

Robust Conjunctive Item–Place Coding by Hippocampal Neurons Parallels Learning What Happens Where

Robert W. Komorowski,¹ Joseph R. Manns,² and Howard Eichenbaum¹

¹Center for Memory and Brain, Boston University, Boston, Massachusetts 02215, and ²Department of Psychology, Emory University, Atlanta, Georgia 30322

Previous research indicates a critical role of the hippocampus in memory for events in the context in which they occur. However, studies to date have not provided compelling evidence that hippocampal neurons encode event–context conjunctions directly associated with this kind of learning. Here we report that, as animals learn different meanings for items in distinct contexts, individual hippocampal neurons develop responses to specific stimuli in the places where they have differential significance. Furthermore, this conjunctive coding evolves in the form of enhanced item-specific responses within a subset of the preexisting spatial representation. These findings support the view that conjunctive representations in the hippocampus underlie the acquisition of context-specific memories.

Introduction

Recent theories about the functional organization of the medial temporal lobe memory system have focused on distinct cortical streams of “what” and “where” information converging within the hippocampus, which combines this information to generate representations of salient items (“what”) in the places (“where”) they occurred (Davachi, 2006; Manns and Eichenbaum, 2006; Diana et al., 2007; Eichenbaum et al., 2007).

Several studies have reported that hippocampal neurons fire in association with combinations of specific memory cues and the locations in which they are presented (for review, see Eichenbaum, 2004). However, the appearance of neurons that encode item–place conjunctions has not been directly related to the learning of item and place associations. In addition, the prevalence of item–place conjunctive activity is typically quite low compared with spatially specific firing, leading many to the alternative view that the hippocampal item coding is coincidental to a primary representation of maps and routes (O’Keefe, 2007).

Functional imaging studies have also identified selective hippocampal activation related to memory for item–place associations (Henke et al., 1997; Davachi et al., 2003; Hannula and Ranganath, 2008). However, these studies are unable to characterize the specific information that drives this activation, raising the question as to whether the increased activity reflects the formation of specific item–place associations or just generally stronger memory (Squire et al., 2007). Characterization of single neuron activity in the hippocampus during a task in which

item–place associations are learned could determine whether firing rates increase separately to items and to locations, reflecting generally increased memory, or whether firing rates increase specifically to item–place conjunctions, reflecting memory for these associations.

Here we recorded from hippocampal principal neurons in rats learning which of two items is rewarded depending on the environmental context in which they were presented (Rajji et al., 2006). We observed that a large percentage of hippocampal neurons developed representations of task-relevant item–place associations, and their evolution was closely correlated with learning those associations. Furthermore, the item–place representations developed from preexisting spatial representations into enhanced activations when particular items were sampled in specific locations. Conversely, the representation of the items alone was minimal throughout learning, and the representation of places where any object was sampled, although strong, remained unchanged throughout training. These findings join the phenomenology of place cells with learning what happens where and support the hypothesis that the development of conjunctive representations within the hippocampus underlies memories for items in the places where they occur.

Materials and Methods

Five 400–450 g male Long–Evans rats were maintained at a minimum of 85% of normal body weight. Each rat was initially shaped to dig in 10-cm-tall, 11-cm-wide terra cotta pots filled with common playground sand (QUIKRETE Premium Play Sand) for one-quarter Froot Loop (Kellogg’s) pieces. Then animals were trained on a simple discrimination between two pots filled with sand and scented with different oil fragrances [aloe (Jason Natural Products) vs clove (AuraCacia)] placed side by side simultaneously in the home cage. All oil fragrances were diluted at a 2% concentration within vegetable oil (Wesson). The left–right positions of the two stimuli were pseudorandomized, although they were never located in the same positions on more than three consecutive trials. The rat could dig in the aloe-scented pot for one-quarter of a Froot Loop, but digging in the clove-scented pot resulted in a 5 s timeout. Rats achieved a performance criterion of 8 correct of 10 consecutive trials within a single session.

Received March 22, 2009; revised June 22, 2009; accepted July 3, 2009.

This work was supported by National Institutes of Health/National Institute of Mental Health (NIH/NIMH) Grant MH051570 and Silvio O. Conte Center Grant NIH/NIMH MH71702. We thank Zachary Haberman and Carolyn Garcia for help with behavioral testing, Nancy Kopell for helpful discussions regarding data analysis, and Adriano Tort for help with shuffling analyses and rate mapping plots. Christopher Andrews and Murat Okatan contributed programs to assist with animal tracking and video scoring.

Correspondence should be addressed to Dr. Howard Eichenbaum, Center for Memory and Brain, Boston University, Boston, MA 02215. E-mail: hbe@bu.edu.

DOI:10.1523/JNEUROSCI.1378-09.2009

Copyright © 2009 Society for Neuroscience 0270-6474/09/299918-12\$15.00/0

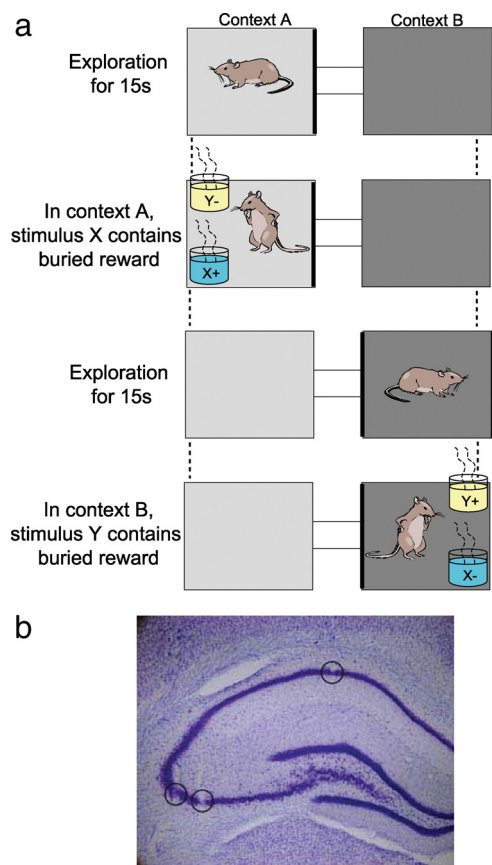


Figure 1. *a*, Conditional discrimination task. The two contexts (represented by different shadings) differed in their flooring (wood vs black paper for the initial conditional discrimination and rubber vs sandpaper for the second discrimination problem) and in the walls (white vs black paper for the initial conditional discrimination and vertical vs horizontal stripes for the second). The stimulus objects (X or Y) differed in odor and in the medium that filled the pots (shown as blue and yellow). *b*, Lesion marks made at tips of tetrodes, one located in CA1 and two others in the CA3 region of hippocampus.

The next day, rats were exposed for 15 min to the environment where conditional discrimination testing would take place and were allowed to forage for food scattered on the floor of the apparatus. The environment consisted of two 37×37 cm boxes connected by a central alley that allowed the rat to shuttle between them (Fig. 1*a*). The entrance to the central alley could be closed with dividers to block the animal within either box. Each box differed substantially in contextual cues that included different flooring (wood vs black paper) and different wallpaper (white paper vs black paper). Rats were trained to alternate between the two contexts by traversing the central alleyway when the dividers were lifted. On each trial of the conditional discrimination task, the rat was allowed to enter a context, after which a divider would close and the animal was permitted to explore the contextual cues for 15 s (Fig. 1*a*). The animal was then blocked within one side of the context using another divider, and two items were placed in different corners of that box. Both items were terra cotta pots, each scented with a different odor (grapefruit or geranium; both from AuraCacia) and filled with a different digging media [white foam pieces (Foamies Rol Darice) or 5 mm metallic purple beads]. The positions of the items were pseudorandomized such that they appeared in each position equally but never in the same position on more than three consecutive trials. In context A, item X (grapefruit–white foam) contained a Froot Loop reward, whereas in context B, item Y (geranium–purple beads) contained the reward. Digging in the correct pot yielded the buried reward, but digging in the incorrect pot resulted in the removal of both pots and a 5 s timeout.

On day 1 of the conditional discrimination, we trained the rat for as many trials as possible (usually 50–60 trials) in blocks of five trials within

the same context, to allow for corrections of that particular discrimination. On day 2, trials alternated between the contexts, always permitting a 15 s exploratory period before item presentation. To ensure that the animal did not simply learn to alternate choices of items X and Y, two successive trials were presented within the same context every 10 trials on average. To ensure that the animal could not simply smell the buried reward, every 10 trials involved a probe in which neither pot contained a reward, and a reward was given only after digging in the correct pot. Initial training on this first conditional discrimination problem required 3–5 d of 80 trials per day until performance reached at least 70% in concurrent 10 trial blocks within each context.

