

Placing Memories in Context: Hippocampal Representations Promote Retrieval of Appropriate Memories

David A. Bulkin,^{1*} L. Matthew Law,^{2,3} and David M. Smith⁴

ABSTRACT: Returning to a familiar context triggers retrieval of relevant memories, making memories from other contexts less likely to intrude and cause interference. We investigated the physiology that underlies the use of context to prevent interference by recording hippocampal neurons while rats learned two conflicting sets of discrimination problems, either in the same context or in two distinct contexts. Rats that learned the conflicting problem sets in the same context maintained similar neural representations, and performed poorly because conflicting memories interfered with new learning. In contrast, rats that learned in different contexts formed distinct ensemble representations and performed significantly better. We also measured trial-to-trial variation in representations and found that hippocampal activity was directly linked with performance: on trials where an old representation was active, rats were far more likely to make errors. These results show that the formation of distinct hippocampal representations is critical for contextually appropriate memory retrieval. © 2016 Wiley Periodicals, Inc.

KEY WORDS: context; interference; memory

INTRODUCTION

Context is a fundamental organizing feature of memory. Information is best remembered when subjects are tested where they learned (Godden and Baddeley, 1975; Balsam and Tomie, 1985; Smith, 1988), a phenomenon so integral to memory that merely providing the instruction to think about the learning context is sufficient to improve recall (Smith, 1979). Context has a particularly important role in mitigating mnemonic interference (Bouton, 1993; Smith, 1988; Smith and Vela, 2001). Subjects that learn two lists of word associations or nonsense syllables experience less interference (i.e., show better recall) when the items are learned in distinct environments than when both lists are learned in the same environment (Bilodeau and Schlosberg, 1951; Greenspoon and Ranyard, 1957; Dallett and Wilcox, 1968; Watts and Royer, 1969).

The neural basis of the link between context and memory almost certainly involves the hippocampus (for review see: Nadel, 2008; Smith, 2008). Context specificity of various kinds of memories is disrupted following lesions of the hippocampus (Kim and Fanselow, 1992; Phillips and LeDoux, 1992; Eacott and Norman, 2004; Smith and Mizumori, 2006a), and an intact hippocampus is critical for the use of context to overcome interference (Butterly et al., 2012). The underlying mechanism seems to involve hippocampal place cells, which respond based on the spatial position that an animal occupies (O'Keefe and Dostrovsky, 1971). These neurons alter their firing properties following a change in context (or “remap”; for review, see: Colgin et al., 2008), and, thus, provide a population representation that is unique to each context.

Two important developments in the study of hippocampal neurophysiology raise questions about the mechanics of hippocampal context representations. First, is the uncovering of a broad spectrum of nonspatial features that govern hippocampal activity. Indeed, O'Keefe and Dostrovsky's initial report of place cells (1971) noted many neurons that responded to a variety of (non-spatial) behavioral factors. Since then, systematic study has indicated a remarkably rich array of tuning properties of hippocampal neurons (e.g., Eichenbaum et al., 1987; Wood et al., 1999; Pastalkova et al., 2008; Muzzio et al., 2009). Second, recent studies have demonstrated that neural activity is exquisitely patterned across ensembles of hippocampal neurons (e.g., Dupret et al., 2013; Jezek et al., 2011; Kelemen and Fenton, 2010) and the clear implication is that critical information is contained in a population code. Ensemble spatial activity seems to remap all at once following a contextual manipulation, providing a statistically independent representation that is ideal for associating with context-specific memories (Smith and Bulkin, 2014).

Does non-spatial activity show a similar type of remapping? And how does the ensemble code support learning and memory? In the present study we trained rats on a task designed to induce proactive interference. Rats learned two sets of conflicting odor discrimination problems, either in the same context or in two distinct contexts. In this task, just like in tasks in which human subjects learn lists of words, rats trained on the new odor problems in a new context are better equipped to resolve interference. Our previous work has shown that rats learning these two sets in two distinct contexts suffer less interference, and so make fewer errors, than rats

¹ Department of Neurobiology and Behavior, Cornell University, Ithaca, New York; ² Barrow Neurological Institute at Phoenix Children's Hospital, Phoenix, Arizona; ³ Department of Child Health, University of Arizona College of Medicine, Phoenix, Arizona; ⁴ Department of Psychology, Cornell University, Ithaca, New York
Grant sponsor: NIH; Grant number: R01 MH083809.

*Correspondence to: David A. Bulkin; Department of Neurobiology and Behavior, Cornell University, Ithaca NY 14853.

E-mail: bulkin@cornell.edu

Accepted for publication 27 February 2016.

DOI 10.1002/hipo.22579

Published online 00 Month 2016 in Wiley Online Library (wileyonlinelibrary.com).

learning both sets in the same context (Butterly et al., 2012; Law and Smith, 2012).

There is an extensive literature on hippocampal remapping. Large scale changes in the environmental context cause a complete orthogonalization of activity patterns (i.e., global remapping; Leutgeb et al., 2005), resulting in unique hippocampal firing patterns for each context the subject encounters (Alme et al., 2014). However, although remapping is ubiquitous and has been exhaustively described (e.g., Muller and Kubie, 1987; Lever et al., 2002; Wills et al., 2005; Leutgeb et al., 2007; Kelemen and Fenton, 2010; Jezek et al., 2011; for review see: Colgin et al., 2008), most studies have not directly examined whether remapping is consequential for memory retrieval. Previous studies that have examined remapping and spatial behavior have not found a clear relationship (Cooper and Mizumori, 2001; Jeffery et al., 2003). We have proposed that ensemble representations of context could become associated with appropriate memories and behaviors such that returning to a familiar context triggers retrieval of relevant memories (Smith and Bulkkin, 2014). Consistent with this hypothesis, recent studies have shown that optogenetic reactivation of a hippocampal representation is sufficient to evoke contextual fear responses (Liu et al., 2012; Ramirez et al., 2013). In the present study, we examined the relationship between endogenous hippocampal representations and context appropriate behavior.

We quantified the difference in hippocampal activity patterns as rats learned the two conflicting problem sets in the same context, and compared it with the difference when the problem sets were learned in different contexts. If distinct representations are important for retrieving context appropriate memories and behaviors, then remapping should be associated with performance. Individual neurons showed a variety of response types, including position and event locked activity and mixtures thereof. Because of the array of responses, the present study afforded a unique opportunity to examine context-triggered remapping of spatial and event activity at the population level. As such, rather than applying a classification strategy, we subjected all neurons to a series of spatial and temporal analyses that were agnostic to response type.

Rats that learned new odor-reward associations in a new context had neural firing that was completely different in the two contexts, while those that learned new associations in the old context showed persistent activity patterns. This remapping yielded an ensemble representation that was statistically independent in the new context, a result that was apparent in both position- and event-locked analyses. Moreover, the trial-to-trial variance in representational similarity was predictive of behavioral performance: rats were least likely to make errors when hippocampal activity was most distinct. Taken together, these results show how hippocampal ensembles work to provide a new and statistically independent code, across a heterogeneity of constituent neuronal response types. This code, a context representation, can prime context-appropriate memories and prevent interference from memories that would be appropriate in other contexts.

MATERIALS AND METHODS

Training Procedures and Apparatus

The training procedure followed our previously reported methods used to probe the use of context to prevent interference in rats (Butterly et al., 2012; Law and Smith, 2012; Peters et al., 2013). Seven adult male Long–Evans rats were trained on a task designed to induce proactive interference. Rats learned two sets of eight odor discrimination problems, either in the same environmental context or in two different contexts. Rats learned the first set until they reached asymptotic performance (>90% correct on two consecutive sessions) and then began learning the second set. Three of the rats learned the second set in a different context and four learned the second set in the same context. All procedures complied with the guidelines established by the Cornell University Animal Care and Use Committee.

Training took place in wooden chambers with a 60 cm by 45 cm floor and a removable divider (Fig. 1B). One side of the chamber served as an intertrial waiting area, the other contained two plastic cup holders for the presentation of stimuli. Contexts were differentiated by a number of multimodal cues: the enclosure's color and substrate (light wood or a black rubber lining), the surrounding walls (black painted walls or white blinds), the frequency content of a 65 dB continuous background masking noise (white noise or pink noise), and the ambient odor left by wiping out the chamber with baby wipes prior to each training session (unscented or scented, Rite Aid, Inc.). Additionally, the rats were transported in covered cages to the experimental area by different methods (via a cart or carried by hand).

