peak spatial rate on the first run to the subsequent within-day change in peak spatial rate. Strikingly, cells in the 5–25 Hz range showed a sharp transition between ~6 and 18 Hz: the peak spatial rate tended to decrease for cells with initial rates <12 Hz and increase for cells with initial rates >12 Hz (Fig. 9*a*). We calculated the likelihood of this transition appearing by chance by shuffling the peak rate/peak rate change relationships for cells in the intermediate group 5000 times and calculating how often a transition whose peak and trough value exceeded those we measured occurred by chance. We found that

this sharp transition was unlikely to arise by chance ($p < 0.01$). These within day dynamics strongly support the claim that the activity of place cells tends to be either enhanced or inhibited based on the strength of spatial tuning. The highly selective nature of these changes also argues against any behavioral explanation for the CA1 results, as there is no evidence for behavioral modulations that affect cells differentially as a function of peak spatial firing rates.

The derived curve represents a system that evolves toward two stable fixed regions. This is illustrated by a diagram of the associated energy surface (Fig. 9*b*) which has one low energy region near zero and another above 20 Hz. To identify the final state of a network that evolves according to this rule we used the actual within-day changes from the middle block to simulate the across-day evolution of the early block population (supplemental Fig. 2, available at www.jneurosci.org as supplemental material). We found that the rate distribution became increasingly different from the initial distribution and converged after ~100 iterations (supplemental Fig. 8, available at www.jneurosci.org as supplemental material). The resulting simulated population distribution was similar to the actual late block distribution and had two main peaks: one centered near zero and the other centered at 31 Hz (Fig. 9*c*). Compared with the early block, the proportion of cells with peak spatial rates <5 Hz increased ($Z = 4.59$, $p < 10^{-5}$), the proportion between 5 and 25 Hz decreased ($Z = 6.98$, $p < 10^{-11}$) and the proportion >25 Hz increased ($Z = 4.06$, $p < 10^{-4}$). This pattern did not arise when we performed the same simulations using randomly scrambled peak rate–rate change relationships from the middle block (supplemental Fig. 9, available at www.jneurosci.org as supplemental material). Thus, using within-day changes, we were able to identify the main pattern of dynamics in the CA1 network and then use these dynamics to predict the actual changes that occur over the course of many days of experience.

To test whether the simulated change in the peak rate distribution predicts the observed decrease in population activity, we made the conservative assumption that the mean rate of each cell changes proportionately to the peak rate. This assumption ignores dynamics related to the sharpening of place fields, in which out-of-field rates decrease faster than within-field rates, and therefore likely causes an underestimate of the magnitude of decrease. Nonetheless, the simulation predicted a decrease in the mean population rate (from 0.84 to 0.62 Hz, rank-sum $p < 10^{-5}$). As might be expected given the numerous factors that are likely to govern neural dynamics across days, the precise quantitative outcomes of the simulation and the actual data were not identical. Despite this, the types of changes seen were very similar. Thus, a rule that selectively suppresses neurons with low spatial specificity and preserves neurons with high spatial specificity captures the key properties of the long term plasticity seen in the CA1 population.

Discussion

We have shown that spatial learning in hippocampal area CA1 invokes a plasticity mechanism that selects a subset of the initially active neurons to remain as part of a longer lasting representation. During the initial exposure to a novel environment, many CA1 cells are active in relatively spatially diffuse patterns, and the population produces more than twice as many spikes as in a familiar environment. This occurs both within the novel environment and during SWRs in a separate rest box. This increase is greater than that seen previously (Nitz and McNaughton, 2004; Csicsvari et al., 2007), perhaps because we included all recorded