After reaching the performance criterion, rats were implanted with a recording head stage above the left dorsal hippocampus centered at 3.6 mm posterior and 2.9 mm lateral to bregma. The head stage contained 12–18 independently movable tetrodes aimed at CA1 and CA3. Each tetrode was composed of four $12.5 \mu\text{m}$ nichrome wires with the tips plated with gold to bring the impedance to $200 \text{ k}\Omega$ at 1 kHz. Animals recovered for 7–10 d, after which the tetrodes were moved down slowly over the course of 1–2 weeks, until the tips reached the pyramidal cell layer of CA1 or CA3 and the animal's performance had again met criterion level on the initial conditional discrimination problem. The locations of these tetrodes was estimated *in vivo* using driver turn counts as well as electrophysiological events, including the appearance of complex spikes, theta-modulated spiking, and the presence of theta and high-frequency ripples in the local field potentials (LFPs). At the end of data collection, electrode location was confirmed by passing a $25 \mu\text{A}$ current for 20 s through each tetrode immediately before perfusion to create a lesion visible after histological processing with Nissl stain (Fig. 1*b*).

Once the tetrodes were in the desired locations, recordings were taken as the rats continued to perform the initial conditional discrimination problem. We defined overtraining sessions as sessions in which the animal's performance had exceeded 80% for 3 consecutive previous testing days. After these overtraining sessions, the animal was introduced to a novel environment with the same general configuration but with new flooring (rubber or sandpaper) and new wallpaper (vertical or horizontal stripes) defining each context. After 15 min of exposure to this environment, we began testing the animal on a second conditional discrimination problem using pots with new scented oils and digging media (patchouli–straws vs mint–buttons; both from Aromaland). Rats learned this second conditional discrimination problem within a single recording session (see Results). On subsequent days, recordings were taken during overtraining on the second problem after the animal had again reached the criterion of 80% correct performance for 3 consecutive testing days.

During all recording sessions, spike activity was amplified ($10,000\times$), filtered (600–6000 Hz), and saved for offline analysis using the software Spike (written by Loren Frank, University of California, San Francisco, San Francisco, CA). Cells from each tetrode were analyzed from only one learning session and one overtraining session to avoid the inclusion of the same neuron more than once in the each type of session. Clusters of single-unit activity were isolated offline and determined to be stable pyramidal units using various three-dimension projections (spike peak, valley, principal components, and timestamps) provided by Offline Sorter (Plexon). In addition, behavior was recorded with digital video (30 frames/s) that was synchronized with the acquisition of neural data, and the animal's location was tracked with one or two light-emitting diodes located on the recording head stages. The onset of stimulus sampling was defined by scanning the video and manually marking the frame on which the rat's nose crossed the rim of the pot. Timestamps for the onset of stimulus sampling and for spikes were imported into Matlab for subsequent data analyses.

We quantified that rat's behavior leading up to and during stimulus sampling with several measures. The angle of approach to the pots was measured for the 1 s preceding stimulus sampling by calculating the angle between the animal's location at the start and end of this time period. To examine each rat's behavior during item sampling, we calculated head direction every 0.1 s during the 1 s of stimulus investigation. The angles for both measures were then converted to linear measurements by calculating the difference between the measured angle and the horizontal

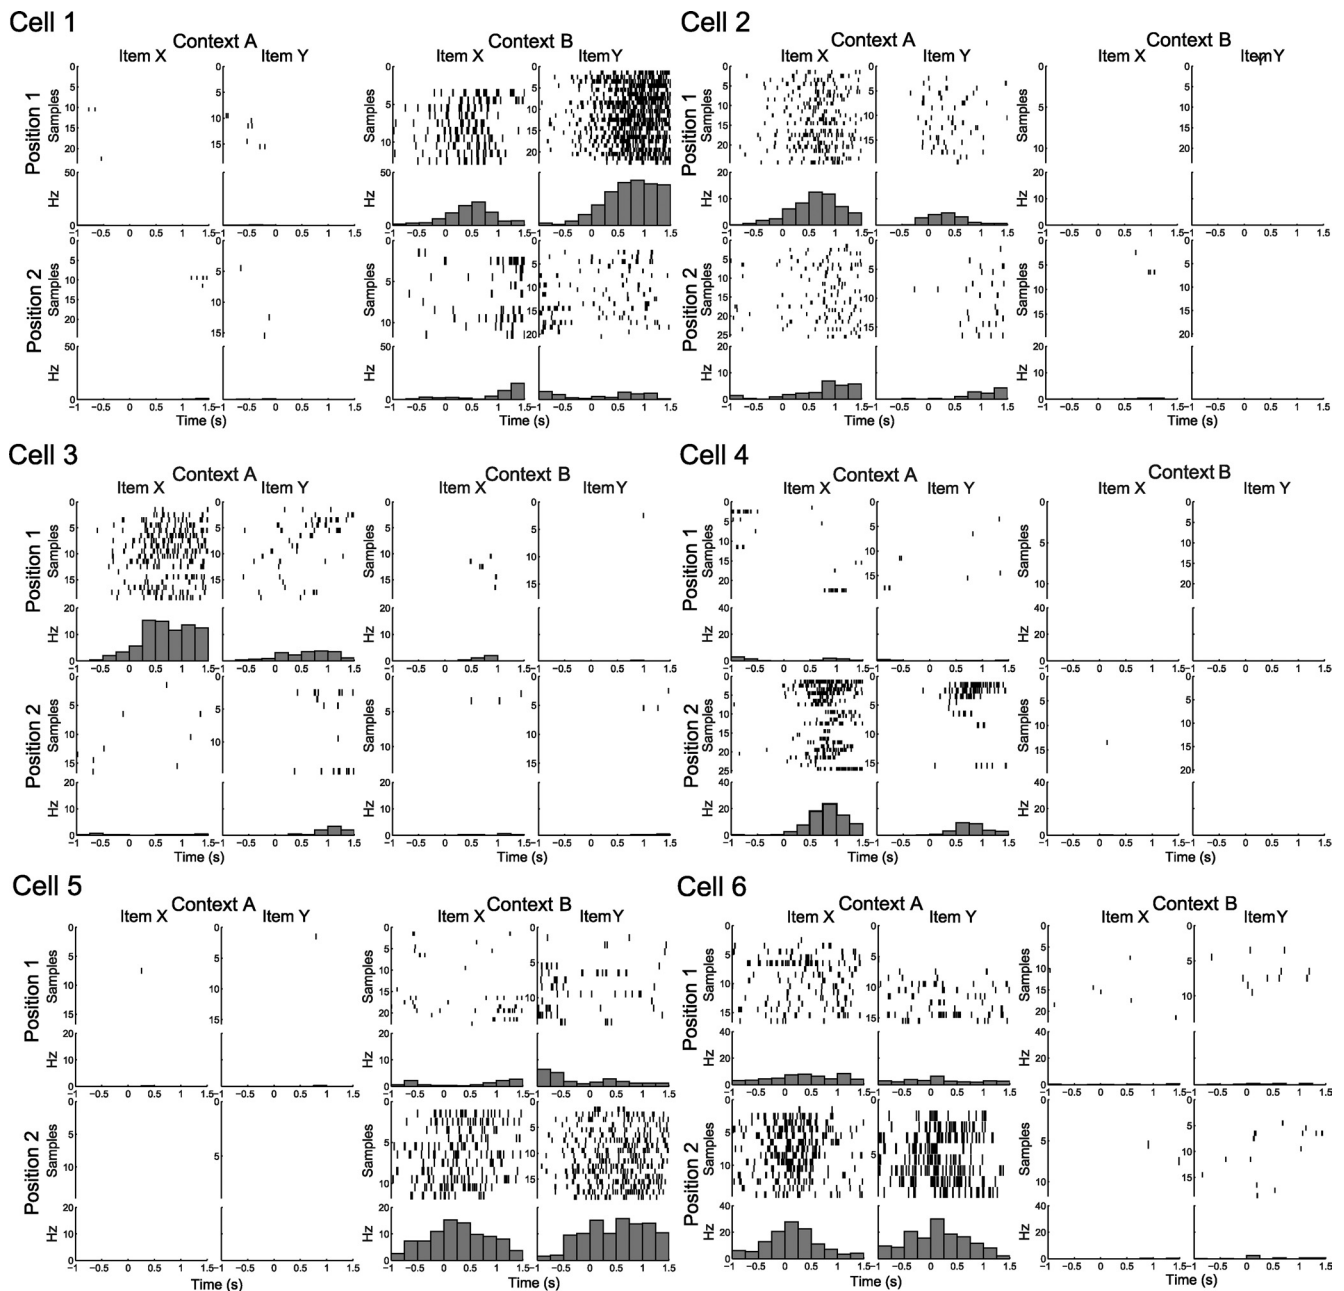


Figure 2. Example cells recorded from a single animal. Rasters and perievent histograms are plotted for stimulus sampling events for each item (X or Y) at each position (1 or 2) within each context (A or B). Each line of the raster represents a single stimulus sample ordered from 0 to the maximum number of samples for each particular item–position combination for all trials, correct and incorrect. Time point 0 denotes the time when the animal's nose crosses the edge of the stimulus pot, and each bar of the histogram represents the average activity in hertz for a 250 ms time window. Cells 1 and 2 are examples of item–position cells recorded during an overtraining session, whereas cells 3 and 4 are examples of item–position cells during a learning session. Cells 5 and 6 are examples of position cells during the same learning session.

plane, which allowed us to use standard N -way ANOVAs. In addition, measurements of the animal's location during stimulus sampling were calculated as the average distance between the animal's location and the center of the pot at each location.