Trials began when the experimenter lifted the divider and the rat entered the odor presentation area. The rat was allowed to approach the cups and search for the reward. Subsequently, the rat was returned to the intertrial waiting area and the divider was replaced while the experimenter prepared the cups for the next trial. Trials were marked as errors if the rat dug in the unbaited cup, any displacement of bedding was considered a digging response. Training sessions continued until the rat reached a behavioral criterion of 90% correct choices, and a minimum of four training sessions had elapsed. All rats achieved criterion performance by the fifth day of training. Sessions were generally conducted daily; however, 1–3 days without training were interleaved in order to maximize the opportunity to obtain large and stable neuronal populations by the time of the final training session on the problem set.

Once the rat reached the performance criterion, training began on a second problem set (Fig. 1A; Supporting Information Fig. 1A). This set contained eight odors from the first set, with their predictive values reversed, paired with eight novel odors. To prevent rats from adopting a strategy based on odor novelty, half of the new odors were rewarded and the other half were unrewarded. Training on the second set continued for five sessions regardless of performance. After the last set rats were given a pellet detection test in which two

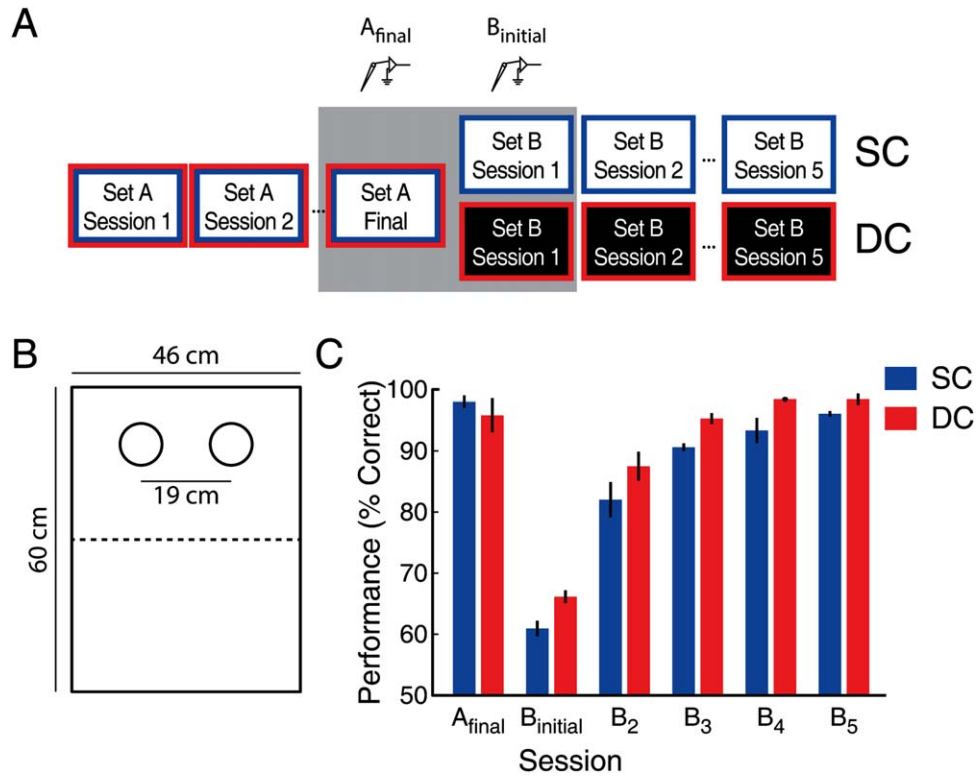


FIGURE 1. Proactive interference task. (A) Rats were trained on a proactive interference task in which they learned two sets of odor discrimination problems. They first learned to discriminate within eight pairs of odors over the course of four to five sessions, until they reached a criterion performance level of 90% correct choices. Next they learned a new set of problems either in the same environmental context (SC; blue) or in a different environmental context (DC; red). Recordings were taken throughout training, but analyses focused on a set of neurons that were stably recorded as the rats switched from

one set to the next (A_{final} and $B_{initial}$; shaded region in A). (B) A diagram of the training apparatus. An inter-trial waiting area (bottom) was separated from the discrimination training area (top) using a removable divider (horizontal broken line). Rats that learned the new set in the new context performed better than rats that learned in the same context. (C) average performance on the final session of Set A and on each session of Set B. Error bars indicate SEM across rats. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

cups containing identical odors were presented, but only one was baited (32 trials). If the rats could detect the sucrose pellets, they would be expected to perform above chance. As with our previous studies (Butterly et al., 2012; Law and Smith, 2012; Peters et al., 2013), none of the rats performed above chance.

A third, neutral context (a 1-m square PVC box, with a unique configuration of context cues in which the rats foraged for chocolate sprinkles), was used to probe for activity as electrodes were lowered into hippocampus before training began. Recordings were also taken in this context after each session and were used to aid spike sorting, specifically to identify cells that were silent during the main experiment on one of the two critical sessions.

Surgery and Electrophysiological Procedures

Moveable electrode arrays containing 16 insulated platinum iridium tetrodes (composed of four 17 μm wires; California Fine Wire, Grover Beach, CA) were implanted bilaterally just above the dorsal hippocampus (3.5 mm posterior and 2.5 mm

lateral to bregma). Following recovery from surgery, tetrodes were slowly lowered into the CA1 cell layer and rats began training on the first set of odors. Because the focus of the study was on changes in activity across two critical sessions (A_{final} and $B_{initial}$; Fig. 1A), tetrodes were advanced over initial training on the first odor set but then left in place at least 24 hours before the final recording session on the first odor set (A_{final}) in order to maintain recordings as the rats began the new set ($B_{initial}$). Multiunit recordings were sorted into constituent units using standard clustering techniques and were matched across sessions by applying cluster boundaries from one session to the other with manual adjustments to account for drift (Mankin et al., 2012).

Data Analysis

All analyses were performed using custom software written in the numerical computing environment Matlab (Mathworks, Natick, MA). Analyses were restricted to cells with an average firing rate of less than 3 spikes/second across the session (i.e., discarding any putative interneurons). Separate analyses with

the inclusion of these neurons, or with higher thresholds for firing rate did not qualitatively change any of the differences between SC and DC conditions.

Spatial firing maps were constructed by calculating firing rates in 4.5-cm square bins spanning the apparatus. The data were smoothed by convolution with a 4 bin Gaussian kernel with unity sum. Spatial bins that contained less than one second of occupancy (following smoothing) were discarded. For display purposes only, firing rate maps were interpolated linearly with three points between each sampled data value.

Peri-event time histograms (PETHs) were calculated by binning and averaging the data in 100 ms bins centered on three events during the trial: the start and end of the trial (defined by the moment the rat crossed an imaginary line at the position of the removable divider) and odor sampling (defined by the moment the rat's nose was directly above the odor cup, scored manually using video recordings of the experiment). Odor sampling times were identified by manual flagging of raw video data sampled at 30 fps. Trials lasting longer than 30 seconds were discarded from all ensemble and event related analyses, however performance was still included for these trials (these trials always contained errors, and the errors were always made within the first 30 seconds of the trial). For display only, the data shown in Figures 3C and D were smoothed with a 1 second Gaussian window, and normalized to z units.

Pearson correlations were computed between the spatial firing maps across sessions, using only those bins that were visited in both contexts. For units that displayed spatial sensitivity in both sessions we computed the center of mass (COM) of spatial firing rate as the weighted average of the binned position of the rat, with weights determined by binned firing rate. For a unit to be considered spatially sensitive it had to have a spatial firing map with a contiguous region (i.e., place field) with at least twice the firing rate within the field as outside of the field, a field area less than 30% of the apparatus, and at least 100 total spikes within the field. To estimate chance levels for shifts in COM, we computed the average pairwise distance between COM for different neurons.

Analogous correlation and COM analyses were applied to event-related firing. PETH correlations were computed between sessions for the trial start and trial end aligned PETHs, separately for the two events. The event COM analysis was performed by fitting a Gaussian curve to each PETH using a nonlinear least-squares approach (Matlab Curve Fitting Toolbox):

$$f(t) = Ae^{-\left(\frac{t-\mu}{c}\right)^2} + d$$

The function $f(t)$ describes the firing rate as a function of time relative to the event (t). In order to ensure identification of a (potential) peak near the event time, the parameter A , which indicates the amplitude of the Gaussian, was restricted to positive values, and the parameter μ (which marks the peak time of the Gaussian) was restricted to the range of the histogram (i.e., ± 2 seconds). Initial values for these coefficients were taken from the amplitude and time of maximum of the average binned firing rates. The parameters c and d , which

specify the width and offset of the curve, were unrestricted. The difference in the parameter μ across sessions was taken as the shift in event COM, the 95% confidence interval on this parameter is plotted in Supporting Information Figure 5B.