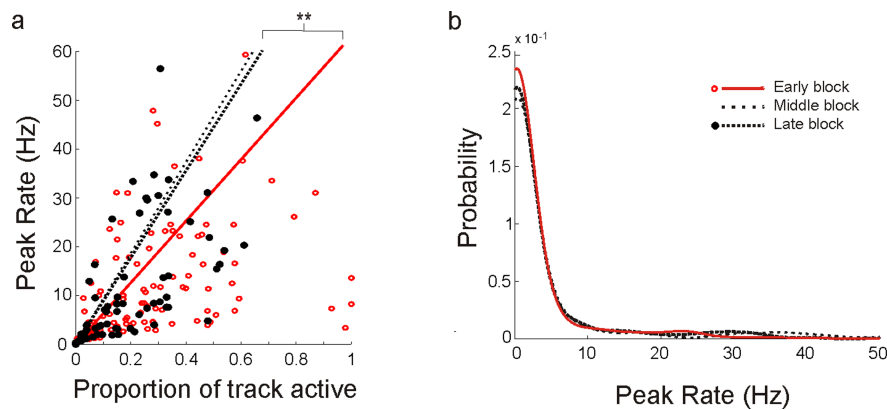


Figure 8. Smaller changes in spatial firing in CA3. *a*, Peak spatial rates versus proportion of the track on which each neuron was active. Regression lines were constrained with a zero y -intercept. Points from middle block were omitted for clarity. There was a significant increase in specificity from the early to the middle block as illustrated by the increase in the slope of the regression line ($p < 0.01$) but not subsequent decrease from the middle to the late block. Overall, cells remained active across large fractions of the track in both the early and the late blocks, and there was no obvious increase in the number of high peak rate neurons. *b*, Peak rate distributions for the early, middle, and late novelty blocks. The peak rate distributions of CA3 neurons changed only slightly, with a decrease in the proportion of peaks near zero and a slight increase in the proportion of high rate cells (KS test, early vs middle, N.S, early vs late, $p < 0.05$). ** $p < 0.01$.

neurons rather than excluding lower rate neurons based on a cutoff. Over time the population firing rate decreases to familiar levels with a time course similar to that of task learning. In the midst of this global suppression, however, the spatial specificity of some neurons increases. A simple rule governs the dynamics: cells that are weakly spatially tuned are suppressed whereas strongly spatially tuned activity is enhanced. The result is a large group of minimally active neurons and a smaller group of well tuned cells with high peak spatial firing rates. The same process is not evident one synapse upstream in CA1, in which the same fraction of neurons is active throughout the animal's experience and there is no evidence for selection based on spatial tuning.

The active selection of high peak spatial firing rate neurons is conceptually similar to the sorts of activity dependent plasticity that increase the specificity of neural tuning during development (Hensch, 2004). Our finding is, to our knowledge, the first demonstration that the strength of tuning predicts changes in neural activity during the formation of representations in the adult. This finding is consistent with previous work showing that the induction of LTP causes larger changes in low peak rate compared with high peak rate place fields (Dragoi et al., 2003), and indicates that CA1 is geared toward creating and preserving high peak rate neurons.

These results cannot easily be explained by any known behavioral modulations of place cell activity. The increases in CA1 population rates were not a result of differences in head direction or activity during ripples. The higher rates also persisted into SWR events seen during the rest after behavior where animals were in a separate rest box. Furthermore, previous results have demonstrated that CA1 firing rate is positively correlated with velocity (McNaughton et al., 1983), which predicts that rates would be lower during initial, low velocity exploration of a novel environment. Thus, our results are the opposite of what would have been expected from the known velocity dependence of place cell firing.

Previous work has demonstrated that CA1 place cells undergo rapid plasticity when an animal enters a new environment (Wilson and McNaughton, 1993; Frank et al., 2004; Leutgeb et al., 2004). These papers suggested that place field plasticity returned to baseline levels after at most 1–2 d of experience in the new

place. Our results confirm the existence of another, slower process that continues after the initial rapid formation of place fields (Austin et al., 1993). The alterations in rate and spatial specificity between the middle and late novelty blocks indicate that the hippocampal representation was dynamic for at least 4–6 d. Although long time scale changes have been observed in the context of exposure to multiple similar environments (Lever et al., 2002), these presence of these effect in a single static environment were less clear.

As a simple within-day rule could capture the essence of the across-day changes, our results also suggest that there is a slow, cumulative plasticity process in CA1. Given the overall similarity between the time course of firing rate change and the time course of learning, we suggest that continued CA1 plasticity is associated with an ongoing learning process. If so, then other within-day changes such as those

seen in the context of order learning (Manns et al., 2007) could in principle accumulate across days to yield a sparse, informative representation of relevant information.

What mechanism could account for the CA1-specific changes we observed? We know that the firing rate of CA1 interneurons drops in novel environments (Nitz and McNaughton, 2004) and then increases as the environment becomes more familiar (Frank et al., 2004). Changes in inhibition likely contribute to the changes we saw, but a simple increase in inhibition would decrease the firing of all principal neurons. Our observation of selective increases in peak spatial rate seen for neurons with peak rates from ~12–18 Hz and the preservation of neurons with peak rates >18 Hz suggests that some other process is involved.