Spatial firing rate maps were estimated using the total number of spikes that occurred when the rat was at a given location (2×2 cm bins) divided by the total time spent in that bin. The smoothed value for each bin was then calculated as the mean for each bin and all bins within 5 cm, each weighted by the distance from the central bin as determined by a two-dimensional Gaussian kernel.

Ripples during stimulus sampling were identified using methods similar to Foster and Wilson (2006). Specifically, the LFPs from tetrodes recorded in CA1 or CA3 were between 100 and 400 Hz. The mean and

SDs for this trace were calculated with a threshold set at 3 SDs above the mean. Threshold crossings <150 ms of each other were identified as a single ripple event.

Item and position selectivity for each cell was measured for 30 trial blocks using a selectivity index (SI) calculated as

$$SI = (n - \sum_{i=1}^n (\lambda_i / \lambda_{\text{pref}})) / (n - 1),$$

where n is the total number of possible stimulus sampling events (two in the case of items and four in the case of positions), λ_i is the firing rate of the neuron within a block to the i th possible event, and λ_{pref} is the firing rate of the neuron to its preferred item, or place, defined event within the

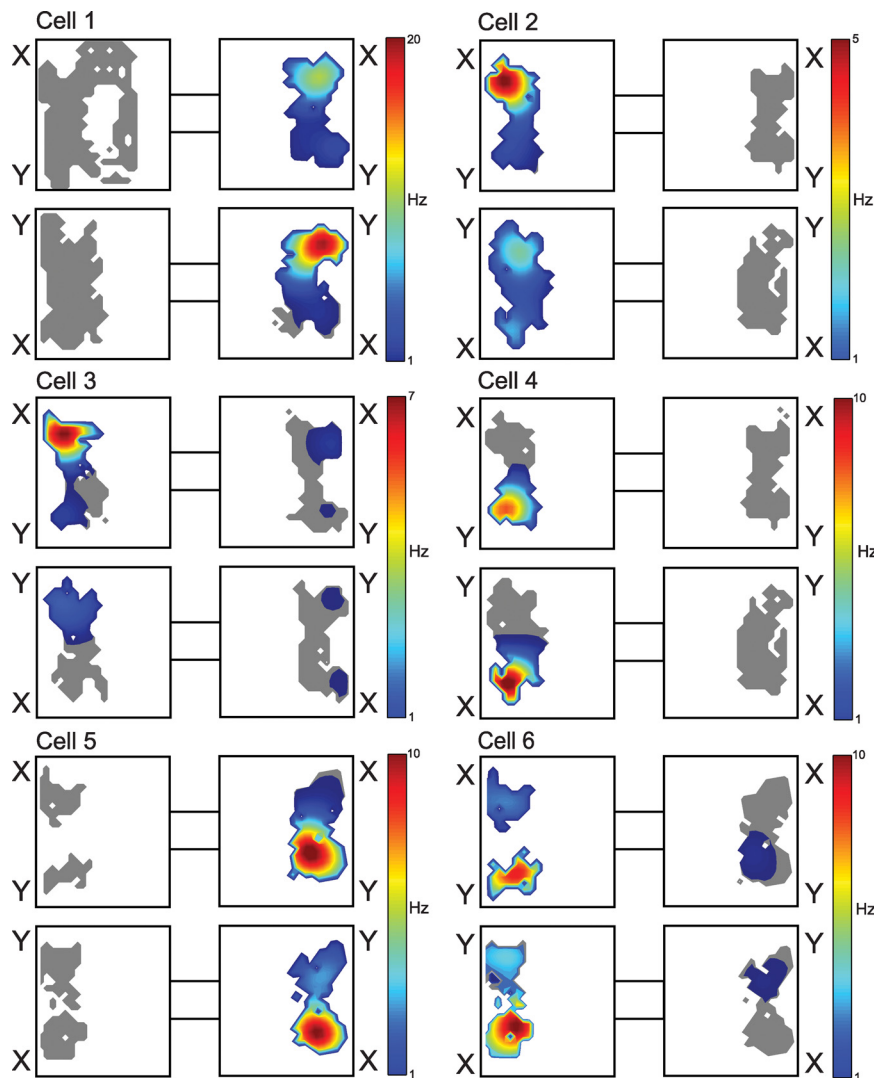


Figure 3. Spatial distributions of firing rates for the same cells for which raster plots are shown in Figure 2. Each panel includes spike activity recorded from the onset of the first stimulus sampling event until 1 s after the onset of the last stimulus sampling event in the trial, during which the rat can sample each stimulus multiple times before digging for the reward. Stimulus identities are indicated just outside their locations within the environment. Color-coded firing rates are indicated in the legend to the right of the plots for each cell. Gray indicates visited areas associated with no neural activity.

same block of trials (Moody et al., 1998; Wirth et al., 2003). The preferred item or place of a cell was defined as the item or place that elicited the highest firing within the trial block. An item SI = 1 if a cell fired only in response to one of the two items and did so only at the preferred stimulus location of that cell. A position SI = 1 if a cell fired only when the rat sampled stimuli at one of the four positions. Conversely, item or position SI = 0 if the cell fired equally for both stimuli at the preferred location or at all positions. To test whether the SI values were larger than that expected by chance, we compared each observed item and position SI values against a distribution of 10,000 SI scores in which the item or position identities of the same dataset were randomly shuffled; then a z-test was used to determine the significance of each observed SI value.

Firing rates during stimulus sampling were normalized into z-scores to compare average normalized firing rate of specific neural populations between conditions. For each cell, a z-score was calculated as the difference between a particular condition and the average firing rate across all stimulus samples, divided by the SD.

Results

Extracellular spike activity of CA1 and CA3 pyramidal neurons was recorded from five rats as they performed a conditional dis-

crimination task that required them to select one of two items (X or Y) within each of two distinctive spatial contexts (A or B) in which both stimuli were presented. Although odor was a primary defining feature, the stimuli differed in multiple modalities (see Materials and Methods), and thus we will refer to them more broadly as “items.” Specifically, item X was rewarded when it appeared within either of two positions within spatial context A, and item Y was rewarded when it appeared within either of two positions within spatial context B (Fig. 1). Rats reached the performance criterion of 70% correct in each spatial context within a 20 trial block on average by trial 58 (range, 42–71 trials). For most subsequent analyses, the learning session was divided into 30 trial blocks: the first 30 trials, 30 trials centered on the middle of the criterion trial block, and the last 30 trials. Average performance accuracy improved from 45.5 ± 3.6% on the first 30 trials, to 79.4 ± 3.4% on the middle 30 trials, to 84.8 ± 3.4% on the final 30 trials.

Hippocampal neural activity was also recorded in sessions in which the animals were already highly overtrained on the second conditional discrimination. These overtraining sessions were defined as sessions in which the rat’s performance had been >80% for at least three consecutive preceding sessions on the same problem. A total of five overtraining sessions were analyzed across three animals, three of which involved a preliminary conditional discrimination problem given before the learning session and two of which occurred after multiple days of training on the learning problem (see Materials and Methods). No differences in performance were noted between these overtraining sessions, so the data from all of them were combined. Overtraining sessions were also divided into 30 trial blocks, but here the middle 30 trials were centered at the middle of the session (because there was no “learning” phase on which to center the middle block). Average performance accuracy was 86.1 ± 3.8% on the first 30 trials, 97.2 ± 1.4% on the middle 30 trials, and 97.8 ± 1.1% on the final 30 trials of overtraining sessions. A repeated measures ANOVA showed that performance increased significantly in both learning and overtraining sessions ($F_{(2,8)} = 103.93, p < 0.001$ for learning sessions; $F_{(2,8)} = 12.11, p = 0.004$ for overtraining sessions) and more so for the learning sessions (two-way ANOVA interaction of performance with trial block, $F_{(2,16)} = 23.25, p < 0.001$).

A total of 198 pyramidal neurons were isolated among five learning sessions in five rats, composed of 141 CA3 cells and 57 CA1 cells. An additional 189 pyramidal neurons were recorded in five overtraining sessions from three rats, composed of 161 CA3 cells and 28 CA1 cells. We failed to find any difference between firing patterns in CA1 and CA3 neurons, so the data were combined for current analyses.

Figure 2 shows raster plots and perievent histograms showing the firing patterns of six example cells. For each cell, separate raster and histogram panels include data for one of the eight combinations of one of the two items sampled in one of two positions within one of the two contexts. Each panel plots spike activity centered around the moment when the rat's nose crossed the rim of a stimulus pot (indicated by time 0 on the *x*-axis), each raster line represents a single stimulus sampling event, and there could be multiple sampling events of either or both items on a trial. Examination of these analyses identified two major types of cells. One type, which we refer to as item–position cells, increased firing when rats sampled one of the two items at one or both positions in one of the contexts. For example, in an overtraining session, cell 1 fired maximally when the rat sampled item Y at position 1 within context B, whereas cell 2 fired maximally when the rat sampled item X at position 1 within context A. We also observed that many item–position cells developed their selectivity for one of the items over the course of learning. Inspection of the sequence of rasters (scan from top to bottom) for cell 3 at position 1 within context A suggests equivalent activation for the first several trials as the rat samples items X and Y, but the cell evolves in subsequent trials to fire selectively during the sampling of item X. Cell 4 shows a similar pattern at position 2 within context A, in which in early trials the cell fires to both items X and Y and then comes to prefer item X. The other class of cell, which we call position cells, fired equivalently when rats sampled either item at one or both positions within one of the contexts. For example, in an overtraining session, cell 5 fired equivalently when a rat sampled either item at position 2 within context B; cell 6 fired equivalently when the rats sampled either item at position 2 within context A.