To calculate normalized information scores, an information score was assessed as described in (Markus et al., 1994):

$$I = \sum P_i \frac{R_i}{R} \log_2 \frac{R_i}{R}$$

For spatial information, P_i was the probability for occupancy of bin i , R_i was the mean firing rate for bin i , and R was the overall mean firing rate. For event related information content, P_i was fixed at the reciprocal of the number of trials (each bin of the PETH was "visited" the same number of times), values for R_i and R were calculated in the same way as for spatial data except that they were calculated using temporal bins (± 2 seconds around each event, 100 ms bins) rather than spatial bins. In both the spatial and temporal analyses, normalized values were calculated using an iterative process (Markus et al., 1994): a distribution of 500 pseudo information values was calculated by randomly offsetting spike times using uniform random values ranging from 5 to 100 seconds. The values used for comparison between sessions were z transformed using this distribution (i.e., the number of standard deviations from the mean of the distribution of pseudo information values).

Classification was performed using linear discriminant analysis with uniform priors. For epoch classification each trial was divided into four time windows: a 1-second period preceding the trial, the period between the start of the trial and the arrival at the first cup, the following 400 ms, and a period extending until 400 ms after the arrival at the final cup (Supporting Information Fig. 6). Firing rates were divided into 100 ms bins, and each bin was labeled with the epoch. An analogous approach was applied to spatial classification, dividing the apparatus into four quadrants (to match the number of epochs in the event related classification). Data were divided into 1-second bins, and each bin was labeled with the quadrant identity. Bins in which the rat occupied multiple quadrants were discarded from analysis.

Accuracy of the event locked classifiers was tested by creating 10,000 subsets of the data, each time selecting half of the Set A trials to train the classifier, and the remaining trials as well as half of the Set B trials to test the classifier. Using subsets allowed us to prevent spuriously high success rates due to overfitting and to estimate confidence intervals on performance. Taking randomly selected trials to form subsets rather than individual bins allowed us to avoid spuriously high success rates due to the high dependency of firing rate from one bin to the next. To select training and testing subsets for the spatial classifiers, each session was divided into 100 evenly distributed blocks and, just like with the event locked classifiers, half of the blocks were used to train the classifier and the remaining half (as well as half of the Set B blocks) were used to measure performance. The performance of the classifier was considered above chance if 95% of the distribution of classifier hit rates

were above a chance value (taken as the reciprocal of the number of classes).

Similarity between representations was calculated using a distance index (MDI):

$$\text{MDI}(\vec{x}) = \frac{(\vec{x} - \vec{\mu})S^{-1}(\vec{x} - \vec{\mu})}{n}$$

For each trial, the average firing rate of each unit was tabulated to form a population vector for that trial. The covariance across A_{final} population vectors was calculated. Next, the distance to the mean A_{final} population vector ($\vec{\mu}$) was calculated for each B_{initial} trial (\vec{x}), in units of A_{final} covariance (S). This Mahalanobis distance (i.e., squared distance normalized by covariance) was then additionally normalized by the number of units (n) in the sample (for comparison across rats with different sample sizes).

Histology

Following the experiment the rats were deeply anesthetized and perfused transcardially with 4% paraformaldehyde. Their brains were removed, sectioned at 40 μm , mounted on slides, and stained with cresyl violet. The position of the electrodes was confirmed to be in the cell layer of CA1 in dorsal hippocampus.

RESULTS

We trained seven rats on a task designed to induce proactive interference. Rats sequentially learned two conflicting sets of odor discrimination problems. On each trial, a removable divider was lifted, and rats encountered a pair of cups filled with odorized digging medium. Rats quickly learned that one odor in each pair was associated with a reward, and dug in the appropriate cup for a buried sucrose pellet, rejecting the unbaited cup. Once the rats reached asymptotic performance on the first set of odor problems, training commenced on a new set, which contained some of the odor cues from the first set with their predictive value reversed (Fig. 1A; Supporting Information Fig. 1A). Four rats learned the new set in the same environment as the first (same context, SC), and three rats learned in a new environment (different context, DC). Rats in the SC group showed a higher error rate on the first session of training on the second set of odor problems (Fig. 1C; $t_{(5)} = 5.05$, $P < 0.01$), indicating that they were more susceptible to interference. Thus, learning the two problem sets in different contexts provided a significant advantage. Although the magnitude of the performance enhancement seen in the DC rats was modest, it persisted through several training sessions and was remarkably consistent. In the present study, every DC rat performed better than every SC rat, and similarly enhanced performance has been reported in four previous stud-

ies using this task (Butterly et al., 2012; Law and Smith, 2012; Luu et al., 2012; Peters et al., 2013).

To investigate the nature of hippocampal representations and their utility in interference prevention, we recorded the activity of 125 dorsal CA1 neurons (99 putative pyramidal units included for analysis, see methods) as rats switched from the well learned initial problem set to the new, conflicting problem set. We then compared neuronal ensemble activity in the SC and DC conditions. Specifically, we focused on the final session of the first problem set and the first learning session of the second problem set (A_{final} and B_{initial} , respectively; the shaded region in Fig. 1A). Using a variety of analyses, we found that ensembles of neurons recorded from DC rats formed a statistically independent representation on B_{initial} while neurons in SC rats continued to fire as they did in A_{final} .

Individual neurons showed a variety of response profiles. While many units showed clear spatial sensitivity, others showed discharge patterns that were tightly locked to important temporal events that occurred during the trial, and many were driven by a combination of spatial and event-related factors. Event responsive neurons in the hippocampus have been described previously (Wood et al., 2000), and event responses seem to undergo remapping comparable to spatial responses (Smith and Mizumori, 2006b). However, many experiments do not have salient events for neurons to respond to (e.g., open field studies), while in others (e.g., maze studies) the design precludes disentanglement of spatial and event-related responses because important task events always occur at the same location. This can make a response appear locked to space when it is actually governed by temporal factors (Kraus et al., 2013). Because both spatial and temporal factors influence activity, and the context manipulation led to changes in both domains, we conducted parallel analyses wherein we binned the data based on spatial location and two important task events, the start and end of the trials. We first treated all of the neurons as spatial, and quantified the extent of change across the critical two sessions. Then, we treated all of the neurons as event-responsive, and applied an analogous set of analyses. This approach allowed us to consider the remapping of neuronal activity while remaining agnostic to the factors that shaped the responses of each individual neuron.

Remapping in a Spatial Framework

When the new learning took place in the same context, spatial response patterns persisted (e.g., Fig. 2A; Supporting Information Fig. 2A). Neurons in DC rats, however, showed a complete reorganization of activity (Fig. 2B; Supporting Information Fig. 2B). The neuron depicted in Figure 2B showed clear spatial tuning but its peak activity shifted to a different location in B_{initial} . During A_{final} , this neuron fired preferentially on the left side of the apparatus, just behind the divider (indicated by the gray broken line). When the rat entered a new context, in B_{initial} firing was mainly on the right hand side of the apparatus, far from the divider. Other neurons showed a binary response pattern, with vigorous spatial firing in one




FIGURE 2. Remapping in a spatial framework. A and B show waveforms (top; the four colors indicate the waveform recorded on the four wires of the tetrode) and firing rate as a function of space (bottom) for example units on the final session of Set A (A_{final} ; left) and the initial session of Set B ($B_{initial}$; right). The unit depicted in A was recorded from a rat in the SC group, and showed elevated activity in the same location in both sessions, whereas the unit depicted in B was recorded from a rat in the DC group, and fired in different locations in the two sessions. The gray overlay on the spatial firing map in A indicates a diagram of the apparatus. (C) Spatial firing rates were well correlated in SC rats but not DC rats. (D) Remapping was also evidenced by a shift in center of mass of spatial firing (COM) across the two sessions in DC rats where

the distance between COM was not different from a chance COM shift, which was estimated by calculating the average COM distance between neurons (*broken line*). Note that while low values in C and E indicate dissimilar activity patterns, low distances between COMs indicate similar activity (the COM was in the same location). E compares the correlation between normalized spatial information across the two sessions for SC and DC rats. The extent to which a neuron exhibited spatial tuning was related across the two sessions in SC rats: neurons that showed high spatial information in one session were likely to show high spatial information in the other. In DC rats, there was no relationship between a given neuron's spatial information across sessions. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

session, but virtually no activity in the other (e.g., see Supporting Information Fig. 2B). In contrast to the remapping seen in DC rats, spatial tuning in SC rats was unchanged across the sessions. The spatial firing maps in Figure 2A and Supporting Information Figure 2A show tuning patterns that were stable, these neurons fired in the same locations during both sessions. Interestingly, while neurons recorded from SC rats did not change their preferred firing locations, many cells showed large differences in firing rate (i.e., rate remapping). This may have occurred in response to the new problem set or merely the passage of time (see Mankin et al., 2012).