One candidate for such a process was discovered recently in the amygdala. Han et al. (2007) showed that CREB levels during the initial encoding of a fear memory are predictive of which neurons participate in the memory during a subsequent retrieval as measured by expression of the immediate early gene *Arc*. Higher initial levels of CREB during encoding were associated with higher probabilities of activity during retrieval. Given that neural activity leads to CREB expression (Hardingham et al., 2001), we suggest that the amygdala and CA1 could employ similar mechanisms to select highly informative neurons to participate in mnemonic representations. Further, previous work documenting the importance of NMDA receptors (Kentros et al., 1998) and protein synthesis (Agnihotri et al., 2004) for place field stability also suggests that LTP and/or LTD plays a role in the formation of a stable place cell representation. More specifically, our finding suggests that the differences in firing patterns at the 12 Hz boundary may push synapses onto CA1 cells toward either potentiation or depression.

Increases in activity as a result of exposure to new stimuli have also been seen in electrophysiological and imaging studies focused on other brain areas (Li et al., 1993; Knight, 1996; Knight and Nakada, 1998; Kirchhoff et al., 2000). We may then wonder why CA3 activity does not show a similar high firing rate state during the initial stages of spatial learning. We suggest that the relative stability of CA3 in firing rates and place field properties is a result of the role of CA3 in the rapid formation of new representations (Nakazawa et al., 2003; Lee et al., 2004a). Across a

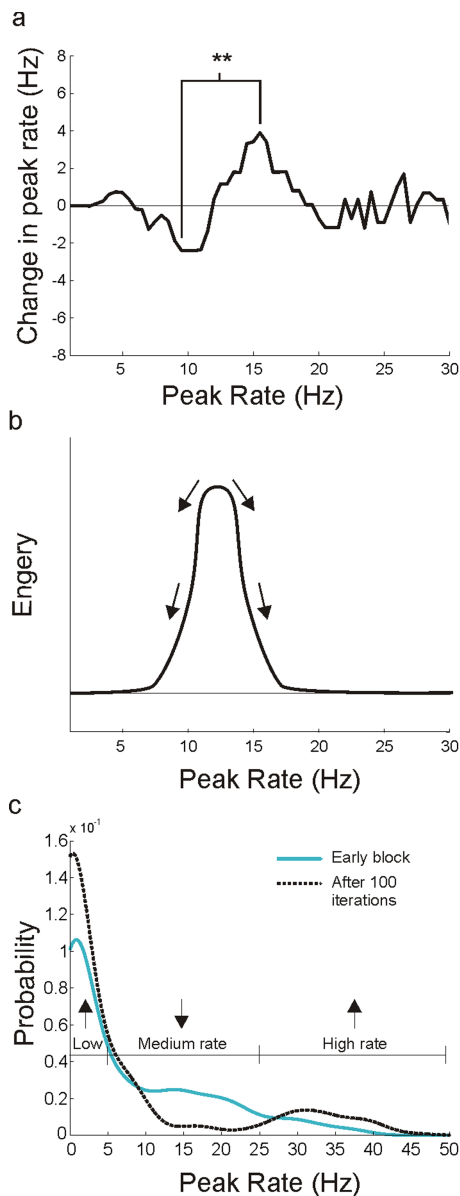


Figure 9. Within-day dynamics predict cross-day evolution. *a*, Peak rate versus median change in peak rate from the first run to the second run in 1 d. Each point represents the median rate change for all data points in a 5 Hz sliding window along the peak rate axis. There was a clear and significant transition from decreasing to increasing peak spatial rates for neurons with peak spatial rates between ~6 and 18 Hz. *b*, A diagram of the associated energy surface, illustrating the dynamics expected from the measured peak rate changes. *c*, Peak rate distributions for the early novelty block and the 100th iteration of the rate dynamics simulation. The final distribution is similar to the distribution of the late block data. $**p < 0.01$.

single exposure to a new place, CA3 place fields develop more slowly than those in CA1 (Leutgeb et al., 2004). Nonetheless, after 20–30 min in the new place, CA3 fields are largely stable. Our data indicate that after this initial period, CA3 remains relatively static compared with CA1. Thus, we conclude that CA3 develops spatial representations quickly and shows only minimal changes during subsequent experiences.

These minimal changes in CA3 imply that CA1 receives approximately the same number of spikes from CA3 throughout the 5–6 d during which CA1 rates fall. In contrast, the subiculum and the entorhinal cortex (EC), the two main targets of CA1 outputs, will receive twice as many spikes from CA1 in a novel compared with a familiar environment. Spikes are metabolically

costly (Laughlin et al., 1998), so this type of transient increase in neural output is likely to serve a functional role. We therefore propose that the high rate of spiking in CA1 could serve to signal the presence of new spatial information and to facilitate plasticity in downstream structures. Theoretical models have suggested that CA1, by virtue of its inputs from CA3 and the EC, is well positioned to “compare” an internal hippocampal representation from CA3 with processed sensory information from the EC (Lisman and Otmakhova, 2001; Kumaran and Maguire, 2007). These models predict that CA1 signals the presence of a new representation from CA3 that is not yet associated with EC inputs; however, the nature of such a signal was not previously known. Our results show that higher firing in CA1 could signal the presence of new information. This could be the physiological basis of the hippocampal “novelty signal” whose strength is predictive of later recall (Kirchhoff et al., 2000; Kumaran and Maguire, 2006).