To visualize the activity patterns of these cells from the perspective of the well documented spatial firing properties of hippocampal principal cells (O'Keefe, 2007), Figure 3 provides standard spatial firing distributions that include the period from the outset of stimulus sampling until 1 s after the response choice is made for each of the cells described in Figure 2. For each item–position cell (cells 1–4), the highest firing rate was consistently centered over the location of the stimulus pot associated with maximal activation in the raster displays and perievent histograms. These spatial plots indicate that firing rates were lower when the less preferred stimulus was at the same location. These cells were largely inactive when the rat was in the nonpreferred environment. For position cells (cells 5 and 6), high firing rates were observed at one pot location regardless of which stimulus was present. These findings

Learning

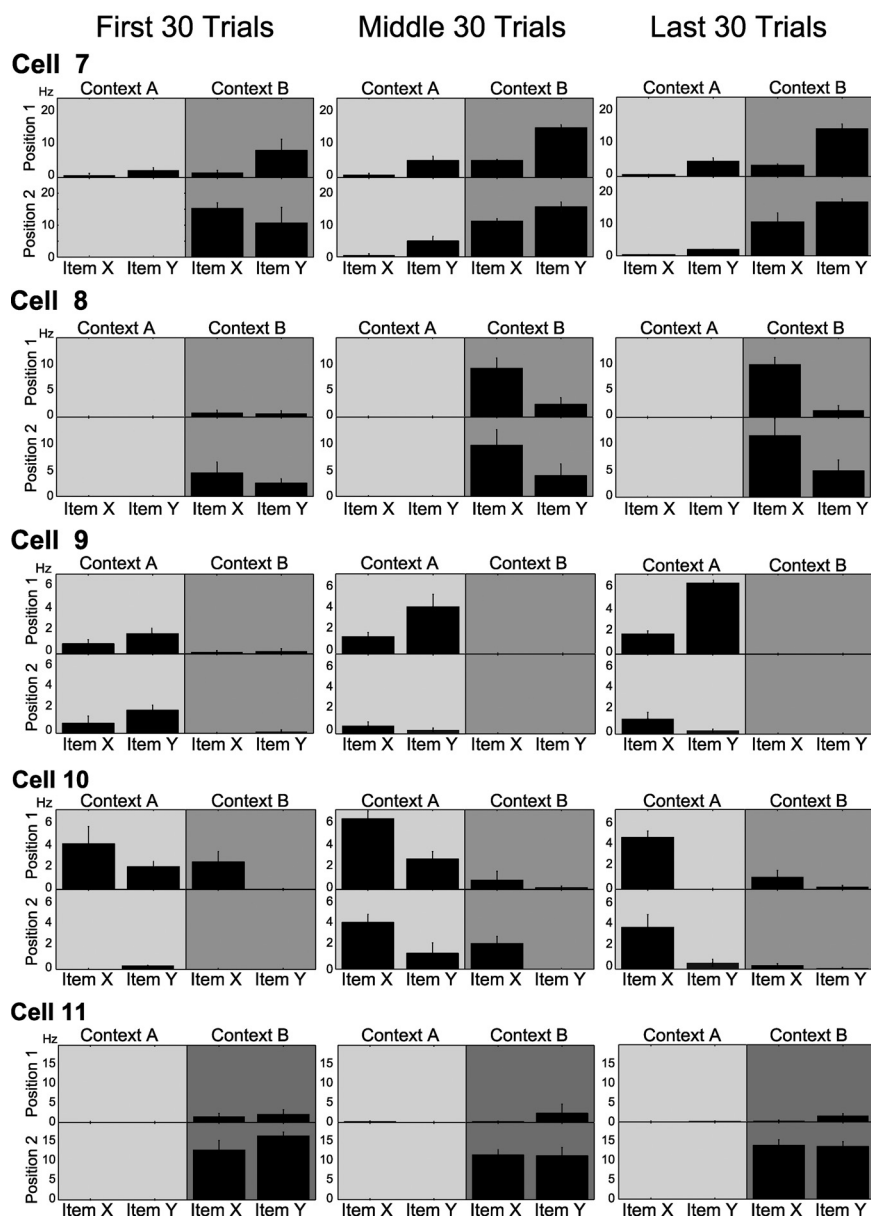


Figure 4. Example cells recorded during a learning session. Panels in each row show the average \pm SE firing rate of a cell during stimulus sampling for early, criterion, and late trial blocks. Each panel shows the average firing rate during sampling of items X and Y at each of two positions within a context (represented with different shades of gray). Cells 7–10 were classified as item–position cells in the last 30 trial block but not in the first 30 trials. Cell 11 was classified as a position cell in all trial blocks.

are entirely consistent with the data shown in the raster plots and confirm the distinctions between item–position and position cells as described above.

Our visual inspections of the raster displays indicated that maximum activation for both cell types always occurred during the 1 s period after the rat's nose crossed the rim of a stimulus pot. In addition, during this period, the animal's behavior was consistent, such that they were relatively immobile before digging for rewards or before moving away from the pot. Therefore, we selected this 1 s period as the time window for comparisons of neural responses across conditions. Typically, rats sampled each stimulus one to two times per trial, and we considered all neurons that generated at least 10 spikes total during all stimulus sampling

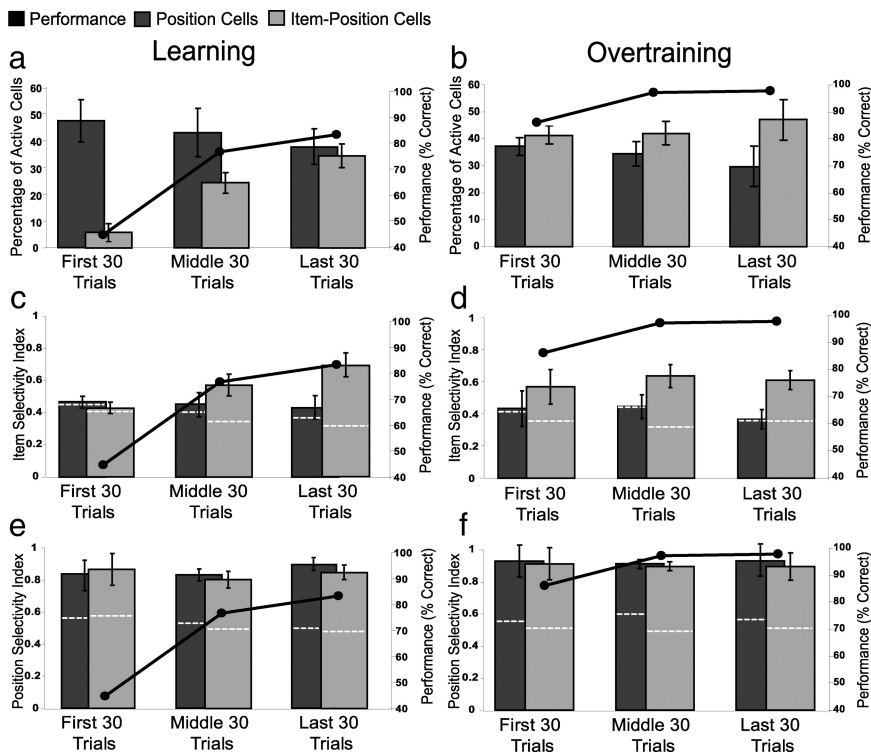


Figure 5. Changes in proportions of item–position and position cells in learning (**a**) versus overtraining (**b**) sessions. Changes in selectivity for items during learning and overtraining sessions (**c**, **d**) and positions for the same sessions (**e**, **f**). The selectivity index was calculated using firing rates for each cell averaged across all stimulus samples. Bars represent the average \pm SE proportions of cells in each 30 trial block, whereas the line shows average performance over the same trial blocks. The dotted white lines indicate average values calculated from the shuffling analyses.

events. Based on this criterion, 56 neurons were active during stimulus sampling in learning sessions, whereas 52 were active during stimulus sampling in overtraining sessions. Two-way ANOVAs were used to compare the firing rates of neurons that were active during stimulus sampling with the identity of the item (X or Y) and its position (two positions within each context) as main factors with a significance criteria of $p < 0.05$. Activity was compared across the blocks of 30 trials as defined above, for both learning and overtraining sessions. These quantitative analyses revealed only one cell that fired differentially to the two stimuli absent an interaction with position in any of the 30 trial blocks. Thus, just one cell appeared to encode only the item itself, and so this class of cell will not be considered further.

