

HHS Public Access

Author manuscript *Hippocampus*. Author manuscript; available in PMC 2021 March 22.

Published in final edited form as: *Hippocampus.* 2021 March ; 31(3): 244–260. doi:10.1002/hipo.23272.

Hippocampal spatial memory representations in mice are heterogeneously stable

Samuel J Levy, Nathaniel R Kinsky, William Mau, David W Sullivan, Michael E Hasselmo Center for Systems Neuroscience, Boston University, 610 Commonwealth Ave., Boston, MA 02215; Graduate Program in Neuroscience, Boston University

Abstract

The population of hippocampal neurons actively coding space continually changes across days as mice repeatedly perform tasks. Many hippocampal place cells become inactive while other previously silent neurons become active, challenging the idea that stable behaviors and memory representations are supported by stable patterns of neural activity. Active cell replacement may disambiguate unique episodes that contain overlapping memory cues, and could contribute to reorganization of memory representations. How active cell replacement affects the evolution of representations of different behaviors within a single task is unknown. We trained mice to perform a Delayed Non-Match to Place (DNMP) task over multiple weeks, and performed calcium imaging in area CA1 of the dorsal hippocampus using head-mounted miniature microscopes. Cells active on the central stem of the maze "split" their calcium activity according to the animal's upcoming turn direction (left or right), the current task phase (study or test), or both task dimensions, even while spatial cues remained unchanged. We found that, among reliably active cells, different splitter neuron populations were replaced at unequal rates, resulting in an increasing number of cells modulated by turn direction and a decreasing number of cells with combined modulation by both turn direction and task phase. Despite continual reorganization, the ensemble code stably segregated these task dimensions. These results show that hippocampal memories can heterogeneously reorganize even while behavior is unchanging.

Introduction

The idea that stable behaviors and reliable memory representations are supported by stable elements of neural circuits (Barnes et al. 1997; Thompson & Best, 1990) has been challenged by many findings that neural circuit components across the brain are unstable over time. Circuit instability is notable in the continual replacement of active cells with previously silent cells (Kinsky et al., 2018; Mau et al., 2018; Ziv et al., 2013), but is also observed in the impermanence of dendritic spines and axonal boutons (Attardo et al. 2015; Pfeiffer et al. 2018; Grutzendler et al. 2002; De Paola et al. 2006). How circuit instability

Corresponding Author: Samuel Levy, samjl@bu.edu.

Contributions: Study was designed by Howard Eichenbaum and SJL. WM performed surgeries and histology. SJL performed experiments and analysis. NRK, WM and DWS wrote TENASPIS. SJL and MEH wrote the manuscript.

Conflict of interest: The authors declare no conflicts of interest.

may affect neural function is a topic of much debate (Chambers & Rumpel, 2017; Rule et al. 2019).

In the hippocampus, a hub for episodic memory and spatial navigation, change is observed in the patterns neuronal of activity and the set of currently active cells. In behaving animals, single neurons become more sensitive to task demands during training and change their firing properties to more precisely encode task demands (Kobayashi et al. 2003; Komorowski et al. 2009; Lever et al. 2002). Hippocampal memory representations are also unstable even during over-trained behaviors, exhibiting a decorrelation in ensemble activity relative to the elapsed time between recordings (Mankin et al. 2015; Mankin et al. 2012; Rubin et al. 2015; Ziv et al. 2013). These decorrelations result both from remapping of firing locations exhibited by continuously active single neurons that is unrelated to changes in behavior (Mehta et al. 2000; Poe et al. 2000; Lee et al. 2006; Law et al. 2016), and from population dynamics that include the continual inactivation of active cells and their replacement by previously silent cells (Mau et al., 2018; Ziv et al., 2013). However, these changes have primarily been observed during learning or during performance of foraging tasks. How changes occur during stable performance of a multi-dimensional memory task remains an open question. Previous studies have linked the long term stability of a neuronal activity to different spatial locations and different task behaviors (Kentros et al., 2004; Kinsky et al., 2020; Taxidis et al., 2018). We sought to expand on these studies by examining how different demands on long term memory influence the evolution of hippocampal memory representations during a task where mice pass through the same spatial location under multiple different task conditions.

To study the reorganization of hippocampal representations over time, we used *in vivo* calcium imaging to monitor the activity of hundreds of neurons across multiple sessions in mice performing a Delayed Non-Match to Place task on a figure-eight maze. We first confirmed that neurons modulate their activity on the central stem according to the animal's upcoming turn direction and the current task phase (Griffin et al., 2007; Wood et al., 2000). We show that the distribution of these single unit responses among the reliably active population changes over time, resulting in an increased number of turn direction-modulated neurons and a decrease in the number of neurons modulated by both the current task phase and upcoming turn direction. These changes primarily result from the unequal recruitment of previously inactive cells to different neuron coding types. While the distribution of single unit activity was unstable among reliably active cells, population analyses revealed a stable separation of task variables in the collective ensemble at extended lags between recordings. These results demonstrate that behavior and population output can remain stable while single neuron responses are unevenly reorganized.

Methods

Surgical Procedures

4 male, naïve mice (C57BL6, Jackson Laboratory) underwent two stereotaxic surgeries to prepare for calcium imaging. All procedures presented here were approved by the Institutional Animal Care and Use Committee (IACUC) at Boston University. Mice were given 0.05mL/kg buprenorphine as a pre-surgical analgesic, and were anesthetized with ~1%

isofluorane delivered with oxygen. The first surgery was to infuse virus to express GCaMP6f. A small craniotomy was made above the dorsal hippocampus at AP –2.0mm, ML +1.5mm relative to bregma, and the infusion needle was lowered at this site to DV –1.5mm. 350 nL of the viral vector AAV9-Syn-GCaMP6f (University of Pennsylvania Vector Core, obtained at a titer of ~4×10e13GC/mL and diluted it to ~5–6×10e12GC/mL with 0.05M phosphate buffered saline) was infused at 40nL/min and allowed to diffuse for 15 minutes before the infusion needle was slowly removed.