We quantified the extent of change between A_{final} and $B_{initial}$ in neurons recorded from SC and DC rats. First we compared

the average firing rates in each two-dimensional spatial bin between the sessions. In the SC condition, the spatial response patterns of neurons were highly correlated across sessions (Fig. 2C). In the DC condition however, there was no relationship between the spatially binned activity across the two sessions (SC compared with DC: $t_{(90)} = 12.95$, $P < 0.01 \times 10^{-19}$; SC r compared with 0: $t_{(45)} = 20.62$, $P < 0.01 \times 10^{-21}$; DC r compared with 0: $t_{(45)} = 1.37$, $P = 0.18$).

Low spatial correlation values in DC-rats could arise from either a change in place field location or the place field could disappear altogether with the change in environment. To investigate remapping that was driven by a change in the preferred firing location, we compared the center of mass of spatial firing

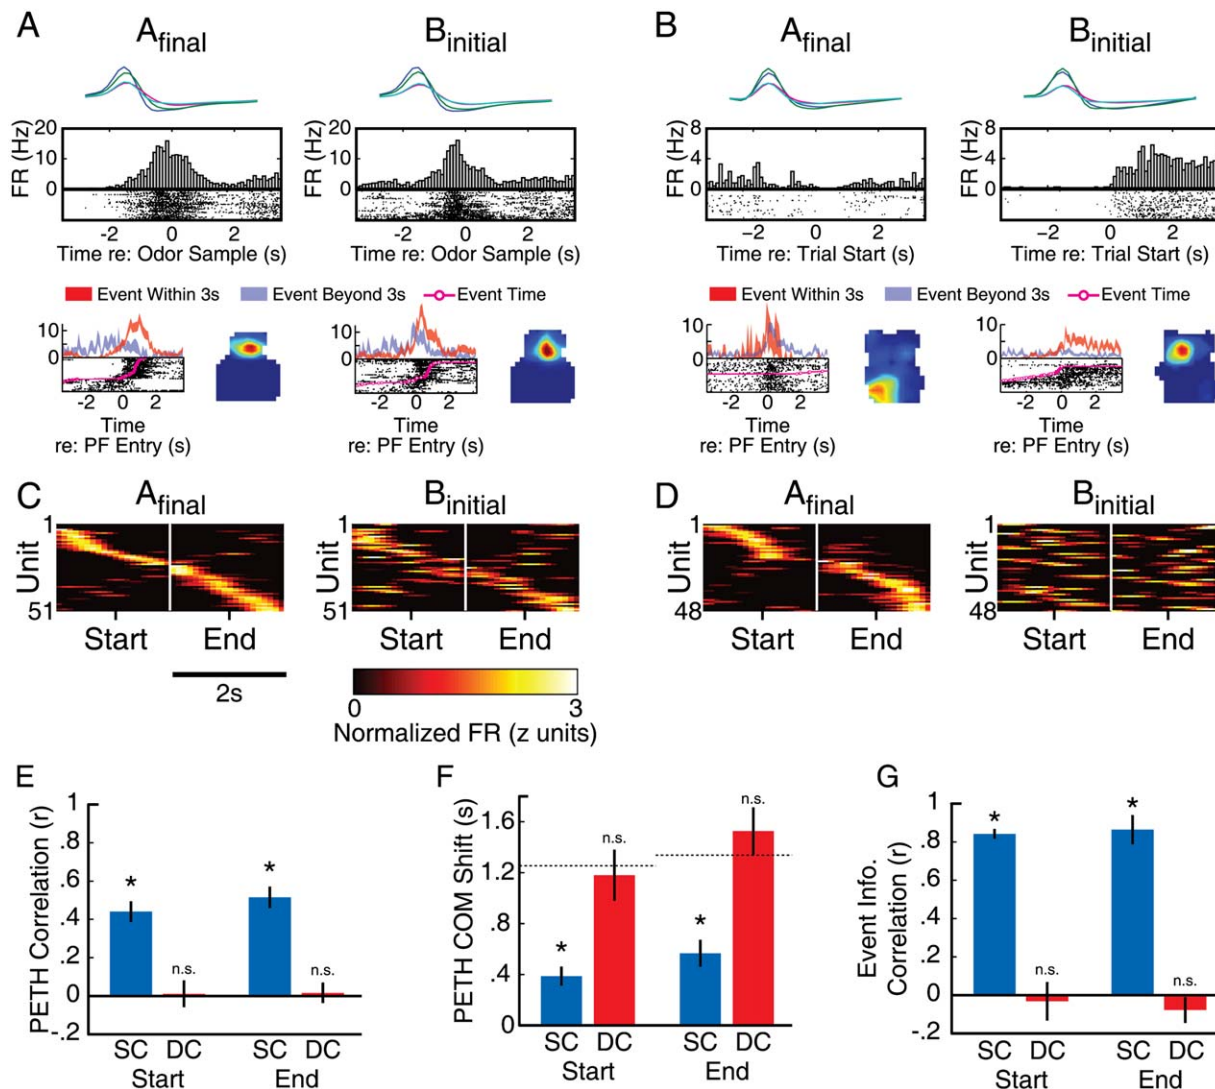


FIGURE 3. Remapping in an event locked framework. A and B show waveforms and PETHs/rasters for example units on the final session of Set A and the initial session of Set B. The cell activity depicted in A, recorded from an SC rat, showed a vigorous response as the rat sampled odors, a response that was present on both sessions. This neuron showed a response regardless of the odor, and samples from all 16 odors are included in the raster. The cell activity depicted in B, recorded from a DC rat, showed a response at the start of trials, but only on the $B_{initial}$ session. Below A and B are PETHs/rasters and heatmaps of the same data plotted to illustrate the differential contributions of spatial location and events. In the rasters, the data are aligned on the rats entry into the apparent place fields at time zero, and sorted from top to bottom by the relative time of the event in question (odor sampling in A and trial start in B), which is indicated by the magenta line. The accompanying histograms indicate firing rate when the event was within 3 seconds of place field entry (red) or more than 3 seconds from place field entry (blue, shading indicates SEM). Greater firing for passes through the apparent place field that coincided with the event suggests that firing is modu-

lated by the occurrence of the event and not solely by spatial location. In contrast, firing that consistently occurs on entry to the apparent place field regardless of the temporal proximity of the event suggests that firing is modulated primarily by spatial location (e.g., red and blue lines are similar in the lower left histogram in plot B). C and D show ensemble activity from all SC (C) and DC (D) units, each row shows the normalized PETH of an individual neuron in a ± 2 s period surrounding the trial start (left) and trial end (right). The rows are sorted based on the time of peak firing relative to the trial start and trial end during A_{final} and the same order is maintained for the data of $B_{initial}$. E-G show remapping of event-locked firing using analogous analyses to those presented in Figures 2C-E. PETHs were well correlated in SC, but not DC rats (E), and showed a shift in the PETH COM in DC, but not SC rats (F). Note that (as in Fig. 2) low values in F indicate a response that was consistent across the two sessions. The extent to which a neuron exhibited event tuning was related across the two sessions in SC rats, while in DC rats there was no such relationship G. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

rates (COM; Leutgeb et al., 2005; Mehta et al., 1997) of neurons across sessions. To be considered for this analysis, cells had to show spatially restricted activity in both sessions (i.e., a

contiguous region with at least twice the firing rate within the field as outside of the field, with an area <30% of the apparatus, and at least 100 total spikes within the field; 39 units in

SC rats, 24 units in DC rats). Importantly, we did not exclude apparent place fields that may have arisen due to event related firing which happened to coincide with a particular location (e.g., odor driven responses occurring at the cup location, see below). The distance in space between COMs from A_{final} to B_{initial} was larger in DC rats than SC rats (Fig. 2D; $t_{(61)} = 6.70$, $P < 0.01 \times 10^{-6}$). Furthermore, the COM shifts in rats in the DC condition were as large as the average distance between the place fields of two different neurons (one sample t -test; DC $t_{(24)} = 0.92$, $P = 0.37$; SC $t_{(38)} = 17.51$, $P < 0.01 \times 10^{-17}$). In short, in DC rats, the preferred firing location of a neuron in one context was completely unrelated to the preferred location in the other context.