Item–position cells develop during training, whereas the number and selectivity of position cells is consistent throughout training

Our quantitative analyses showed that the amount of item–position coding increased dramatically during learning. The first four rows in Figure 4 show cells that eventually fired differentially associated with stimulus and position. Cell 7 fired during stimulus sampling within context B regardless of the stimulus in the first 30 trials. However, as performance improved in the middle and last 30 trials, the same cell developed a higher firing rate during the sampling of item Y in context B. Cell 8 also fired during stimulus sampling in context B regardless of the stimulus in the first 30 trials. However, in the middle and last 30 trials, the cell became selective to item X at both positions within context B. Cell 9 began firing at higher rates in context A regardless of the stimulus but came to prefer item Y in position 1 of that same

context in the middle and last 30 trials. Conversely, cell 10 gradually developed a preference of item X, across learning at both positions within context A. However, position cells (e.g., cell 11) fired strongly and similarly whenever the animal sampled either stimulus in a position (in this case position 2 of context B), and their specificity did not change during learning. Combining these analyses from all active cells, the percentage of item–position cells (i.e., cells with a significant interaction in the ANOVA) observed across trial blocks increased dramatically during learning sessions, from an average of 6.4% in the first 30 trial block to 31.3% at the end of the learning session ($F_{(2,8)} = 9.39$, $p = 0.008$) (Fig. 5a). In contrast, the percentage of position-only cells (i.e., cells with a significant effect of position but not an interaction in the ANOVA) did not change significantly during the course of learning ($F_{(2,8)} = 0.218$, $p = 0.809$).

As an additional approach to quantifying item and position coding, we calculated both an item and a position SI for each of the early, criterion, and final 30 trial blocks in learning and overtraining sessions for both of the cell types identified in the previously described ANOVAs as item–position and position cells in the final block (see Materials and Methods).

In the first 30 trial block, when the animals were performing at chance level, item–position cells failed to show significant item selectivity when compared with data in which the item identities were shuffled ($z = -0.093$, $p = 0.926$) (Fig. 5c) (see Materials and Methods). However, item–position cells developed significant item selectivity in the middle 30 trials when the performance criterion was attained ($z = 3.882$, $p < 0.001$), and they maintained this strong selectivity through the final 30 trial block at the end of the session ($z = 6.479$, $p < 0.00001$). Conversely, position cells failed to show significant item selectivity throughout the learning session (first 30 trials, $z = -0.078$, $p = 0.932$; middle 30 trials, $z = 0.639$, $p = 0.522$; last 30 trials, $z = 0.845$, $p = 0.397$). Therefore, the two categories of cells behave differently during the course of learning (two-way repeated measures ANOVA interaction, $F_{(2,16)} = 5.89$, $p = 0.012$). *Post hoc* tests revealed that, whereas position cells maintained the same level of stimulus selectivity across all three trial blocks ($F_{(2,8)} = 0.380$, $p = 0.695$), item–position cells adopted a higher level of stimulus specificity in parallel with the increase in performance ($F_{(2,8)} = 7.11$, $p = 0.017$).

In contrast to item selectivity, both types of cells were equally (one-way ANOVA, $F_{(1,8)} = 0.146$, $p = 0.712$) and significantly (position, $z = 6.489$ – 9.892 for all trial blocks, all p values < 0.0001 ; item–position, $z = 6.267$ – 9.617 for all trial blocks, all p values < 0.0001) selective for position. In addition, the average position SIs did not differ across trial blocks during learning (two-way repeated measures ANOVA, $F_{(2,8)} = 0.701$, $p = 0.511$) (Fig. 5e). These findings on item and position selectivity confirm the observations from the ANOVAs, and together the two analyses show that the observations of the appearance of item–position cells associated with learning, contrasted with the stability of Position coding, is robust across methods of analysis.

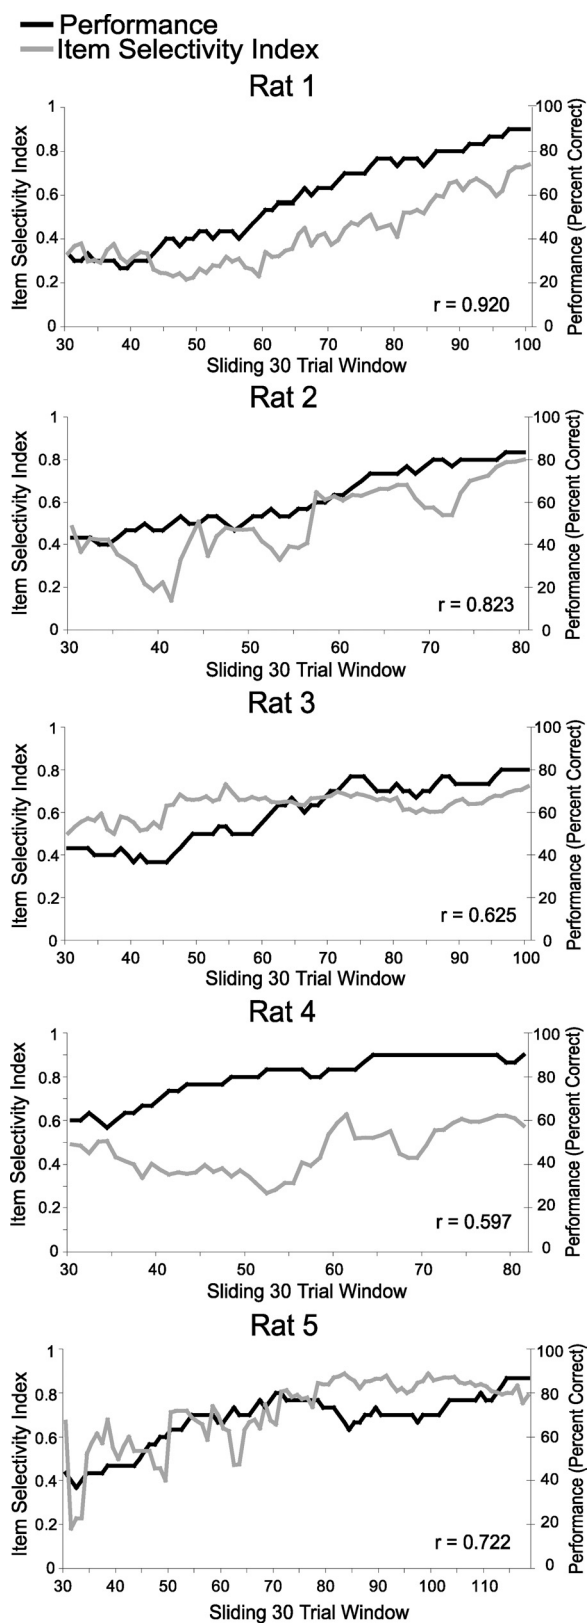


Figure 6. Changes in item selectivity during learning for individual animals. Average item selectivity index and performance accuracy are plotted over sliding 30 trial windows.

The selectivity of item–position cells parallels learning and their activation predicts accuracy

We further examined the degree of stimulus selectivity in item–position cells as accuracy increased during learning. In Figure 6,

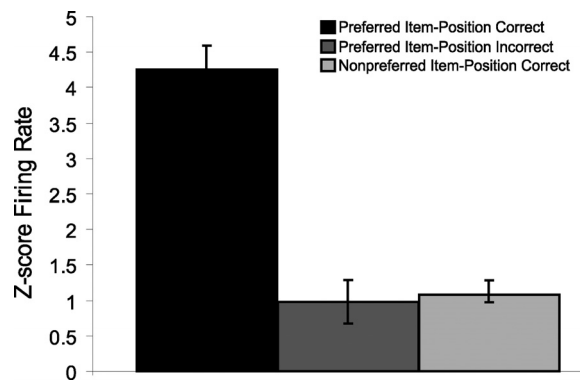


Figure 7. Average \pm SE z-score firing for all item–position cells during item samples for the preferred item–position combination of each cell on correct and incorrect trials compared with averaged z-score firing to nonpreferred item samples at that same position on correct trials.

performance accuracy for each trial was calculated as the percentage correct over a sliding 30 trial window ending with that trial, and the average item SI for all item–position cells simultaneously recorded was also calculated over the same trial window. Each animal showed a steady, gradual improvement in performance accuracy, although the course of learning varied for each animal. The item SI also gradually improved in each animal and correlated strongly (Pearson’s $r = 0.722$ – 0.920) with performance in three of the animals and moderately well (Pearson’s $r = 0.597$ and 0.625) in the other two animals (Fig. 6). Because the scores for each window are not independent, we calculated the significance of the correlations against the distribution of r values calculated from 10,000 shuffles of trial and item SI order. These observed that r values were all significantly higher than the shuffled distribution ($z = 4.22$ – 6.30 , all p values < 0.0001). The degree of correlation between item SI and performance was impressive given the number of item–position cells recorded in each animal was six or less.

We also asked whether the activation of item–position cells predicted accuracy of behavioral responses. To measure activation of the population of item–position cells, we normalized responses as z-score firing rates during the sampling period for each cell and compared average responses to the preferred stimulus at the place associated with maximal responses in correct versus incorrect trials of learning sessions. Item–position cell responses were higher on correct trials than on incorrect trials (Fig. 7) ($t_{(332)} = 5.26$, $p < 0.0001$). In contrast, the responses during incorrect trials to the preferred item–position combination were no different than for nonpreferred items in the same position ($t_{(340)} = 0.46$, $p = 0.644$). Thus the activity of item–position cells strongly predicted subsequent accuracy of the behavioral response.