The second surgery, to implant a gradient-index (GRIN) lens for imaging, was performed three weeks later to allow for viral infection and GCaMP6f expression. A 2mm diameter circular craniotomy was made at AP-2.25mm, ML +1.8mm, and the neocortex was aspirated until rostral-caudal fiber tracts of the alveus were visible. Near-freezing 0.9% saline solution and GelFoam (Pfizer) were used continuously to control bleeding and to dry the base of the craniotomy prior to lens implantation. The GRIN lens (1mm diameter, 4mm length, Inscopix) was slowly lowered stereotaxically to 200 um dorsal to the infusion site of the virus, measured relative to the skull surface. The lens was then fixed in place using a non-bioreactive silicone polymer (Kwik-Sil, World Precision Instruments) to entirely cover the craniotomy, which was then covered with Metabond dental cement (Parkell) to anchor the lens to the skull. The lens was covered with a temporary cap made from Kwik-Cast (World Precision Instruments) until the baseplate was attached.

After allowing a week of recovery from the lens implantation surgery, mice were again anesthetized and placed in the stereotaxic holder. The baseplate was magnetically attached to the imaging microscope camera, which was then aligned parallel to the GRIN lens by adjusting until the edge of the lens was entirely in focus in the nVista recording software (Inscopix). The camera with baseplate was then lowered until GCaMP6f-expressing cells were optimally in focus, and then raised by 50 um to allow for shrinkage of the dental cement used to affix the baseplate. The baseplate was then fixed in place to the existing metabond around the GRIN lens with Flow-It ALC Flowable Composite (Pentron), and cured with ultraviolet light. Gaps in the dental cement were filled in with Metabond, the camera was removed, and a cover attached to the baseplate.

Maze Description

The maze was constructed from wood and the internal floor area measured 64.5 cm long by 29.2 cm wide, and walls were 17.75 cm high. Middle maze walls separated this area into a central hallway (Center Stem) and left and right Return Arms. Each hallway was 7.5 cm wide. This resulted in low variability of the animals' left/right position within a hallway, although it did not prevent the animals from occasionally running with their head turned towards one side. Rewards were delivered through ports at the maze walls at floor level of the side arms 12 cm from the delay-end of the maze. To dictate turn direction on Study Trials (see below) and to contain the mouse during the delay period, arm barriers were used that were made of transparent plastic. The delay barrier was made of wood. In this manuscript we only consider data from the central stem and return arms.

For analysis of the central stem, we chose a region starting ~8 cm in front of the delay barrier and extending 30cm to end ~5 cm before the choice region at the end of the middle

maze walls; this region was selected to encompass the region where the mouse was running similarly between study and test task phases and left and right turn directions (Figure 1d, purple). Left and right variability in the animals' head position at the end of this region was less than 2.5 times the standard deviation of the animals' left/right variability for the first half of the stem, and was usually indistinguishable by visual observation in behavioral recordings. We divided this 30cm long region into 8 spatial bins each 3.75 cm in length. For the return arms (Figure 1d, orange), we chose a 30cm stretch that started after the animals had fully entered the return arms and ended before they reached the reward zone, also separated into 8 bins each 3.75 cm.

Behavior pre-training and recording sequence

Mice were trained to run on a Delayed Non-Match to Place (DNMP) task shown in Figure 1. This involved extensive pre-training in order to obtain performance at the criterion of 70% correct.

After fully recovering from surgeries, mice were extensively handled for ~15 min/day for 5 days. They were simultaneously food restricted to 80% of free feeding body weight, and acclimated to consuming chocolate sprinkles. Over the next two weeks, mice were given time to explore the maze, and were slowly shaped to run in a single direction through the maze and to receive reward, with inserted walls to block paths and guide them. In the last few days of pre-training, mice were guided with blocking walls to alternate between the two reward arms and given experience with continuous and delayed alternation.

Mice were recorded performing two tasks. In the Delayed Non-Match to Place (DNMP) task (Griffin et al., 2007), mice alternated between Study and Test trials. On Study trials, mice were placed in the center stem in front of the delay barrier, ran to the choice point, where a removable barrier forced them to take a path down one return arm where they received a reward of one chocolate sprinkle. They then moved to the delay area, waited through a 20-second delay, and the delay barrier was lifted to start the Test trial. On a test trial, mice again ran to the choice point but there was no barrier and mice had to go down the return arm opposite to the preceding study trial in order to receive a reward. They then moved to the delay area, from which they were removed to their home cage to wait through a 15–25 second inter-trial interval while the next Study trial was prepared. Mice completed between 25 and 40 Study-Test trial pairs per session.

A second task, termed the Forced-Free task, was used on other days for a different study question not addressed here. On each trial in the Forced-Free task, mice were placed in front of the delay barrier, proceeded to the choice point and were either forced down a particular return arm or were free to choose which arm. On all trials mice received a reward regardless of which arm they entered. After consuming the reward, mice entered the delay area and were immediately returned to their home cage for a 15–25 second inter-trial interval while the next trial was prepared. Mice typically completed 40 trials per session. Forced and free trials were pseudo-randomly interleaved, as was turn direction on forced trials.

The full recording sequence involved blocks of DNMP days interleaved with Forced-Free sessions and 0–2 day breaks. The full sequence is as follows: FF-D-D-FF, break, FF-D-D-

D-FF, break, FF-D-D-D-D-FF. In this manuscript, we present data from DNMP recording sessions when mice made the correct turn direction on 70% of test trials (opposite of preceding study trial). Only correct Test trials and their preceding Study trials are included. Additional sessions were excluded where cell registration could not be performed.