To quantify the extent to which neurons changed their degree of spatial sensitivity we computed a spatial information score for each unit/session and compared these values for individual units across sessions. This approach asks whether neurons that had responses which were tightly linked with space in A_{final} continued to have strong spatial sensitivity in B_{initial} . Spatial information was calculated based on the spatially binned firing rates, and then z -transformed using the methods described in (Markus et al., 1994). Neurons recorded from SC rats showed similar spatial information scores across sessions (Fig. 2E; r compared with 0: $t_{(3)} = 13.42$, $P < 0.001$): those units that displayed highly spatially sensitive responses during A_{final} continued to display highly sensitive responses in B_{initial} , and neurons that were uninformative about space remained as such. The DC rats did not show a relationship between spatial information scores of neurons across contexts (Fig. 2E; r compared with 0: $t_{(2)} = 2.27$, $P = 0.11$). In other words, neurons that were spatially tuned in one context were equally likely to be tuned or untuned in the other.

Remapping in an Event Locked Framework

Many neurons showed responses linked to temporal events that occurred during the trial, and these responses showed a similar pattern of changes as the spatial responses described above. The neuron depicted in Figure 3A, recorded from an SC rat, showed stable responses time-locked to odor sampling in both sessions. In contrast, the activity pattern shown in Figure 3B, recorded from a DC rat, indicated event-locked firing in only one of the two sessions (B_{initial}). Distinguishing these responses as event-driven, rather than position-driven, was challenging because the rat visited a restricted set of locations at the times when events occurred. For instance, neurons that fired as the rats sampled odors (like the cell depicted in Fig. 3A) formed an apparent place field around the location of the cups containing odor stimuli. Importantly, the activity was not purely determined by the location of the rat. To confirm the specificity of the cells we defined a place field based on the apparent spatial sensitivity. We then constructed raster plots (shown in the insets below Figs. 3A,B) using each journey the rat took through the field. We finally sorted the rasters based on the temporal proximity of the event in question. Often (e.g., Fig. 3; additional examples in Supporting Information

Fig. 4), the rasters showed that activity was more closely locked to the timing of events (magenta line) than the entry into the place field (time 0). This stands in contrast to the pattern seen in the more purely spatial response shown in the left panel below Figure 3B: this cell increased its activity as the rat entered the place field, regardless of the temporal proximity of the trial start.

To investigate changes in an event locked frame of reference we constructed peri-event time histograms (PETHs) by binning unit response data aligned on the start and end of each trial. These two events were selected because they were highly salient (the divider was removed and replaced at the start and end, see Fig. 1), we could precisely identify them from the video record (when the rat crossed into and out of the odor sampling area), and because other key events reliably occurred in close temporal proximity. For example, because the rats immediately approached the cups, initial odor sampling always occurred shortly after trial start. Similarly, because the rats returned to the ITI side of the box after the trial, the rat always obtained the reward just before the trial end. Thus, by centering our analyses on these two events we were able to incorporate several key events known to affect hippocampal firing (e.g., cup approach, odor sampling, reward, etc.; see Eichenbaum et al., 1987; Smith and Mizumori, 2006b) all within the same temporal reference frame. Figures 3C,D show ensemble firing rates (the average PETH of each neuron), on both sessions, sorted based on the time of peak activity during A_{final} . When rats learned the new problem set in the same context (Fig. 3C) the pattern was largely preserved: neurons that showed elevated activity before/after the start/end of trials continued to show elevated activity at similar times. Neurons recorded from DC rats (Fig. 3D), however, changed their event-related firing. When the PETHs were sorted by peak firing time on A_{final} , and the same sorting order applied to B_{initial} , the temporal structure was completely abolished.

Similar to the position-locked analyses, which included apparent place sensitivity that likely arose from event sensitive firing, the event locked analyses included apparent event responses that likely arose from spatial firing. For instance, a cell with a place field near the divider (like the cell depicted in Fig. 2A) would show an apparent event response as the rat passed through the field around the start of the trials, and again at the end of the trials. Indeed, this sort of activity can be seen in Figure 3C: several of the rows show firing that is symmetrical around the midline of the plot. By applying both sets of analyses to all cells, ignoring the ‘cause’ of firing, we were able to measure remapping without categorizing cells based on the factors shaping their receptive fields.

To quantify the extent to which event-locked activity changed across the critical sessions in the two groups, we applied an analogous set of analyses to those which we used for space, but anchored to the timing of events rather than the location of the rats. We first applied a binwise correlation, analogous to the spatially binned correlation shown in Figure 2C. For each neuron we compared the average activity in each of the twenty 100ms bins surrounding the start and end of

trials across the two critical sessions (Fig. 3E). Event related firing was well correlated in neurons recorded from SC rats (one sample t -test compared with 0; trial start: $t_{(45)} = 8.18, P < 0.01 \times 10^{-7}$; trial end: $t_{(44)} = 9.22, P < 0.01 \times 10^{-9}$) and showed no relationship in DC rats (trial start: $t_{(38)} = 0.16, P = 0.87$; trial end: $t_{(45)} = 0.30, P = 0.76$). These data indicate that when neurons were considered in an event-locked reference frame (as in the spatial frame), the ensemble as a whole underwent a complete reorganization of activity for rats that learned the new odor set in a new context.

Just as cells with spatially tuned responses could remap either by changing the location of peak activity or the degree to which they were tuned to space, event sensitive neurons could change either their peak activity time (with respect to an event) or the degree to which they showed event locked activity altogether. As such, we investigated the extent to which units shifted the event-related firing time, using an analysis that was analogous to the spatial COM shift shown in Figure 2D. We defined a PETH COM by fitting Gaussian curves to the PETH data (aligned on trial start and end, ± 2 seconds, 100 ms bins), and identifying the time of the peak (Supporting Information Fig. 5). To be included in the analysis, both sessions had to be best fit by a Gaussian that had positive amplitude (i.e., a peak in response). The shift in peak times was smaller in SC rats than DC rats (Fig. 3F; trial start: $t_{(61)} = 4.09, P < 0.001$, trial end: $t_{(65)} = 4.60, P < 0.0001$). As with the spatial COM, the PETH peak shifts for neurons recorded from rats trained in the DC condition were as large as the average shift between the peaks of two different randomly selected neurons (one sample t -test, trial start: $t_{(26)} = 0.37, P = 0.71$; trial end: $t_{(31)} = 1.02, P = 0.32$). In contrast, COM shifts for rats in the SC condition were significantly smaller than chance (trial start: $t_{(35)} = 11.57, P < 0.01 \times 10^{-9}$; trial end: $t_{(34)} = 7.30, P < 0.01 \times 10^{-5}$).

For each neuron we calculated an event related information score for the trial start and trial end using a method analogous to the approach used for spatial information (i.e., considering how informative firing was about the time of an event, rather than location). This approach asks whether neurons that have responses which are tightly linked with an event in one session continue to be associated with an event in the next. Similar to spatial information scores, the information about events was highly correlated across sessions in neurons recorded from SC rats (Fig. 3G; r compared with 0; trial start: $t_{(3)} = 32.55, P < 0.0001$; trial end: $t_{(3)} = 11.32, P < 0.01$) but not DC rats (r compared with 0; trial start: $t_{(2)} = 0.32, P = 0.78$; trial end: $t_{(2)} = 1.15, P = 0.37$). Thus, neurons that had firing that was informative about proximal events retained this information in SC rats, but in DC rats the degree event-related information in one session was unrelated to the degree of event-related information in the other.

Ensemble Predictions of Behavior

The above analyses showed that both spatial and event representations persisted in SC rats and underwent complete remap-




FIGURE 4. Population vector prediction of location and temporal epoch within and across sessions. Linear discriminant analysis was used to classify neural ensemble responses associated with one of four spatial quadrants or one of four temporal epochs (an example is illustrated in Supporting Information Fig. 6). A and B show the accuracy of the spatial and epoch classifiers, respectively, when trained on half of the data in A_{final} , and then tested on the remainder (orange bars), and when the same classifiers were tested on $B_{initial}$ (green bars). To measure the performance of the classifiers, training and testing subsets were selected randomly using an iterative process (see methods). Heights of the bars indicate the median, and error bars range from 5% to 95%, of the iterated distribution. Error bars not overlapping with 0.25 (horizontal broken line) indicate that the classifier performed above chance. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

ping in DC rats. However, recent work has highlighted just how much of the ensemble signal is either inaccessible, or obscured, when measured at the level of individual units (e.g., Kelemen and Fenton, 2010; Jezek et al., 2011; Dupret et al., 2013). We next compared ensembles as a whole, on the same critical days as the single unit analyses, using metrics sensitive to interactions between neurons (i.e., the covariance of the ensemble). Ensembles consisted of all of the neurons recorded from each rat (number of neurons from SC rats: 7, 10, 15, 19; number of neurons from DC rats: 16, 15, 17).