Both item–position and position cells were stable during the course of overtraining

In contrast with the observations on learning sessions, cells that fired selectively associated with stimulus and position in overtraining sessions typically maintained the same pattern of selectivity across the entire session. Figure 8 uses the same format as in Figure 4 to show different example cells recorded in overtraining sessions. Cells 12–15 each show patterns of responses that are specific to stimulus identity and position across all trial blocks during high, stable performance in overtraining. We also observed position cells that fired consistently when the animal sampled either stimulus at one position in one of the contexts (e.g., cell 16), and these cells also had stable spatial selectivity across the

Overtraining

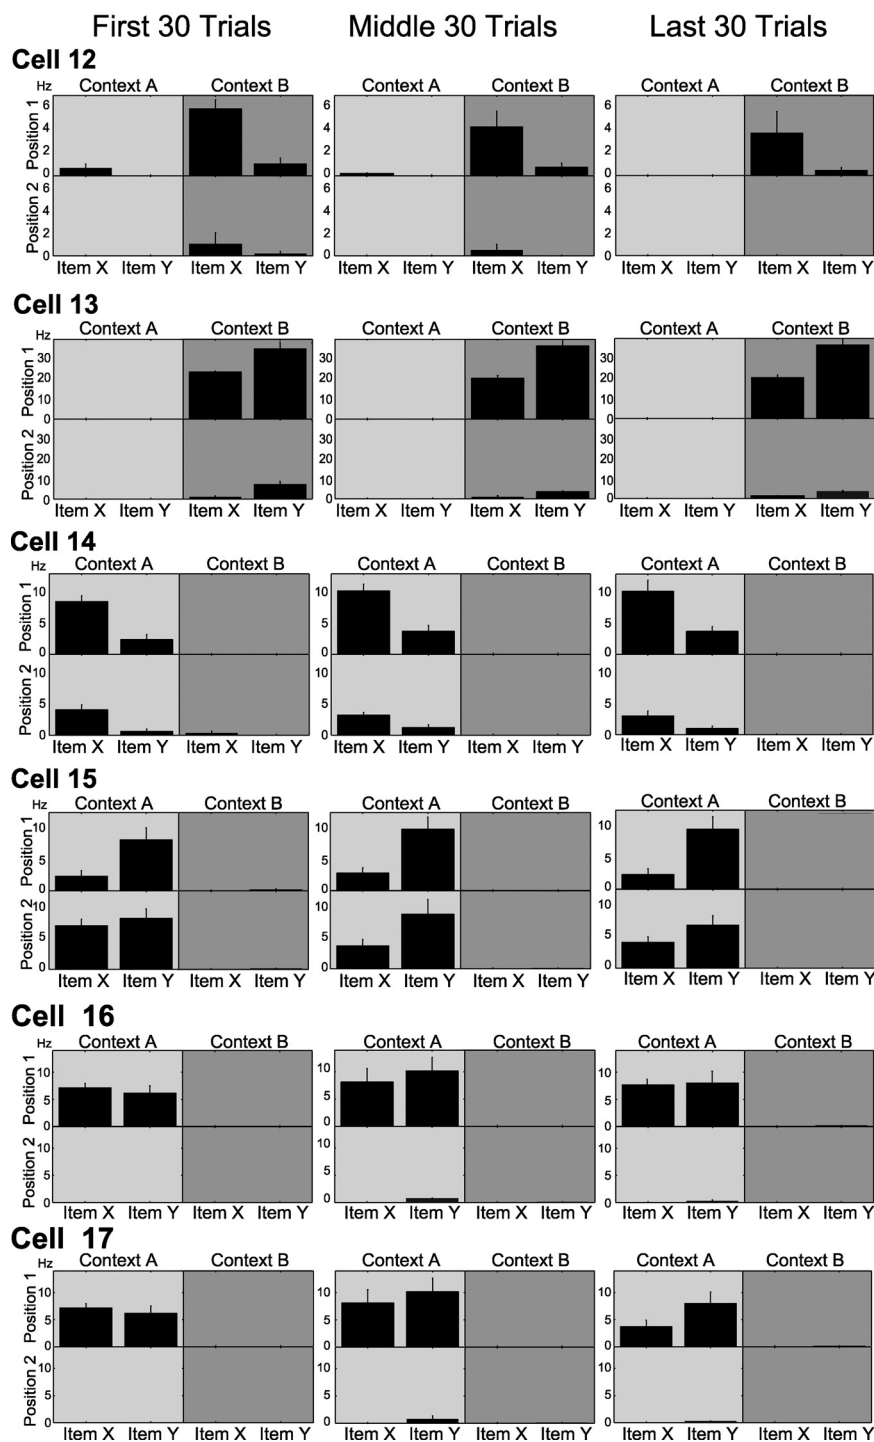


Figure 8. Example cells recorded during an overtraining session. Format is the same as in Figure 4. Cells 12–15 were classified as item–position cells in all trial blocks, whereas cell 16 was classified as a position cell over the same three intervals. Cell 17 shows a significant interaction of stimulus identity with position in the last 30 trials of the session but fails to show this effect in the first 30 trials.

entire session. Not all cells were stable during overtraining. For example, cell 17 showed only spatial selectivity for position 1 in context A in the early and middle blocks and fired differentially during sampling of item Y in that position in the final 30 trials. Combining all of the quantitative analyses, the percentages of both item–position and position cells were stable across trial bins

in overtraining sessions (two-way repeated measures ANOVA; item–position, $F_{(2,8)} = 1.4, p = 0.3$; position, $F_{(2,8)} = 0.570, p = 0.587$) (Fig. 5*b*), such that both remained high throughout overtraining. Consequently, item–position and position cells had divergent patterns of appearance associated with the course of learning (two-way repeated measures ANOVA interaction, $F_{(2,16)} = 4.20, p = 0.034$) but not with overtraining (two-way repeated measures ANOVA interaction, $F_{(2,16)} = 1.85, p = 0.189$). In addition, when compared with shuffled item SI values, item–position cells showed significant item selectivity in all three blocks of trials ($z = 3.405–5.798$ for all trial blocks, all p values < 0.001), whereas position cells failed to show item selectivity ($z = -0.976–1.593$ for all trial blocks, all p values > 0.1). Therefore, item SIs for both cell classes remained stable in overtraining sessions (two-way repeated measures ANOVA; item–position cells, $F_{(2,8)} = 0.817, p = 0.475$; position cells, $F_{(2,8)} = 0.185, p = 0.835$), and item–position cells had a significantly higher item SIs than position cells (one-way ANOVA, $F_{(1,8)} = 44.581, p < 0.001$) (Fig. 5*d*). Both types of cells were equally (one-way ANOVA, $F_{(1,8)} = 0.146, p = 0.712$) and significantly (position, $z = 6.489–9.892$ for all trial blocks, all p values < 0.0001 ; item–position, $z = 6.267–9.617$ for all trial blocks, all p values < 0.0001) selective for position. The average position SIs were also stable during overtraining (two-way repeated measures ANOVA, $F_{(2,8)} = 0.952, p = 0.407$) (Fig. 5*f*).

Development of item–position cell selectivity is not explained by changes in behavior during stimulus sampling, ripple activity, or reward, but information about reward status was incorporated into the item–position representations

Two analyses of stimulus sampling behavior indicated that there were no systematic changes in behavior over the course of learning that accounted for the evolution of item–position coding. In each analysis, we reasoned that item–position encoding could be secondary to a systematic behavioral differences only if the behavioral difference developed over trial blocks and in the particular locations where the item selectivity was observed; such an effect would be detected in a three-way (stimulus \times location \times trial block) interaction in an ANOVA. First, to examine whether item–position coding could be explained by changes in the duration of stimulus sampling, we compared the distance of the animal’s head from the center of each pot at all four stimulus sampling locations, across trial

blocks. If the rat ceased sampling earlier over the course of learning or for a particular stimulus, the average distance should be larger for that stimulus. Average stimulus sampling distances did not differ by stimulus identity or location across trial blocks in the learning sessions for four of the five rats (three-way ANOVA interaction, all p values >0.13). For one rat, a significant difference in stimulus sampling distance between the two stimuli emerged in the final trial block at two of the positions where pots were sampled. However, item–position cells were observed both at the locations where the sampling distances differed and where sampling differences did not differ; in all the other rats, many item–position cells were observed at locations where sampling distances were consistent across locations and trial blocks. Therefore, differences in sampling distance do not explain the observed learning-related appearance of item–position-selective neurons.

Second, we examined head direction during the animal's approach to the stimulus pots, as well as head directions during the stimulus sampling period (see Materials and Methods). None of these analyses distinguished head direction during the approach or stimulus sampling behavior between the items across trial blocks during learning (three-way ANOVA interactions, all p values >0.13). Therefore, head direction during approach or sampling of the stimuli does not explain the observed learning-related appearance of item–position-selective neurons.