How then, would the increase in rate cause changes in processing in downstream structures? Previous work has shown that novelty-dependent dopamine release in the nucleus accumbens requires hippocampal activity (Legault and Wise, 2001). Thus, increased CA1 output during both behavior and subsequent rest could trigger broadly distributed modulation to facilitate learning across the brain. Further, our finding that individual CA1 neurons that were highly active during the run are also highly active during SWRs in the rest period shows that this high rate activity persists after the experience. Given the association between SWRs and memory consolidation (Buzsáki, 1986; Buzsáki et al., 1992), our results provide a mechanism whereby newly learned information could be preferentially consolidated after the experience. This complements previous work showing that CA1 cells are reactivated together both during and after a novel experience (Cheng and Frank, 2008; O’Neill et al., 2008). Thus, we suggest that this greater total activity has the potential to increase the influence of CA1 on downstream structures during and after new experiences, facilitating memory storage.

References

- Agnihotri NT, Hawkins RD, Kandel ER, Kentros C (2004) The long-term stability of new hippocampal place fields requires new protein synthesis. *Proc Natl Acad Sci U S A* 101:3656–3661.
- Alvernhe A, Van Cauter T, Save E, Poucet B (2008) Different CA1 and CA3 representations of novel routes in a shortcut situation. *J Neurosci* 28:7324–7333.
- Austin KB, White LH, Shapiro ML (1993) Short- and long-term effects of experience on hippocampal place fields. *Soc Neurosci Abstr* 19:325.6.
- Battaglia FP, Sutherland GR, McNaughton BL (2004) Local sensory cues and place cell directionality: additional evidence of prospective coding in the hippocampus. *J Neurosci* 24:4541–4550.
- Best PJ, White AM, Minai A (2001) Spatial processing in the brain: the activity of hippocampal place cells. *Annu Rev Neurosci* 24:459–486.
- Buzsáki G (1986) Hippocampal sharp waves: their origin and significance. *Brain Res* 398:242–252.
- Buzsáki G, Horváth Z, Urioste R, Hetke J, Wise K (1992) High-frequency network oscillation in the hippocampus. *Science* 256:1025–1027.
- Cheng S, Frank LM (2008) New experiences enhance coordinated neural activity in the hippocampus. *Neuron* 57:303–313.
- Csicsvari J, O’Neill J, Allen K, Senior T (2007) Place-selective firing contributes to the reverse-order reactivation of CA1 pyramidal cells during sharp waves in open-field exploration. *Eur J Neurosci* 26:704–716.
- Dragoi G, Harris KD, Buzsáki G (2003) Place representation within hippocampal networks is modified by long-term potentiation. *Neuron* 39:843–853.
- Eichenbaum H, Cohen NJ (2001) From conditioning to conscious recollection. New York: Oxford UP.
- Fox SE, Ranck JB Jr (1981) Electrophysiological characteristics of hippocampal complex-spike cells and theta cells. *Exp Brain Res* 41:399–410.