Imaging

Imaging data were acquired using a commercially available miniaturized head-mounted epifluorescence microscope (Inscopix). Microscopes were attached on awake, restrained mice, and optical focus, LED gain and intensity adjusted for each individual mouse but kept stable across days. Videos were captured at 20 Hz with a resolution of 1440×1080 pixels, spatially downsampled 2x to 720×540 pixels. Dropped and corrupted frames were replaced with the preceding good frame, and lost frames were excluded from analysis. Mosaic (Inscopix) was used to pre-process recordings for motion correction and cropping (exclude pixels without GCaMP6f activity), and to generate a minimum projection of the final video (image which has the same height and width of each frame and each pixel is the minimum of that pixel for the entire video) to be used during ROI extraction.

To extract neuron regions of interest (ROIs) and calcium event times, pre-processed videos were then passed through custom-made MATLAB-based image segmentation software (Mau et al., 2018; Kinsky et al., 2018) (TENASPIS, software available at https://github.com/ SharpWave/TENASPIS; see D.W. Sullivan et al., 2017, Soc. Neurosci., abstract). Briefly, TENASPIS applies an adaptive thresholding process on a frame-by-frame basis to a band-pass filtered video to identify discrete regions of fluorescent activity (blobs). Blobs are then identified as likely cells based on expected shape and size, and the software aligns these blobs together over successive frames. Dynamics in calcium activity, including event duration, distance traveled over successive frames, and probable spatial origin, are used to identify putative neuron ROIs. Fluorescence of neuron ROIs is refined into events based on the rising phase of calcium activity. Neuron ROIs with significant spatial overlap and high correlations in calcium activity are merged into single cells.

This neuron ROI and calcium transient event detection algorithm does not exclude overlapping cells, and includes steps to isolate events to individual cells. The image segmentation step is designed to minimize type 1 errors in calcium transient detection, but this comes with a tradeoff of increased type 2 errors (i.e., increased probable transient event rejection at the expense of false event detection). We correct these type 2 errors by summing pixel intensity within the ROIs over time to create fluorescence traces, and then detect missed calcium transients from sharp peaks in the traces. Because a calcium transient in a single neuron often causes a peak in the fluorescence trace of any other overlapping ROIs, the correct ROI origin of any peak in the fluorescence trace must be determined. To achieve this, we calculate averaged pattern of pixel intensities during a given detected fluorescence peak, and calculate the correlation of those intensities to the average pattern of intensities of segmentation-detected transients of all ROIs that overlap with the ROI in question. The ROI that produces the highest correlation (Spearman rho) is then considered to be the correct origin of that calcium transient. See example in Supplementary Figure 1b.

Cells were registered across sessions using a semi-automated procedure with custom software developed in MATLAB that is available along with the rest of our analysis code. For each animal, each session was first aligned to the same 'base' session, selected from the middle of the recording schedule. To align sessions, a set of 25-40 'Anchor' cells was chosen based on the relative positions of neuron ROIs in the base session and each other session (Supplementary Figure 1a-b). Centers of these 'anchor' cells were used to compute an affine geometric transformation ('fitgeotrans' function in MATLAB) and then align the entire set of ROIs in the sessions being registered with the base session ('transformpointsforward' function in MATLAB). Cells with centers within 3um (translated to pixels) were identified as the same cell, and when there was more than one match within that radius, the registered cell with the higher spatial correlation to the base cell was chosen (Supplementary Figure 1c). Cells from a registered session that were not partnered to the base session were added to the set of unique footprints alongside base session cells so that cells in successively registered sessions could be paired to them in turn. Alignment maps were validated by visual inspection: this included looking at the relative alignment with other cells in the field of view, and orientation of putatively mapped cells across sessions. Cells that were not aligned by the automated procedure based on center-to-center distance but that shared orientation and relative alignment to neighboring cells were registered manually (Supplementary Figure 1e, green cell). When looking at the relationship for all cell pairs across all sessions, the correlation of ROIs and distances between centers formed a cluster near the top of the distribution for all cell pairs (Supplementary Figure 1d). The TENASPIS algorithm is designed to discriminate between partially overlapping cells, which gives rise to in many pairs of cells that have high ROI correlations and low center-to-center distances, but remain unregistered because a better matched pair was found using the procedures above; in Supplementary Figure 1d, this manifests in the black points mixed in among the red registered cell pairs.

Behavioral Tracking

Animal position was recorded using an overhead video camera and CinePlex V2 tracking software (Plexon). Tracking was performed at 30 Hz, and was synchronized with a TTL pulse to the imaging data acquisition through nVista software. Tracking was validated manually and errors were corrected using custom software written in MATLAB. Position was then interpolated to the 20 Hz imaging time stamps. We did not filter our data by the animal's running speed, but the animals moved consistently through the central stem and we did not observe incidents where the animal would entirely stop moving or reverse direction while traversing the central stem.

Histology

Mice were perfused transcardially with 10% phosphate buffered saline until outflow ran clear and then with 10% phosphate buffered formalin. Brains were then extracted and post-fixed in formalin for 2–4 days, and then transferred to 30% sucrose solution in phosphate buffered saline for 1–2 days. Brains were then frozen and sliced into 40 um sections on a cryostat (Leica CM 3050S), mounted, and coverslipped with Vectashield Hardset mounting medium with DAPI (Vector Laboratories). Slides were then imaged using a Nikon Eclipse

Ni-E epifluorescence microscope at 10x and 20x to verify viral expression and location and GRIN lens location relative to the CA1 cell layer.

Quantification and Statistical Analysis

Event likelihood—Calcium events were detected and analyzed to compute the likelihood of calcium events occurring at a given location. The analysis software, TENASPIS, (see above) defines an event as the time during the rising phase of a spike in calcium fluorescence in a cell which exceeds a local threshold of that cell's session average of fluorescence activity. This returns a binary output for each cell which describes whether that cell was or was not, at every imaging frame, exhibiting a calcium event. We calculated event likelihood by pooling data from the set of trials of interest for each cell (e.g., Study trials on the stem), and then, for each spatial bin, dividing the number of frames for which an event was occurring by the number of frames when the mouse was in that bin in that set of trials. This produces an output between 0 (an event never occurred in that spatial bin) and 1 (an event always occurred when the mouse was in that spatial bin).