In addition, because the incidence of ripples can increase when animals are placed in novel environments (Foster and Wilson, 2006; Cheng and Frank, 2008), we also examined whether the increase in item–position coding was related to differential hippocampal ripple activity during the 1 s stimulus sampling period (see Materials and Methods). However, in no case did our analyses reveal a difference in ripple incidence between items or positions across the trial blocks (three-way ANOVA interaction, all p values >0.1). Therefore, the incidence of ripples does not explain the observed learning-related appearance of item–position-selective neurons.

The appearance and selectivity of item–position cells is also not explained as increased firing at places associated with reward, because equivalent numbers of cells showed conjunctive item–position coding that preferred rewarded and nonrewarded stimulus–position combinations. Thus, 25 cells fired selectively during sampling of item X in context A or item Y in context B, both of which signaled reward and 30 cells fired selectively during sampling of item Y in context A or item X in context B, both of which signaled nonreward. These proportions do not significantly differ ($p = 0.09$, binomial distribution test). In addition, when only rewarded stimulus sampling events are considered, item–position cells still respond more strongly to preferred items in the optimal position than to nonpreferred items in the optimal position (preferred item mean z -score, 5.51 ± 0.51 vs nonpreferred item, 1.67 ± 0.49 ; $t_{(182)} = 3.86$, $p < 0.001$). Nevertheless, item–position responses are greater in cells that prefer rewarded items than cells that prefer nonrewarded items (rewarded preferred item mean z -score, 5.51 ± 0.51 vs nonrewarded preferred item, 1.33 ± 0.50 ; $t_{(230)} = 3.96$, $p < 0.0001$), so a reward signal is incorporated into the representation as increased activation, along with information about the stimuli themselves.

How are item–position cells generated?

Most of the item–position cells (67%) identified at the end of learning evolved out of cells that were identified as position cells earlier in the session. Subsequently, all item–position conjunctive cells retained their specificity for the same context across all three trial blocks. In contrast, 43% of the position cells identified in the

last 30 trials failed to show any significant preference for position earlier in the session and were recruited from previously inactive cells. Thus, because the number of position cells was constant throughout learning, as position cells adopted stimulus specificity, more position cells were added to maintain a relatively constant number of location-selective cells.

Finally, we asked whether the increase in item selectivity by the hippocampal population reflected an overall increment or decrement in the responses to one of the stimuli, or a combination of both increased responses to the preferred stimulus and decreased responses to the nonpreferred stimulus. To address this question for the population of item–position cells, we compared z -score firing rate responses to the preferred and nonpreferred stimulus at the place associated with maximal responses between the first and last 30 trial blocks during the learning session. This analysis revealed that the overall z -score firing rate for item–position cells was higher at the end of learning (one-way ANOVA, $F_{(1,42)} = 5.10$, $p = 0.029$) (Fig. 9*a*). Furthermore, this increased responsiveness can be attributed to a strong and selective increase in the magnitude of responses to the preferred stimulus ($t_{(21)} = -3.35$, $p = 0.003$) without a significant change in response to the nonpreferred stimulus ($t_{(21)} = 1.60$, $p = 0.124$). In contrast, the normalized responses of position cells to stimulus positions showed a qualitatively different pattern (Fig. 9*b*). These cells fired strongly when stimuli were sampled at the most preferred location and very weakly at the least preferred location (one-way ANOVA, $F_{(1,27)} = 114.18$, $p < 0.001$). Furthermore, in contrast to the responses item–position cells, there was no significant change between the first and last 30 trial blocks in the response magnitudes to the preferred position ($t_{(14)} = -0.358$, $p = 0.725$) or to the nonpreferred position ($t_{(13)} = -0.832$, $p = 0.420$).

Discussion

The conditional item–context learning paradigm allowed us to track the firing patterns of single neurons as animals gradually learned which of two items signaled reward in each of two distinctive environments. At the outset of training, we observed strong and prevalent spatial representations of the locations where the items were sampled. In contrast, early in training, almost no neurons fired selectively during the sampling of one of the items, and few cells fired differentially during the sampling of a particular item in one place. However, as learning progressed, the number of cells that demonstrated conjunctive item–place coding grew, and the strength of differential firing also increased, both closely corresponding to the course of learning. By the end of learning, the proportion of cells with conjunctive item–position coding was equivalent to that for location alone. The importance of conjunctive representation to performance is further shown by the absence of increased firing to the preferred stimulus on incorrect trials. Thus, a robust conjunctive activation parallels learning and is critical for the appropriate identification of an item in its context.

Furthermore, item–place conjunctive coding was derived from an initial preference for specific locations at the outset of training and was characterized by increased responses to one of the two items sampled at those locations. By the end of training, more than half (52%) of the position cells that fired during stimulus sampling had converted to item–position conjunctive cells. In contrast, the number of position cells was maintained by recruitment from previously inactive cells, and the magnitude of location selectivity was stable. In overtraining, item–place coding remained prevalent and stable. Importantly, the growth of this conjunctive coding cannot be explained by differences in behav-




Figure 9. Changes in firing rate to preferred and nonpreferred stimuli and positions. **a**, Average z-score firing rate (see Materials and Methods) for item–position cells increases selectively in response to preferred stimuli after learning but remains unchanged in response to nonpreferred stimuli. **b**, Average z-score firing rate for position cells remains unchanged in response to both preferred and nonpreferred positions.

ior or hippocampal ripple activity. In addition, hippocampal neural activity encoded rewarded and nonrewarded item–place combinations equally and showed item–position preferences on equivalently rewarded item samples, indicating that these firing patterns cannot be solely characterized as spatial firing that is enhanced by reward or reduced by nonreward expectancy. These findings provide the first compelling evidence for robust hippocampal representation of items and place conjunctions that are formed during the course of learning about those items and the spatial context in which the events occur.

Many previous studies of hippocampal neuronal activity in behaving rats have identified neurons that fire associated with particular cues and the places where they are experienced (Wible et al., 1986; Muller and Kubie, 1987; Rolls et al., 1989; Wiener et al., 1989; Young et al., 1997; Wiebe and Stäubli, 1999; Wood et al., 1999; Moita et al., 2003; Lenck-Santini et al., 2005). However, in

all of these experiments, conjunctive coding was incidental to the task demands, which involved remembering only the stimuli and their reward contingencies independent of location, and the development of item in place coding was not associated with learning. Rolls et al. (1989) (also Cahusac et al., 1989) reported sparse object–place coding by hippocampal neurons in monkeys performing an object–place recognition task, but the predominant response was a decline in firing rate when an object reappeared in a particular place, and this type of coding was not related to memory performance. Wirth et al. (2003) (see also Cahusac et al., 1993) observed a substantial fraction of hippocampal neurons that increased or decreased the magnitude of their stimulus-evoked responses in parallel with learning specific eye movement responses to the stimuli, and stimulus selectivity increased during the course of learning. This study did not involve learning the place in which stimuli occur as the critical association. Rutishauser et al. (2008) recorded from hippocampal neurons in humans as they remembered specific stimuli and the locations where they had been seen. They reported that hippocampal neurons had greater responses to previously experienced stimuli than novel stimuli and yet greater responses when the subject could remember where the stimulus had been seen. However, the responses were similar to all experienced stimuli, whether they were remembered or not. Furthermore, the responses were greater among all old stimuli and greater yet across all stimuli and locations; so it appears they did not represent specific information about particular items and their locations. Here we trained rats on an item–place association task that requires hippocampal function (Rajji et al., 2006) and observed conjunctive item–place coding that closely paralleled learning about those items and places. Thus, the present findings are consistent with the previous characterizations of hippocampal neuronal activity and demonstrate for the first time the relevance of conjunctive item–place coding in hippocampal-dependent learning about conjunctions of items and places.

The present findings also provide a framework to integrate theories that characterize hippocampal neuronal activity as a representation of the spatial layout of an environment (O’Keefe, 2007) with those that relate it to episodic memory (Eichenbaum et al., 1999). A previous study showed that hippocampal neurons acquire tone-evoked responses within their place fields associated with tone-cued fear conditioning (Moita et al., 2003), and another study reported that almost all hippocampal cells that encoded odors did so in conjunction with their spatial location (Wiebe and Stäubli, 1999). Although these studies did not examine learning of items in places (see these and other examples above), the findings are consistent with the idea that relevant stimuli are encoded within a preexisting spatial framework (O’Keefe and Nadel, 1978). The increment in responses to the salient stimuli in their spatial context observed in the current study is also consistent with the view that hippocampal neurons encode experience-specific information by changes in firing rate within a generally maintained contextual representation (Leutgeb et al., 2005). Here we confirmed that conjunctive coding is derived from preexisting spatial coding in a manner that is directly relevant to and parallels learning about specific items in particular places. Thus, the resolution of the place versus memory controversy may be that the hippocampus encodes events that occur in a particular spatial context, a key feature of episodic memory. However, the specific role of space as the fundamental contextual dimension may need to be expanded. Other recent evidence suggests that “context” may not be strictly restricted to spatial context but rather may also include the temporal context

in which events occur, even in the same place (Manns et al., 2007; Lipton and Eichenbaum, 2008; Pastalkova et al., 2008).