We began with a discriminant classification approach to predict location and temporal epoch relative to task events using each rat's neural ensemble responses (Fig. 4, see also Supporting Information Fig. 6). We first divided the apparatus into four spatial quadrants, and in a corresponding manner divided the trials into four epochs surrounding the time of the events. We binned unit activity to form population vectors containing the firing rates of simultaneously recorded

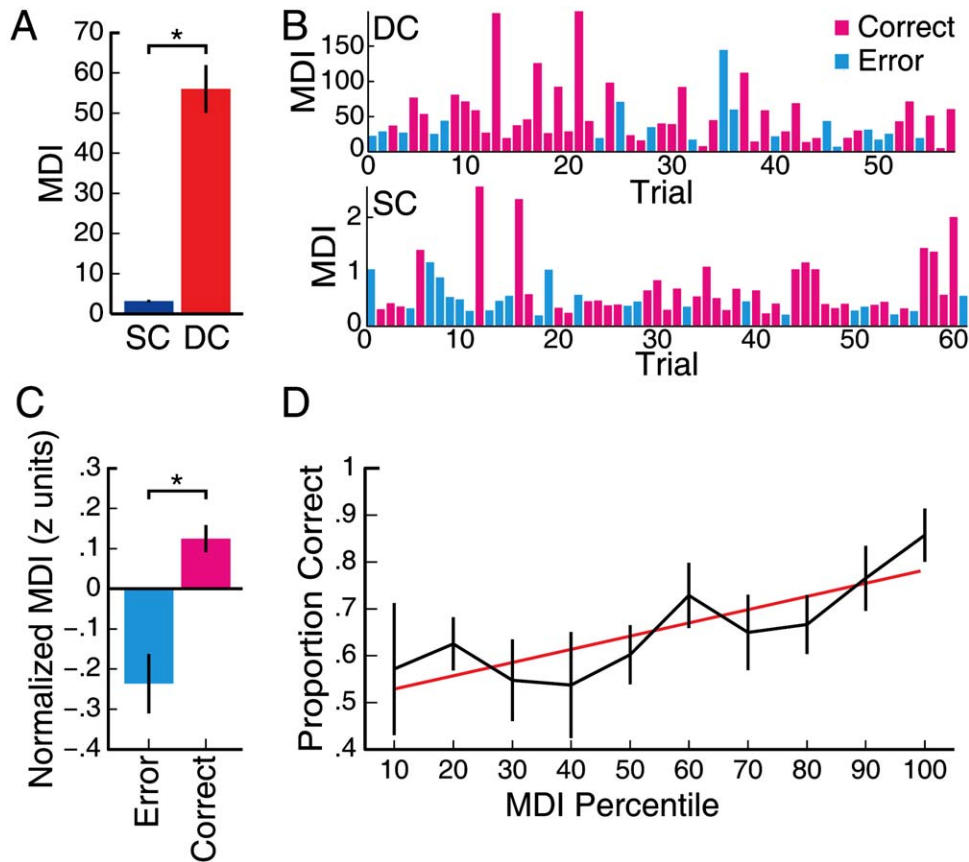


FIGURE 5. Ensemble distances between A_{final} and B_{initial} trial representations. Dissimilarity, quantified as the distance between each B_{initial} population vector (i.e., the average firing rate of each neuron throughout entire trials) and the set of A_{final} population vectors was calculated using Mahalanobis distance, and then normalized to the number of neurons to yield a distance index (MDI). (A) Ensembles showed far greater dissimilarity in DC rats (red bar) than SC rats (blue bar). Error bars indicate SEM across rats. (B) MDI for each trial in an example DC rat (top) and SC rat (bottom). Correct trials are indicated by magenta bars and error trials in cyan. On trials with a high MDI, the rat rarely

made errors. Note that the large difference in the scale of the ordinate: though both rats show the same pattern, the distribution of MDI values spanned a larger range in DC rats. C shows the average z-scored MDI across rats (both SC and DC) was lower for error than correct trials. Error bars indicate SEM across rats. D shows the proportion of correct trials, across rats (both SC and DC), for each decile of MDI. Each rat was most likely to make an error when MDI was at the low end of that rat's distribution. The red line indicates a linear regression fit to the data ($r^2 = 0.12$; $P < 0.005$). [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

neurons, labeled these vectors with the associated quadrant/epoch, and then trained linear discriminant classifiers using a subset of vectors from A_{final} . To avoid errors associated with overfitting, and to estimate confidence intervals on the performance of the models, we used an iterative approach (see methods).

When the classifiers were trained and tested using data from the same session, they generally performed well above chance (orange bars in Fig. 4). In SC rats, both the spatial (Fig. 4A) and epoch-based (Fig. 4B) classifiers continued to perform well when trained using data from A_{final} and tested on the responses in B_{initial} . In DC rats, however, the performance of the classifiers fell to chance when applied across sessions. These results indicate that the population codes for space and temporal epoch persisted in SC rats across sessions, but in DC rats the code was rendered useless. Note that in Rat SC2 the temporal

epoch classifier performed at chance when tested within A_{final} , as such it was unsurprising that it continued to perform at chance when tested across sessions.

To measure the dissimilarity between ensemble representations on a trial-to-trial basis, we computed whole trial length population vectors marking the average firing rate of each neuron on each trial. We then computed the distance between each B_{initial} trial to the average vector for the set of A_{final} trials using Mahalanobis distance, which was normalized to the number of neurons in the sample to define a distance index (MDI). Mahalanobis distance quantifies the distance of an n -dimensional point (the population vector from a B_{initial} trial) to the average of a group of points (the vectors from all A_{final} trials), scaled by the covariance of the group. This provides a metric that is sensitive to changes both in global rate as well as interactions between neurons, and it has been used previously

for quantifying dissimilarity of ensemble responses (Hyman et al., 2012; Manns et al., 2007; Sheinberg and Logothetis, 1997).

The average MDI was far greater in DC rats than SC rats (Fig. 5A; $t_{(5)} = 3.35$, $P < 0.05$). The finding that this metric was sensitive to the remapping of ensemble activity is actually somewhat surprising: trial durations were long and variable (mean = 10.0 seconds, SD = 4.5 seconds). Furthermore, the trajectory of the rat, the specific odors encountered and even the number of odor sampling events, varied from one trial to the next. The differential effects of the SC and DC conditions on hippocampal firing patterns cannot be attributed to large scale differences in behavior in the two conditions. For example, running speed did not differ across groups or across the switch from B_{final} to A_{initial} (Supporting Information Fig. 1B). Nevertheless, the switch from B_{final} to A_{initial} did cause a large increase in the error rate and, consequently, an increase in bouts of odor investigation and digging. However, the number of errors (and related behaviors) was greatest in the SC condition where the hippocampal representations were stable and conversely, the biggest changes in hippocampal representations were seen in the DC condition where there were significantly fewer errors. Thus, the changes in hippocampal representations could not have been driven by these changes in behavior.

Interestingly, in both context conditions, MDI was smaller on those trials in which rats made errors. Figure 5B shows each trial's MDI in a B_{initial} session from a DC rat (top) and SC rat (bottom). When the MDI was high (i.e., the representation was most distinct), the rats rarely made errors. When the MDI was low (i.e., the current representation was more similar to the old, A_{final} representation), interference dominated and the rats performed near chance (Figs. 5C,D). Interestingly, the precise quantity of representational dissimilarity alone was not sufficient to prevent interference: SC rats showed distributions of MDI values that were far lower than DC rats. However, when considering the relative MDI (either using z-transformed data or percentile division) it was clear that correct performance was associated with more distinct representations than incorrect performance (Figs. 5C,D; $t_{(6)} = 3.38$, $P < 0.05$). Taken together, these results show how the remapping of hippocampal ensemble representations provides a context signal that can be used to alleviate the effects of interference.