Reliably Active/Included Cells

For single unit analyses, cells are included on a given day when they exhibited a calcium event on at least 25% of trials in a single trial type (e.g. Study-Left) (Supplementary Figure 2). These criteria were chosen before analysis was conducted based on intuition that some reliability criteria would be necessary to exclude neurons which might pass certain statistical criteria (e.g., splitter identification via permutation test, see below) but which might be unconvincing given observation of a neuron's single-trial calcium activity (see example cell in Supplementary Figure 2d). An analysis of the choice of this effect on results can be found in Supplementary Figure 4. In the population analyses, we included all cells that were successfully registered to the sessions being compared.

Splitter Identification

Splitter neurons are cells that exhibit a significant bias in their firing activity on the central stem for trials of a particular upcoming turn direction (Left versus Right) or task phase (Study versus Test) (Figure 2). Thus, each cell is a member of one of four mutually exclusive categories, depending on whether its calcium activity is modulated by either task dimension, both, or neither: turn splitter neuron, task phase splitter neuron, turn+phase splitter neuron, or non-splitter. Note that turn+phase splitter neurons refer to cells splitting both turn direction and task phase.

To identify whether each cell's activity was significantly modulated by task variables, we used a permutation test to measure the significance of the difference in event activity likelihood against a shuffled distribution. This was repeated separately to measure activity bias for turn direction or task phase. We first separated epochs when the mouse ran through the central stem according to the given task dimension (i.e. left and right turn trials, or study and test trials), and computed the event likelihood (see above) for these sets of trials. Then took the difference in likelihood scores by subtracting the Right trial event likelihood in each spatial bin from that for Left trials, or Test trial from Study. We then repeated this for all 1000 sets of shuffled trials, which were generated by shuffling the trials between trial types

accordingly, to get a shuffled difference distribution. Cells were determined to "split" the dimension of interest if their original event likelihood difference was greater than 95% of the shuffle differences in any spatial bin.

In the supplemental data, this procedure was repeated in the same fashion for epochs when the mouse ran down the return arms to measure selectivity for the separate (Right or Left) return arms and for Study and Test task phases while on the return arms.

Population Vector Correlations

Population vector correlations were computed in a manner similar to that described by Leutgeb et al. (2005) (Figure 3a). We generated three sets of correlations: 1) withincondition: trials of the same type (e.g. Study-Left vs. Study-Left); 2) Left vs Right, and 3) Study vs. Test. First, trials were grouped for the comparison of interest and then each group was split so that within condition comparisons would have the same number of trials as the other two correlations. For a given half-set of trials, we computed the event likelihood in each spatial bin with the method described above. We then took these spatial bin event likelihoods for the set of cells included and computed a Spearman correlation for each spatial bin against the event likelihoods in the same spatial bin for the trials in the different comparisons listed above. For correlations computed across days, we computed all day-pair combinations for each self-comparison and for each comparison between study and test trials and between left and right turn trials, for example between left turn trials on day 1 and right turn trials on day 4. Cells included were those present (successfully registered) on both days for each comparison (Similar results were achieved using several other cell inclusion criteria, data not shown).

Statistics

All statistical tests were done with Spearman rank correlations, Wilcoxon rank-sum tests (Mann-Whitney U tests), Wilcoxon signed-rank tests, sign tests, or permutation tests with threshold set at >95% of shuffles for the given test. These tests were used because data were often not normally distributed. Statistics for results in individual animals can be found in Supplementary Table 5.

Results

Heterogeneous changes in daily distribution of single-cell task-related responses

We recorded calcium activity in neurons in dorsal area CA1 as mice performed a delayed non-match to place (DNMP) task over several days. In the DNMP task, mice first run a Study trial where they are forced to turn into one side arm to receive reward. After a 20second delay, mice begin the Test phase and must choose to go down the opposite arm to receive a reward (Figure 1a). We used this task because mice traverse the same section of the maze (the central stem, Figure 1d purple) under each combination of the current Task Phase and upcoming Turn Direction. This allows us to examine hippocampal representations of the same space under four different behavioral conditions: Study-Left, Study-Right, Test-Left, Test-Right. Performance in this task is measured by the number of Test trials on which the mouse correctly chooses the side arm opposite that from the immediately preceding Study

trial; we include only sessions where the mouse made 70% correct alternations (38 sessions in 4 male mice: 9 days in 3 mice, 11 days in 1, spanning up to 18 calendar days, Figure 1b). 4 sessions were excluded for performance below 70%. Performance did not change over the experiment (only days above threshold: rho=-0.031, p=0.852; all days recorded: rho=0.198, p=0.210; Spearman rank correlation).

We recorded activity using the virally-delivered fluorescent calcium indicator GCaMP6f and head-mounted miniature microscopes (Figure 1c), and extracted 8256 unique cell ROIs, cumulative from all sessions in all animals, using custom software (example ROIs in Figure 1e–f, bottom; see Methods) (Kinsky et al., 2018; Mau et al., 2018) (see also Supplementary Figure 1a–c). The number of cells found in each animal stayed consistent over the course of recordings, ranging from 500–1600 (Supplementary Figure 2a). On average, each cell was successfully registered for 3.45 sessions (Supplementary Figure 1f–j, 2b), and cells often displayed stable activity profiles across sessions (see examples in Figure 1e–f).

Single cells often modulate their spatial firing activity according to context-dependent task dimensions such as upcoming turn direction or current task phase. Turn direction responses are thought to represent specific spatial trajectories (Frank et al. 2000; Wood et al. 2000; Ferbinteanu and Shapiro 2003), while a task phase-modulated response profile reflects the (presumably) different network activity states for encoding during the study phase and retrieval during the test phase (Griffin et al. 2007). We assessed whether task variables were encoded by neurons in our recordings by measuring the bias of calcium events towards one task variable using a permutation test which shuffled trial types (see Methods).