The present findings provide strong and direct support for the hypothesis that the hippocampus integrates “what” and “where” information in the service of episodic memory (Davachi, 2006; Manns and Eichenbaum, 2006; Diana et al., 2007; Eichenbaum et al., 2007). Recent theorizing on the functional organization of the primate medial temporal lobe has suggested that information about the specific items and events is initially processed in the ventral visual pathway, as well as other specific sensory pathways, and arrives in the medial temporal lobe within the perirhinal cortex. At the same time, information about spatial locations and other contextual features of events is processed by the dorsal visual pathway, as well as other multimodal areas, and arrives in the parahippocampal cortex (Eichenbaum et al., 2007). The segregation of “what” and “where” streams are largely preserved through the entorhinal cortex, such that the perirhinal cortex projects predominantly to lateral entorhinal cortex and the parahippocampal cortex projects predominantly to the medial entorhinal cortex, although there are some interactions between these cortical areas.

According to this view, the “what” and “where” information converge within the hippocampus. Functional imaging studies on humans have shown that the hippocampus is activated associated with memory for items and their spatial context (Henke et al., 1997; Davachi et al., 2003; Hannula and Ranganath, 2008). However, based on a meta-analysis of the data from several studies, Wais (2008) concluded that functional imaging cannot determine whether hippocampal processing is selective to conjoint item–context representations or just increased for stronger and more detailed memories (Squire et al., 2007). The current study addresses this issue by identifying qualitative as well as quantitative changes in the responses of hippocampal neurons associated with learning about items in context. Overall responsiveness of hippocampal neurons increased during learning, consistent with reports of greater responses to previously experienced stimuli (Rutishauser et al., 2006, 2008) and to better remembered stimuli (Kirwan et al., 2008; Shrager et al., 2008). However, we also observed that increased neuronal responsiveness associated with learning could be attributed solely to development of robust representation of specific items in the context in which they have differential significance, whereas the number of, the selectivity of, and the magnitude of the responses to the items or contexts alone did not change. Previous recording studies have shown that hippocampal neurons can encode individual items in which the same items are identified across several contexts (Wood et al., 1999; Kreiman et al., 2000; Hampson et al., 2004). The present findings suggest that, in situations in which specific item–context associations must be remembered, stronger hippocampal activity reflects the specific encoding of conjunctive “what–where” representations.

References

- Cahusac PM, Miyashita Y, Rolls ET (1989) Responses of hippocampal formation neurons in the monkey related to delayed spatial response and object-place memory tasks. *Behav Brain Res* 33:229–240.
- Cahusac PM, Rolls ET, Miyashita Y, Niki H (1993) Modification of the responses of hippocampal neurons in the monkey during the learning of a conditional spatial response task. *Hippocampus* 3:29–42.
- Cheng S, Frank LM (2008) New experiences enhance coordinated neural activity in the hippocampus. *Neuron* 57:303–313.
- Davachi L (2006) Item, context and relational episodic encoding in humans. *Curr Opin Neurobiol* 16:693–700.
- Davachi L, Mitchell JP, Wagner AD (2003) Multiple routes to memory: distinct medial temporal lobe processes build item and source memories. *Proc Natl Acad Sci U S A* 100:2157–2162.
- Diana RA, Yonelinas AP, Ranganath C (2007) Imaging recollection and familiarity in the medial temporal lobe: a three-component model. *Trends Cogn Sci* 11:379–386.
- Eichenbaum H (2004) Hippocampus: cognitive processes and neural representations that underlie declarative memory. *Neuron* 44:109–120.
- Eichenbaum H, Dudchenko P, Wood E, Shapiro M, Tanila H (1999) The hippocampus, memory, and place cells: is it spatial memory or a memory space? *Neuron* 23:209–226.
- Eichenbaum H, Yonelinas AP, Ranganath C (2007) The medial temporal lobe and recognition memory. *Annu Rev Neurosci* 30:123–152.
- Foster DJ, Wilson MA (2006) Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature* 440:680–683.
- Hampson RE, Pons TP, Stanford TR, Deadwyler SA (2004) Categorization in the monkey hippocampus: a possible mechanism for encoding information into memory. *Proc Natl Acad Sci U S A* 101:3184–3189.
- Hannula DE, Ranganath C (2008) Medial temporal lobe activity predicts successful relational binding. *J Neurosci* 28:116–124.
- Henke K, Buck A, Weber B, Wieser HG (1997) Human hippocampus establishes associations in memory. *Hippocampus* 7:249–256.
- Kirwan CB, Wixted JT, Squire LR (2008) Activity in the medial temporal lobe predict memory strength, whereas activity in the prefrontal cortex predicts recollection. *J Neurosci* 28:10541–10548.
- Kreiman G, Koch C, Fried I (2000) Category specific visual responses of single neurons in the human medial temporal lobe. *Nat Neurosci* 3:946–953.
- Lenck-Santini PP, Rivard B, Muller RU, Poucet B (2005) Study of CA1 place cell activity and exploratory behavior following spatial and nonspatial changes in the environment. *Hippocampus* 15:356–369.
- Leutgeb S, Leutgeb JK, Barnes CA, Moser EI, McNaughton BL, Moser MB (2005) Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. *Science* 309:619–623.
- Lipton PA, Eichenbaum H (2008) Complementary roles of hippocampus and medial entorhinal cortex in episodic memory. *Neural Plast* 2008:258467.
- Manns JR, Eichenbaum H (2006) Evolution of declarative memory. *Hippocampus* 16:795–808.
- Manns JR, Howard MW, Eichenbaum H (2007) Gradual changes in hippocampal activity support remembering the order of events. *Neuron* 56:530–540.
- Moita MA, Rosis S, Zhou Y, LeDoux JE, Blair HT (2003) Hippocampal place cells acquire location specific responses to the conditioned stimulus during auditory fear conditioning. *Neuron* 37:485–497.
- Moody SL, Wise SP, di Pellegrino G, Zipse D (1998) A model that accounts for activity in primate frontal cortex during a delayed matching-to-sample task. *J Neurosci* 18:399–410.
- Muller RU, Kubie JL (1987) The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *J Neurosci* 7:1951–1968.
- O’Keefe J (2007) Hippocampal neurophysiology in the behaving animal. In: *The hippocampus book* (Andersen P, Morris R, Amaral D, Bliss T, O’Keefe J, ed), pp 475–548. New York: Oxford UP.
- O’Keefe J, Nadel L (1978) *The hippocampus as a cognitive map*. New York: Oxford UP.
- Pastalkova E, Itskov V, Amarasingham A, Buzsáki G (2008) Internally generated cell assembly sequences in the rat hippocampus. *Science* 321:1322–1327.
- Rajji T, Chapman D, Eichenbaum H, Greene R (2006) The role of CA3 hippocampal NMDA receptors in paired associate learning. *J Neurosci* 26:908–915.
- Ranganath C, Yonelinas AP, Cohen MX, Dy CJ, Tom SM, D’Esposito M (2004) Dissociable correlates of recollection and familiarity within the medial temporal lobes. *Neuropsychologia* 42:2–13.
- Rolls ET, Miyashita Y, Cahusac PM, Kesner RP, Niki H, Feigenbaum JD, Bach L (1989) Hippocampal neurons in the monkey with activity related to the place in which a stimulus is shown. *J Neurosci* 9:1835–1845.
- Rutishauser U, Mamelak AN, Schuman EM (2006) Single trial learning of novel stimuli by individual neurons of the human hippocampus-amygdala complex. *Neuron* 49:805–813.
- Rutishauser U, Schuman EM, Mamelak AN (2008) Activity of hippocampal and amygdala neurons during retrieval of declarative memories. *Proc Natl Acad Sci U S A* 105:329–334.

- Shrager Y, Kirwan CB, Squire LR (2008) Activity in both hippocampus and perirhinal cortex predicts the memory strength of subsequently remembered information. *Neuron* 59:547–553.
- Squire LR, Zola-Morgan M, Clark RE (2007) Recognition memory and the medial temporal lobe: a new perspective. *Nat Rev Neurosci* 8:872–883.
- Wais PE (2008) fMRI signals associated with memory strength in the medial temporal lobes: a meta-analysis. *Neuropsychologia* 46:3185–3196.
- Wible CG, Findling RL, Shapiro M, Lang EJ, Crane S, Olton DS (1986) Mnemonic correlates of unit activity in the hippocampus. *Brain Res* 399:97–110.
- Wiebe SP, Stäubli UV (1999) Dynamic filtering of recognition memory codes in the hippocampus. *J Neurosci* 19:10562–10574.
- Wiener SI, Paul CA, Eichenbaum H (1989) Spatial and behavioral correlates of hippocampal neuronal activity. *J Neurosci* 9:2737–2763.
- Wirth S, Yanike M, Frank LM, Smith AC, Brown EN, Suzuki WA (2003) Single neurons in the monkey hippocampus and learning of new associations. *Science* 300:1578–1581.
- Wood ER, Dudchenko PA, Eichenbaum H (1999) The global record of memory in hippocampal neuronal activity. *Nature* 397:613–616.
- Yonelinas AP, Otten LJ, Shaw KN, Rugg MD (2005) Separating the brain regions involved in recollection and familiarity in recognition memory. *J Neurosci* 25:3002–3008.
- Young BJ, Otto T, Fox GD, Eichenbaum H (1997) Memory representation within the parahippocampal region. *J Neurosci* 17:5183–5195.