DISCUSSION

Our results show a complete reorganization of hippocampal activity patterns following a change in context, and this was directly linked to subjects' success in preventing interference. Individual neurons showed responses that were tightly linked to the position of the animal, the timing of critical trial events, or some combination thereof. However, these neurons changed their response when the rat entered a new context, resulting in unique ensemble representations that could be used to identify the current context and distinguish it from other contexts. Interestingly,

this large scale representational shift was observed regardless of the specific response type, and both spatial and event locked analyses indicated statistically independent representations.

The extent of ensemble remapping was predictive of performance. In SC rats the ensemble activity patterns of Set A persisted during Set B, the rats were subject to interference and made more errors. In DC rats a new representation emerged, insulating rats from interference. As such these rats made fewer errors. Moreover, in both groups, rats were most likely to make errors on individual Set B trials that showed ensemble activity similar to Set A patterns. When an old representation was active, the rats were impeded by interference from the associated memories. These data provide the first direct link between endogenously generated patterns of neural activity in the hippocampus and the use of context to prevent interference. The results establish that one function of complete remapping is to insulate memories from interference based on the learning context. Whether or not rate remapping of hippocampal ensembles can also be used to support the separation of context appropriate memories remains unknown. While we observed large changes in firing rate in neurons with stable receptive fields recorded from rats in the SC condition, it was unclear whether these changes were directly related to learning or merely occurred due to the passage of time (Mankin et al., 2012). Nonetheless, rats in both groups were best equipped to resolve interference on trials with the most dissimilar hippocampal representations, suggesting that even in an unchanged environment a distinct representation supports learning of new memories.

Previous experiments using the same proactive interference task showed that temporary inactivation of the dorsal hippocampus severely impaired performance in the DC condition (Butterly et al., 2012). However, inactivation had no effect on performance in the SC condition, suggesting that the hippocampus was not needed for remembering the odors themselves, but instead was critical for the ability to use the context to cue the relevant odor memories. The present results suggest that the hippocampus allows DC rats to perform better by supplying unique ensemble representations so that memories can be linked to the context in which they were learned. Consistent with this idea, recent studies have shown that artificially reactivating hippocampal ensembles associated with a particular context is sufficient to induce a contextually appropriate fear response (Ramirez et al., 2013). Our results are also consistent with recent work showing that hippocampal ensemble representations encode a complex hierarchical set of relationships among olfactory cues, reward, and location, with a code representing context at the top of the hierarchy (McKenzie et al., 2014). Our results demonstrate that these ensemble codes support the capacity to prime contextually appropriate memories and behavior, which is particularly useful for mitigating the effects of interference.

We propose that hippocampal firing patterns influence memories stored in extra-hippocampal, presumably neocortical, locations. Learning is undoubtedly associated with plastic changes in sensory cortical regions (Harris et al., 2001; Ghose, 2004). Moreover, many kinds of memory are preserved following hippocampal damage yet become highly sensitive to hippocampal

damage when a contextual component is present (Eichenbaum et al., 1988; Phillips and LeDoux, 1992; Honey and Good, 1993; Eacott and Norman, 2004; Butterly et al., 2012; for review see: Nadel, 2008). Work in patients with hippocampal pathology has indicated that semantic memory (which is context independent) is preserved while episodic memory (which includes contextual information; see: Tulving and Markowitsch, 1998) is completely abolished. Coupled with the idea that the hippocampus serves as an index that activates memories stored elsewhere (Teyler and DiScenna, 1985, 1986), our results suggest that a key function of the hippocampus is to supply an ensemble representation of context that can serve to prime memories encoded in extra-hippocampal circuits. Because the context-appropriate memories are strongly primed, subjects are less susceptible to interference from memories that belong to other contexts (Smith and Bulkin, 2014).

In addition to group effects, we observed trial by trial variation in the representation that was predictive of performance. On many trials from the new problem set (B_{initial}), the representation of DC rats was similar to the previous day (A_{final}) even though the rat was in an entirely new environmental context (e.g., the lower MDI values in the DC example in Fig. 5B). How can two statistically independent representations, as we observed in DC rats, vary in their similarity? Several recent studies have demonstrated that, at the population level, hippocampal representations can alternate on a moment to moment basis (Kelemen and Fenton, 2010; Jezek et al., 2011; Dupret et al., 2013). In fact, new evidence suggests that the transition from one hippocampal context representation to another is accomplished by rapid “flickering” between the representations, with the old pattern becoming more infrequent as the new pattern becomes more prominent (Jezek et al., 2011). A similar process, with transient appearances of the A_{final} representation, may have been at work in our DC rats. However, the complexity of our task and the fact that firing patterns were influenced by a combination of spatial location and ongoing task events made it impossible to observe these transitions.

Interestingly, SC Rats also showed trial by trial variability in their representations and like the DC rats, more distinct representations were associated with better performance. This may also reflect the early stages of a transition to a new hippocampal representation. Many studies have shown that hippocampal neurons develop new representations in response to changes in non-spatial aspects of the context, including shifts in the task demands, behavioral strategies and motivational state (e.g., Wood et al., 2000; Smith and Mizumori, 2006b; Eschenko and Mizumori, 2007; Kennedy and Shapiro, 2009). In previous studies (Smith and Mizumori, 2006b), we found that new hippocampal representations emerge slowly when subtle changes in task demands can only be detected through trial and error. In the SC condition, the change from a well learned odor discrimination problem set to a new, conflicting problem set may have been enough to trigger the beginning of that process. Alternatively, differences between correct and error trials might have arisen from a representation that was less strongly driven on error trials. Others have suggested that cells which

show robust firing on correct trials fire less on trials in which the rat makes an error (Robitsek et al., 2013). Although we found no direct evidence for systematic differences between the gross firing rates on correct and error trials, a slightly degraded hippocampal context signal could have resulted in less robust retrieval of the old, potentially interfering odor memories, and improved performance. That is, a momentarily poorer hippocampal representation could, paradoxically, result in better performance due to less robust retrieval of interfering memories.

The behavioral and physiological complexity in the present experiment presented a unique set of analytic challenges, but also a particularly rich data set with a variety of neural response profiles. The location of the rat at any given moment, the order and the timing of odor sampling, were all experimentally uncontrolled—we left these decisions to the rat. This approach provided data that highlights the heterogeneity of response properties of individual neurons. We generally applied a strategy of “agnostic” analyses, subjecting all cells to measures of both spatial and temporal response change. Regardless of the framework, ensemble activity was sufficient to identify the context. Our population distance analyses relied exclusively on neural covariance and did not use information about the rats’ position or the timing of task events. The results show that the ensemble code is directly linked to the ability to use the context to prevent interference.

ACKNOWLEDGMENTS

We thank undergraduates Leela Patel, Carly Britton and Luke Grosvenor for assistance with animal training and electrode hyperdrive fabrication. We would also like to thank professors Sheri Mizumori, Joseph Fetcho, Ronald Hoy, and Melissa Warden for feedback on the manuscript. This work was supported by an NIH grant R01 MH083809 to D.M.S.

REFERENCES

- Alme CB, Miao C, Jezek K, Treves A, Moser EI, Moser MB. 2014. Place cells in the hippocampus: Eleven maps for eleven rooms. *Proc Natl Acad Sci USA* 111:18428–18435.
- Balsam PD, Tomie A. 1985. *Context and Learning*. Hillsdale, N.J.: L. Erlbaum Associates.
- Bilodeau IM, Schlosberg H. 1951. Similarity in stimulating conditions as a variable in retroactive inhibition. *J Exp Psychol* 41:199–204.
- Bouton ME. 1993. Context, time, and memory retrieval in the interference paradigms of Pavlovian learning. *Psychol Bull* 114:80–99.
- Butterly DA, Petroccione MA, Smith DM. 2012. Hippocampal context processing is critical for interference free recall of odor memories in rats. *Hippocampus* 22:906–913.
- Colgin LL, Moser EI, Moser MB. 2008. Understanding memory through hippocampal remapping. *Trends Neurosci* 31:469–477.
- Cooper BG, Mizumori SJ. 2001. Temporary inactivation of the retrosplenial cortex causes a transient reorganization of spatial coding in the hippocampus. *J Neurosci* 21:3986–4001.
- Dallett K, Wilcox SG. 1968. Contextual stimuli and proactive inhibition. *J Exp Psychol* 78:475.