We first present the data from all neurons found in our imaging: in each recording session, an average of 34.52±0.47% of cells exhibited at least one calcium event on the central stem during the full set of passes through the central stem, and this proportion did not change over the course of recordings (rho=0.200, p=0.228; Spearman rank correlation of percent cells with any transients vs. recording day number) (Figure 2b). Many neurons passed the permutation test for splitting activity, displaying a functional phenotype described by a modulation of their calcium activity according to the animal's upcoming turn direction (turn splitter neurons), the current task phase (phase splitter neurons), or both (turn+phase splitter neurons) (see examples in Figure 2a). Over all recording sessions, 7,403 out of 11,281 observations of neurons (each recording day treated as a separate observation) which displayed at least one transient were found to have their activity modulated by one or both of the task variables. On each recording day, an average of 14.77±0.91% of neurons active on the stem were turn splitters, significantly less than either the daily mean proportion of task phase splitters at 24.21±0.95% (z=4.59, p=4.333e-06; Wilcoxon signed-rank test) or turn +phase splitters at 26.91±1.39% (z=4.971 p=4.663e-07), which did not differ in proportion from each other (z=1.380, p=0.168); the remaining average of $34.11\pm1.66\%$ did not split either task variable (Non-splitters), and this proportion was significantly greater than both that of phase splitters (z=4.068, p=4.774e-05) and turn+phase splitters (z=2.18, p=0.29) (Figure 2c). Over the course of recordings, these proportions did not change (Turn: rho=0.086, p=0.608; Phase: rho=0.130, 0.438; Turn+phase: rho=-0.183, p=0.272; Nonsplitter: rho=0.036, p=0.832; Spearman rank correlation on percentage of splitter neurons against recording day number) (Figure 2d). These categories of splitter neurons are mutually

exclusive. Note that many cells which display a turn direction-modulated response on Study trials, suggesting that mice could see the turn barrier before having reached it.

When visually inspecting raster plots of calcium activity, we observed many neurons which exceeded the 95% confidence interval determined by shuffling activity in the permutation test but whose activity appeared too unreliable to satisfy our intuition for being a reliable splitter neuron; this was often a result of circumstances related to the low sampling rate for calcium imaging and variable animal behavior. We did not apply a speed threshold to our data, but did not observe trials when the mice entirely stopped moving or reversed their direction while running towards the choice point. Exceptionally long trials can have outsize effects on calcium event likelihoods, and can especially affect neurons that have low activity rates and do not fire on exceptionally long trials (see example in Supplementary Figure 2d). To address these potential issues, we repeated the analyses above with an activity threshold which only included cells that were active on at least 25% of trials of one trial type (reliably active cells).

Across all DNMP recordings, an average of 11.23±0.74% of cells found were above the reliable activity threshold on the stem, and this increased over the course of recordings (rho=0.437, p=0.006; Spearman rank correlation) (Figure 2e). Across recordings, ~90% of reliably active cells displayed a functional phenotype described by a modulation of their calcium activity according to one or both task variables (3,334 passed permutation test/3,654 reliably active on the stem across all recording days, where each day is a separate observation). Among reliably active neurons, there was no difference in the percentages of turn or phase splitter neurons (19.66±1.29% and 18.30±1.20%, respectively, z=1.056, p=0.291, Wilcoxon signed-rank test), but there were more turn+phase splitter neurons than either group (52.72±1.87%, vs. turn: z=5.315, p=1.066e-07, vs. phase: z=5.289, p=1.238e-07) (Figure 2f). We also observed a location bias among different splitting phenotypes of single cells: phase splitter neurons were more likely to have their activity center of mass (event activity pooled across all trial types) closer to the start of the stem (Bin 1) than did turn splitter neurons (p=3.313e-31, Mann-Whitney U test) (Figure 2g). A bias in firing location may indicate that cells tend to fire in proximity to features of behavioral relevance: for phase splitters, this could be whether the trial began in the delay area or by being placed on the maze by the experimenter, while turn splitters encode an upcoming spatial turn direction.

The daily distribution of splitter types among reliably active neurons was not stable: the percentage of turn+phase splitters significantly declined over the course of the experiment (rho=-0.369, p=0.023, Spearman rank correlation), though it remained greater than other splitter types. Meanwhile, the percentage of turn splitter neurons went up (rho=0.357, p=0.028) and there was no change in the percentages of phase splitter neurons (rho=0.120, p=0.472) and non-splitter neurons (rho=0.267, p=0.105) (Figure 2h). The percentages of each type of splitter neuron were not correlated with animals' performance on the DNMP task (all rho absolute value <0.276, all p>0.094) (Supplementary Figure 3a). To examine how our activity criterion affected these results, we repeated this analysis across a range of criteria and find that at the majority of thresholds they at least trend in the same direction (Supplementary Figure 4a,c). The observed changes in reliably active splitter cell

distributions (Figure 2h) could result from an increase in calcium activity over time; however, this does not seem to be the case since the average likelihood of calcium events increases over time in turn splitters but not in turn+phase (Supplementary Figure 4e–f).

These results indicate heterogeneous stability within the population, with the entire population maintaining a distribution of task-related activity but exhibiting change in that representation among more reliably active neurons. These findings replicate a previous result showing phase and turn splitters (Griffin et al. 2007) in a new species and extend that work to suggest that the distribution of task-dimension modulated responses among reliably active neurons is unstable over time, even though behavioral output is reliable. In particular, the number of reliably active turn splitter neurons increases over time, whereas the number of reliably active turn+phase splitter neurons decreases over time, suggesting representations in single-neurons become less experience-specific over time.