- Frank LM, Brown EN, Wilson MA (2000) Trajectory encoding in the hippocampus and entorhinal cortex. *Neuron* 27:169–178.
- Frank LM, Brown EN, Wilson MA (2001) A comparison of the firing properties of putative excitatory and inhibitory neurons from CA1 and the entorhinal cortex. *J Neurophysiol* 86:2029–2040.
- Frank LM, Stanley GB, Brown EN (2004) Hippocampal plasticity across multiple days of exposure to novel environments. *J Neurosci* 24:7681–7689.
- Gelbard-Sagiv H, Mukamel R, Harel M, Malach R, Fried I (2008) Internally generated reactivation of single neurons in human hippocampus during free recall. *Science* 322:96–101.
- Guzowski JF, Knierim JJ, Moser EI (2004) Ensemble dynamics of hippocampal regions CA3 and CA1. *Neuron* 44:581–584.
- Han JH, Kushner SA, Yiu AP, Cole CJ, Matynia A, Brown RA, Neve RL, Guzowski JF, Silva AJ, Josselyn SA (2007) Neuronal competition and selection during memory formation. *Science* 316:457–460.
- Hardingham GE, Arnold FJ, Bading H (2001) Nuclear calcium signaling controls CREB-mediated gene expression triggered by synaptic activity. *Nat Neurosci* 4:261–267.
- Hensch TK (2004) Critical period regulation. *Annu Rev Neurosci* 27:549–579.
- Ishizuka N, Weber J, Amaral DG (1990) Organization of intrahippocampal projections originating from CA3 pyramidal cells in the rat. *J Comp Neurol* 295:580–623.
- Kentros C, Hargreaves E, Hawkins RD, Kandel ER, Shapiro M, Muller RV (1998) Abolition of long-term stability of new hippocampal place cell maps by NMDA receptor blockade. *Science* 280:2121–2126.
- Kirchhoff BA, Wagner AD, Maril A, Stern CE (2000) Prefrontal-temporal circuitry for episodic encoding and subsequent memory. *J Neurosci* 20:6173–6180.
- Knight R (1996) Contribution of human hippocampal region to novelty detection. *Nature* 383:256–259.
- Knight RT, Nakada T (1998) Cortico-limbic circuits and novelty: a review of EEG and blood flow data. *Rev Neurosci* 9:57–70.
- Kumaran D, Maguire EA (2006) An unexpected sequence of events: mismatch detection in the human hippocampus. *PLoS Biol* 4:e424.
- Kumaran D, Maguire EA (2007) Match mismatch processes underlie human hippocampal responses to associative novelty. *J Neurosci* 27:8517–8524.
- Laughlin SB, de Ruyter van Steveninck RR, Anderson JC (1998) The metabolic cost of neural information. *Nat Neurosci* 1:36–41.
- Lee I, Rao G, Knierim JJ (2004a) A double dissociation between hippocampal subfields: differential time course of CA3 and CA1 place cells for processing changed environments. *Neuron* 42:803–815.
- Lee I, Yoganarasimha D, Rao G, Knierim JJ (2004b) Comparison of population coherence of place cells in hippocampal subfields CA1 and CA3. *Nature* 430:456–459.
- Legault M, Wise RA (2001) Novelty-evoked elevations of nucleus accumbens dopamine: dependence on impulse flow from the ventral subiculum and glutamatergic neurotransmission in the ventral tegmental area. *Eur J Neurosci* 13:819–828.
- Leutgeb S, Leutgeb JK, Treves A, Moser MB, Moser EI (2004) Distinct ensemble codes in hippocampal areas CA3 and CA1. *Science* 305:1295–1298.
- Lever C, Wills T, Cacucci F, Burgess N, O'Keefe J (2002) Long-term plasticity in hippocampal place-cell representation of environmental geometry. *Nature* 416:90–94.
- Li L, Miller EK, Desimone R (1993) The representation of stimulus familiarity in anterior inferior temporal cortex. *J Neurophysiol* 69:1918–1929.
- Lisman JE, Otmakhova NA (2001) Storage, recall, and novelty detection of sequences by the hippocampus: elaborating on the SOCRATIC model to account for normal and aberrant effects of dopamine. *Hippocampus* 11:551–568.
- Manns JR, Howard MW, Eichenbaum H (2007) Gradual changes in hippocampal activity support remembering the order of events. *Neuron* 56:530–540.
- McNaughton BL, Barnes CA, 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.
- Nakazawa K, Sun LD, Quirk MC, Rondí-Reig L, Wilson MA, Tonegawa S (2003) Hippocampal CA3 NMDA receptors are crucial for memory acquisition of one-time experience. *Neuron* 38:305–315.
- Nitz D, McNaughton B (2004) Differential modulation of CA1 and dentate gyrus interneurons during exploration of novel environments. *J Neurophysiol* 91:863–872.
- O'Keefe J, 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'Neill J, Senior TJ, Allen K, Huxter JR, Csicsvari J (2008) Reactivation of experience-dependent cell assembly patterns in the hippocampus. *Nat Neurosci* 11:209–215.
- Scoville WB, Milner B (1957) Loss of recent memory after bilateral hippocampal lesions. *J Neurol Neurosurg Psychiatry* 20:11–21.
- Smith AC, Frank LM, Wirth S, Yanike M, Hu D, Kubota Y, Graybiel AM, Suzuki WA, Brown EN (2004) Dynamic analysis of learning in behavioral experiments. *J Neurosci* 24:447–461.
- Squire LR, Zola-Morgan S (1991) The medial temporal lobe memory system. *Science* 253:1380–1386.
- Sutherland GR, McNaughton B (2000) Memory trace reactivation in hippocampal and neocortical neuronal ensembles. *Curr Opin Neurobiol* 10:180–186.
- Wilson MA, McNaughton BL (1993) Dynamics of the hippocampal ensemble code for space. *Science* 261:1055–1058.