- Dupret D, O'Neill J, Csicsvari J. 2013. Dynamic reconfiguration of hippocampal interneuron circuits during spatial learning. *Neuron* 78:166–180.
- Eacott MJ, Norman G. 2004. Integrated memory for object, place, and context in rats: A possible model of episodic-like memory? *J Neurosci* 24:1948–1953.
- Eichenbaum H, Kuperstein M, Fagan A, Nagode J. 1987. Cue-sampling and goal-approach correlates of hippocampal unit activity in rats performing an odor-discrimination task. *J Neurosci* 7:716–732.
- Eichenbaum H, Fagan A, Mathews P, Cohen NJ. 1988. Hippocampal system dysfunction and odor discrimination learning in rats: Impairment or facilitation depending on representational demands. *Behav Neurosci* 102:331–339.
- Eschenko O, Mizumori SJ. 2007. Memory influences on hippocampal and striatal neural codes: Effects of a shift between task rules. *Neurobiol Learn Mem* 87:495–509.
- Ghose GM. 2004. Learning in mammalian sensory cortex. *Curr Opin Neurobiol* 14:513–518.
- Godden DR, Baddeley AD. 1975. Context-dependent memory in two natural environments: On land and underwater. *Br J Psychol* 66:325–331.
- Greenspoon J, Ranyard R. 1957. Stimulus conditions and retroactive inhibition. *J Exp Psychol* 53:55–59.
- Harris JA, Petersen RS, Diamond ME. 2001. The cortical distribution of sensory memories. *Neuron* 30:315–318.
- Honey RC, Good M. 1993. Selective hippocampal lesions abolish the contextual specificity of latent inhibition and conditioning. *Behav Neurosci* 107:23–33.
- Hyman JM, Ma L, Balaguer-Ballester E, Durstewitz D, Seamans JK. 2012. Contextual encoding by ensembles of medial prefrontal cortex neurons. *Proc Natl Acad Sci USA* 109:5086–5091.
- Jeffery KJ, Gilbert A, Burton S, Strudwick A. 2003. Preserved performance in a hippocampal-dependent spatial task despite complete place cell remapping. *Hippocampus* 13:175–189.
- Jezek K, Henriksen EJ, Treves A, Moser EI, Moser MB. 2011. Theta-paced flickering between place-cell maps in the hippocampus. *Nature* 478:246–249.
- Kelemen E, Fenton AA. 2010. Dynamic grouping of hippocampal neural activity during cognitive control of two spatial frames. *PLoS Biol* 8:e1000403.
- Kennedy PJ, Shapiro ML. 2009. Motivational states activate distinct hippocampal representations to guide goal-directed behaviors. *Proc Natl Acad Sci USA* 106:10805–10810.
- Kim JJ, Fanselow MS. 1992. Modality-specific retrograde amnesia of fear. *Science* 256:675–677.
- Kraus BJ, Robinson 2nd RJ, White JA, Eichenbaum H, Hasselmo ME. 2013. Hippocampal “time cells”: Time versus path integration. *Neuron* 78:1090–1101.
- Law LM, Smith DM. 2012. The anterior thalamus is critical for overcoming interference in a context-dependent odor discrimination task. *Behav Neurosci* 126:710–719.
- 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.
- Leutgeb JK, Leutgeb S, Moser MB, Moser EI. 2007. Pattern separation in the dentate gyrus and CA3 of the hippocampus. *Science* 315:961–966.
- 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.
- Liu X, Ramirez S, Pang PT, Puryear CB, Govindarajan A, Deisseroth K, Tonegawa S. 2012. Optogenetic stimulation of a hippocampal engram activates fear memory recall. *Nature* 484:381–385.
- Luu P, Sill OC, Gao L, Becker S, Wojtowicz JM, Smith DM. 2012. The role of adult hippocampal neurogenesis in reducing interference. *Behav Neurosci* 126:381–391.
- Mankin EA, Sparks FT, Slayyeh B, Sutherland RJ, Leutgeb S, Leutgeb JK. 2012. Neuronal code for extended time in the hippocampus. *Proc Natl Acad Sci USA* 109:19462–19467.
- Manns JR, Howard MW, Eichenbaum H. 2007. Gradual changes in hippocampal activity support remembering the order of events. *Neuron* 56:530–540.
- Markus EJ, Barnes CA, McNaughton BL, Gladden VL, Skaggs WE. 1994. Spatial information content and reliability of hippocampal CA1 neurons: Effects of visual input. *Hippocampus* 4:410–421.
- McKenzie S, Frank AJ, Kinsky NR, Porter B, Rivière PD, Eichenbaum H. 2014. Hippocampal representation of related and opposing memories develop within distinct, hierarchically organized neural schemas. *Neuron* 83:202–215.
- Mehta MR, Barnes CA, McNaughton BL. 1997. Experience-dependent, asymmetric expansion of hippocampal place fields. *Proc Natl Acad Sci USA* 94:8918–8921.
- 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.
- Muzzio IA, Levita L, Kulkarni J, Monaco J, Kentros C, Stead M, Abbott LF, Kandel ER. 2009. Attention enhances the retrieval and stability of visuospatial and olfactory representations in the dorsal hippocampus. *PLoS Biol* 7:e1000140.
- Nadel L. 2008. The hippocampus and context revisited. In: Mizumori SJ, editor. *Hippocampal Place Fields: Relevance to Learning and Memory*. New York: Oxford University Press. pp. 3–15.
- 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.
- Pastalkova E, Itskov V, Amarasingham A, Buzsaki G. 2008. Internally generated cell assembly sequences in the rat hippocampus. *Science* 321:1322–1327.
- Peters GJ, David CN, Marcus MD, Smith DM. 2013. The medial prefrontal cortex is critical for memory retrieval and resolving interference. *Learn Mem* 20:201–209.
- Phillips RG, LeDoux JE. 1992. Differential contribution of amygdala and hippocampus to cued and contextual fear conditioning. *Behav Neurosci* 106:274–285.
- Ramirez S, Liu X, Lin PA, Suh J, Pignatelli M, Redondo RL, Ryan TJ, Tonegawa S. 2013. Creating a false memory in the hippocampus. *Science* 341:387–391.
- Robitsek RJ, White JA, Eichenbaum H. 2013. Place cell activation predicts subsequent memory. *Behav Brain Res* 254:65–72.
- Sheinberg DL, Logothetis NK. 1997. The role of temporal cortical areas in perceptual organization. *Proc Natl Acad Sci USA* 94:3408–3413.
- Smith SM. 1979. Remembering in and out of context. *J Exp Psychol Hum Learn Mem* 5:460.
- Smith SM. 1988. Environmental context-dependent memory. In: Davies G, Thomson DM, editors. *Memory in Context: Context in Memory*. New York: John Wiley and Sons Ltd. pp. 13–34.
- Smith DM. 2008. Chapter 4.4 The hippocampus, context processing and episodic memory. In: Ekrem Dere AELN, Joseph PH, editors. *Handbook of Behavioral Neuroscience*. New York: Elsevier. pp. 465–630.
- Smith SM, Vela E. 2001. Environmental context-dependent memory: A review and meta-analysis. *Psychon Bull Rev* 8:203–220.
- Smith DM, Mizumori SJ. 2006a. Hippocampal place cells, context, and episodic memory. *Hippocampus* 16:716–729.
- Smith DM, Mizumori SJ. 2006b. Learning-related development of context-specific neuronal responses to places and events: The hippocampal role in context processing. *J Neurosci* 26:3154–3163.
- Smith DM, Bulkin DA. 2014. The form and function of hippocampal context representations. *Neurosci Biobehav Rev* 40C:52–61.
- Taylor TJ, DiScenna P. 1985. The role of hippocampus in memory: A hypothesis. *Neurosci Biobehav Rev* 9:377–389.
- Taylor TJ, DiScenna P. 1986. The hippocampal memory indexing theory. *Behav Neurosci* 100:147–154.

- Tulving E, Markowitsch HJ. 1998. Episodic and declarative memory: Role of the hippocampus. *Hippocampus* 8:198–204.
- Watts GH, Royer JM. 1969. Stimulus context and retroactive inhibition in free recall. *Psychon Sci* 17:253–254.
- Wills TJ, Lever C, Cacucci F, Burgess N, O'Keefe J. 2005. Attractor dynamics in the hippocampal representation of the local environment. *Science* 308:873–876.
- Wood ER, Dudchenko PA, Eichenbaum H. 1999. The global record of memory in hippocampal neuronal activity. *Nature* 397:613–616.
- Wood ER, Dudchenko PA, Robitsek RJ, Eichenbaum H. 2000. Hippocampal neurons encode information about different types of memory episodes occurring in the same location. *Neuron* 27:623–633.