We applied these same analyses to determine activity modulation according to task variables to neuronal activity during the return arm epochs. Because this analysis is performed in the same way, it can be used to indicate relative distinctiveness in the way neurons code for overlapping spatial trajectories (central stem) as opposed to unique spatial locations (return arms). Many cells displayed a calcium event bias for one arm over the other (place cells, referred to in the text as "place splitters"), and many cells also showed selectivity for one task phase (see Statistics in Supplementary Table 1). The percentages of place and phase splitter neurons on the return arms did not change over time, though there was an increase in the number of cells which were reliably active on the return arms but did not show place or task phase selectivity (non-splitters) (Supplementary Figure 5a-e). Additionally, we repeated these procedures for central-stem activity during the Forced-Free sessions (see Methods). We observed fewer neurons coding for task phase in these sessions than during DNMP sessions, and saw neither a change in the number of splitter neurons over time nor a correlation with spontaneous alternation behavior ("Accuracy," choosing a return arm during a free trial opposite that of the prior trial, even though all choices were rewarded) (Supplementary Figure 5f-i). That we saw a change in distribution of splitter neurons in the DNMP task but not in the Forced-Free task suggests that representations for these tasks do not affect each other. These results also show that changes in the representations of the task and environment among reliably active neurons are modulated by memory load, which is low on the return arms and high in the central stem during the DNMP task.

In summary, by demonstrating that the distribution of task variable responses among single units is unstable, we show that representations for various task dimensions experienced in the same spatial location and during similar behaviors (stem traversal) are heterogeneously stable, with divergent changes based on their coding of the behavioral context.

Population-level separation of task dimensions is stable over experience

We next asked how these patterns of activity manifested in the activity state of CA1 as a whole. This population analysis was designed to measure the similarity in the pattern of activity among the population of neurons within and across recording sessions. We computed Spearman correlations for the calcium event likelihood in each spatial bin from the start of the stem to the choice point for a given trial type using the calcium event

likelihood for each trial type of all cells present in the session pair (Figure 3a) (see Methods). We generated three sets of correlations: 1) trials of the same turn direction and task phase (within-condition; e.g. Study-Left vs. Study-Left), 2) trials of different turn directions (Left vs. Right, abbreviated as LvR), and 3) trials of different task phases (Study vs. Test, abbreviated as SvT).

We found a stable ensemble activity pattern when examining the population vector correlations for trials occurring on the same day. Activity states for trials of the same type were significantly more correlated than those for trials of different direction and for trials of different task phase, thus showing a discrimination in the ensemble-level code for different trial types (see Supplementary Table 2 for detailed statistics). As shown in Figure 3b, the correlations between trials of the same type did not change across spatial bins (rho=0.045, p=0.116; Spearman rank correlation). In contrast, activity states for left and right trials grew more decorrelated as animals approached the choice point (rho=-0.678, p=4.946e-83), and study and test trials were most discriminable (less correlated) at the start of the stem (rho=0.332, p=4.418e-17). The correlation change along the stem follows the center-of-mass distribution for splitter cell firing fields (Figure 2f). This pattern of correlations across spatial bins was stable over the course of recordings (all rho absolute value < 0.313, all p > 0.056; Spearman rank correlation of 2-bin mean for each type of population vector correlation value against recording day number) (Examples for bins 1-2 and 7-8 in Figure 3c-d). This result demonstrates that, in spite of the changing distribution of single-neuron encoding properties (Figure 1e), the population-level distinction between activity states (Figure 3b) and its relationship to spatial position is stable over time (Figure 3c–d). We also evaluated the relationship between population vector correlation and animal behavioral performance, and only found a relationship between increasing performance and increased correlation between left and right trials for bins 1-2 (rho=0.350, p=0.032, Spearman rank correlation; Supplementary Figure 3b). This stability of representations at the population level mirrors that observed in the population as a whole without thresholding for reliably active neurons (Figure 2a-d).

We next assessed the correlations within and between trial types for trials on different days. It may be expected that population activity states would diverge with respect to time (i.e., become less correlated) due to cell replacement and changes in the splitter neuron distribution (Figure 2e). To assess this, we examined the mean population vector correlations at the beginning and end of the stem between sessions recorded 1 to 16 days apart. We observed that all three types of correlations significantly decreased with increasing day lag at both ends of the stem, (Figure 3e–f). However, even as correlations decreased, LvR and SvT correlations were significantly lower than those between trials of the same type for at least a week between sessions and in many cases longer (see detailed statistics in Supplementary Table 3,4). These results show that constant cell turnover minimally impacts the ability of the population to represent different experiences of the same space over many days of recording and that this representational structure is preserved over time. However, the extent to which the neuronal population distinguishes between task dimensions depends on the dimensions being compared, the animal's physical location, and the temporal lag between experiences.

We next assessed the origin of the changes in the distribution of splitter neuron types over time. There are several possible sources of change in the splitter neuron distribution: different splitter neuron types could be persistently active for different amounts of time before becoming silent (variable stability); neurons could change their splitter type (splitter type transition); or previously silent neurons could be preferentially allocated to certain splitter types (unequal allocation of newly active cells). We found no evidence of variable stability: cells were equally likely to be reactivated in later recording days regardless of splitting type (all p>0.05, Wilcoxon rank-sum test between each pair of splitting phenotypes at each day lag) (Figure 4a).

We next tracked the history of all cells to determine the origin or "source" of each splitter neuron in the preceding session. For each splitter neuron from the second included session onwards, we tracked whether that cell was a splitter neuron of any type in the preceding DNMP session or was inactive (neurons below the reliable activity threshold or undetected by our ROI extraction algorithm; note that this includes cells which may have been active 2 or more sessions before the session of interest). We found that previously inactive cells were the largest source category to all types of splitter neurons in 89.26% of recording sessions, and contributed an average of 57.10% of splitter neurons per session (Figure 4b). Turn +phase splitter neurons were the second largest source category to splitter neurons of all types, contributing on average 21.30% of splitter neurons. The percentage of active cells which had been inactive on the prior recording day did not change over time (Figure 4c) (rho=-0.025, p=0.891, Spearman rank correlation). In addition to showing the immediate integration of newly active cells into the coding population, this result suggests that representation of task variables in single units becomes less specific over time, where cells become less likely to encode both task phase and turn direction.

The above result on splitter neuron sources suggests that changes in the distribution of single unit responses are, to a large degree, driven by the splitting type a newly active neuron assumes as opposed to transitions between different splitting types. Indeed, the percentage of splitter types of newly active cells closely matched the distribution of splitter types overall: new cells were more likely to become turn+phase splitter neurons rather than turn-only or phase-only splitter neurons (Turn+phase vs. Turn: z=5.069, p=4.000e-07; Turn+phase vs. Phase: z=5.001, p=5.709e-07; Wilcoxon signed-rank test) (Figure 4d). Additionally, the changes in this distribution of newly active cells over the course of recordings closely matched those observed for all splitter neurons (Figure 2g): while newly active cells on all days were more likely to be turn+phase splitter neurons than other types, this likelihood significantly decreased over time (rho=-0.388, p=0.023; Spearman rank correlation), the percentage of new cells allocated to turn splitter neurons on the stem exhibited a non-significant trend towards increasing (rho=0.328, p=0.058), while those for phase splitter neurons and non-splitters were stable (rho=0.161, p=0.362 and rho=0.189, p=0.285 respectively) (Figure 4e).

Splitter and place neurons on the return arms were also found to be equally stable and primarily derived from newly active cells, but the distribution of cells newly active on the

return arms among splitter types did not change over time, again suggesting the redistribution of splitter neurons is related to memory load (Supplementary Figure 6).

These results show that the changing distribution of single unit responses is primarily attributable to changes in the allocation of new cells to encode task variables, rather than unequal stability of different splitter types.

Discussion

We recorded cells in dorsal CA1 of the hippocampus in mice performing a Delayed Non-Match to Place task over several sessions. In tracking the same populations of cells, we found that there was heterogeneity in the stability of task-related representations. Many single cells exhibited context-dependent modulation in their calcium activity while the animal was in the same spatial location, replicating earlier findings demonstrating that hippocampal place cells encode the behavioral context in addition to spatial position (Griffin et al., 2007). Among cells active on the stem, ~65% were sensitive to turn context or phase context, and the distribution of splitting types among active neurons was stable over the course of recordings. However, we also found that this proportion went up to $\sim 90\%$ among neurons crossing a threshold for reliable activity, a proportion typically higher than seen in previous studies but may be attributable to differences in species, recording methods and statistical detection of splitting (see **Appendix 1** for specific details), and that this distribution of context-dependent responses among reliably active neurons was not stable over the course of recordings: the percentage of task phase splitter neurons was stable, the percentage of turn direction splitter neurons increased, and the percentage of turn+phase splitter neurons decreased. Thus heterogeneous stability is suggested both in a distinction between populations of neurons with different activity rates and among the distribution of task-modulated responses among more reliably active neurons. We found the change of splitter distributions among reliably active neurons was not attributable to variable stability of each splitter type, but instead appeared to be due to how cells which became reliably active were allocated to different splitter types. In spite of cell turnover and changes in the representation of task features among single neurons, ensemble-level population representations for different trial types were stably segregated over many recording sessions. These data demonstrate that the hippocampal representation of ongoing experience can undergo reorganization at the single neuron level while minimally impacting populationlevel coding.

The changes that we observed in the distribution of splitter neurons may be influenced by the use of specific methods and thresholds. The results on distribution of task representations among the population of cells are presented as a percentage of cells which met criteria for reliable activity on a particular recording session, so changes in these percentages will be influenced by the number of cells included. Regarding our reliability threshold, this does seem to be a possible factor in our results since more neurons pass this threshold over time (Figure 2e), although we found that this was consistent across a range of reliability thresholds (Supplementary Figure 4a–d). Since any proportion is related to other proportions in the sample, an increase or decrease in proportion could arise either from an increase of the one quantity of interest (here, number of cells which pass a particular test for splitting) or

from a change in the number of elements sampled (here, number of cells included). To estimate possible changes in activity coding over time, future studies should seek to employ methods which can better track individual neurons across recording sessions in a manner that is not activity dependent, such as two-photon imaging, and by methods which more precisely resolve individual action potentials.

Representations may change in different ways over time during stable behavior based on competing demands on memory reorganization. Generalization emphasizes the similarities across experiences to aid in the transfer of learning across contexts, while orthogonalization makes representations more distinct to mitigate interference between contexts. Both mechanisms are important for spatial navigation and episodic memory (Hasselmo & Wyble, 1997; Kumaran & McClelland, 2012; McNaughton & Morris, 1987; Norman & O'Reilly, 2003; Schapiro et al. 2017; Treves & Rolls, 1994; Winocur et al. 2010), and both processes are observed in the hippocampus in fMRI studies using behavioral tasks with multiple demands (Brown et al. 2010; Brown & Stern, 2014; Chanales et al. 2017). However, the interplay of generalization and orthogonalization in the long term reorganization of memory has not been previously studied at the single neuron level in a dynamically evolving neural circuit. Representations of different trial types may become more orthogonalized and distinct, following the precedent set by many studies on learning (Komorowski et al. 2009; McKenzie et al. 2013; Chanales et al. 2017). Alternatively, representations could become more schematic through generalization as the animals become over-trained on the task, perhaps preserving only those distinctions relevant to performing the task. At the single neuron level, we observed a result consistent with the generalization hypothesis among the population of reliably active neurons: a decreasing number of turn+phase splitter neurons (which encode a single experience: a route to a single destination during a single task phase) and an increasing number of turn splitter neurons (which encode multiple experiences: routes to the same destination during multiple task phases). At the population level, however, we instead observed a highly stable representational structure.

Studies which report orthogonalizing change in hippocampal coding properties typically examine an initial learning phase, comparing data from before and after a subject reaches a performance criterion, and often in a single session (Kobayashi et al. 2003; Komorowski, et al. 2009; McKenzie et al. 2013). Because our recordings began after animals had received considerable experience with the maze environment during behavioral shaping, we may have captured a set of operational demands unlike initial task learning. To reconcile our finding of generalization with previous reports of orthogonalization, we propose that both mechanisms act on the organization of hippocampal memory but at different timescales: orthogonalization dominates an early, fast encoding process which emphasizes the uniqueness of current experiences, while generalization acts as a slower refinement of existing memory representations by finding statistical regularities; both of these processes likely involve regions outside the hippocampus (Ghosh & Gilboa, 2013; Koster et al., 2018; Lewis et al., 2018). This distinction suggests that it is more appropriate for our work to be framed in terms of long-term mechanisms of memory stability, rather than those which are relevant to shaping the initial learning and encoding process.

Divergent expectations for short and long-term memory organization are apparent when comparing our results to a previous report which employed a similar task to ours in which human participants navigated partially overlapping trajectories in a virtual environment (Chanales et al., 2017; see also Brown et al. 2010; Brown and Stern 2014). The authors found that the hippocampal voxel activity patterns for overlapping trajectory segments grew more distinct from each other over the course of learning, while patterns for non-overlapping segments did not change in their representational similarity. Our results parallel this finding in showing that conflicts between behavioral responses in overlapping locations (experienced on the central stem in the DNMP task) can drive changes in the neural representation while representations for non-overlapping segments remain stable (return arms, Supplementary Figure 5). However, unlike Chanales and colleagues, we did not observe a population-level increase in discriminability of overlapping segments, which could be explained by the fact that their study was conducted in a single session while ours ran for multiple weeks.

Prior studies have attributed a working memory role to the hippocampus in DNMP and other alternation tasks. Working memory accounts propose that on short, behaviorally relevant timescales the hippocampus maintains a representation of the previous trial to inform future behavior. This interpretation was prompted by findings that hippocampal lesions produce performance deficits in alternation tasks which involve a delay (Hampson et al. 1999; Dudchenko et al. 2000) and by correspondence between during delay period neural activity and upcoming turn directions (Deadwyler et al. 1996). However, alternation tasks cannot distinguish between prospective and retrospective coding (see Frank et al. 2000 and Ferbinteanu & Shapiro 2003), meaning delay and central stem activity could represent a previous trial or upcoming trajectory.

We suggest instead that continued involvement of the hippocampus in distinctly representing overlapping spatial trajectories may be appropriate for self-localization within an existing spatial memory map (Redish & Touretztky, 1998). It was previously assumed that task splitter neurons reflected respective encoding and retrieval demands for Study and Test trials (Griffin et al. 2007); the self-localization interpretation suggests instead that task phase splitters instead encode immediate history of the stem traversal, whether the current trial began by being placed in the maze by the experimenter (Study) or being released from the delay area (Test). Self-localization assumes neither that the animals are sensitive to our conception of the task nor that encoding and retrieval "modes" be expressed as measurably different patterns of activity in CA1. The lack of neurons that code exclusively for Task Phase on the return arms (Supplementary Figure 5), where the trial-start behavioral cue is less salient, is consistent with this hypothesis. The strictest interpretation of task phase splitting as self-localization suggests it acts as a code to distinguish slightly different routes to the same reward destination (Grieves et al. 2016). Task phase splitting (Figure 2) and delay period splitting (Deadwyler et al. 1996) could together contribute to self-localization within a cognitive map of the task that links longer sequences of events through the maze, wherein overlapping trajectories begin on the central stem, pass down one side arm, linger in the delay area, and then pass again through the stem and onto the other side arm (Hasselmo, 2008). Task phase splitting on the central stem is similar to many other findings of contextdependent place-cell activity (Ferbinteanu & Shapiro, 2003; Frank et al., 2000; Hasselmo,

2008; Sun, Yang, Martin, & Tonegawa, 2019). Disambiguating the working-memory and self-localization accounts of splitter neuron activity will require designing tasks that use behavioral and spatial cues that are highly consistent across distinct but overlapping behaviors.

We observed that neurons which became active (above the reliability criterion) over the course of the study often displayed a high activity level and task-modulated activity patterns even on the first day that they were apparent in our recordings. The immediate integration of newly active cells suggests an attractor-like mechanism, constrained by the activity and connectivity within an area and its inputs, which defines the "activity coding space" that new cells are likely to function within, manifested as a low-dimensional manifold that remains stable despite cell turnover over days and trial-to-trial variability of activity over minutes (Low et al., 2018; Rubin et al., 2015). In the context of this study where no new learning occurred, we suggest that new cells are activated by intra-cell factors that raise a cell's excitability, such as the level of CREB (Han et al., 2007; Yiu et al., 2014; Zhou et al., 2009), coupled with changes in inputs to CA1 neurons from upstream contacts. Higher excitability and raised input activity together would cause activation via previously silent synapses to drive new CA1 neurons with task-modulated activity. This would result in LTP of synaptic inputs that would enable these new CA1 neurons to remain active for many days. Drift incidentally coupled across regions thus preserves task representations in low dimensional activity states and over time causes the activity in a given structure to reflect the dominant coding scheme of its upstream partners, preventing a wholesale disconnection between associated regions. New hippocampal learning, the highly orthogonal and relational explicit recombination and extraction of features from upstream activity patterns, would not occur on this timescale and could instead be driven at short timescales by mismatched activity patterns (Hasselmo and Schnell 1994; Hasselmo and Wyble 1997; Hasselmo and Eichenbaum 2005; Lisman and Otmakhova 2001; Hasselmo 2005) and signals such as dopaminergic input (Kempadoo et al., 2016; Kentros et al., 2004; Takeuchi et al., 2016) or cholinergic input (Hasselmo & Schnell, 1994; Hasselmo & Wyble, 1997).

To address our results specifically, we observed that reliably active single neurons encoded fewer task variables over time. Lipton et al. (2007) previously found that there were more turn-direction splitter neurons in the medial entorhinal cortex (MEC) than in CA1: if this is additionally true in the DNMP task and there are more neurons coding for task phase in the MEC than in CA1, the hippocampus may over long periods of time emphasize those dimensions in the activity of single neurons, slowly "undoing" the highly orthogonal coding hypothesized to occur during episodic memory encoding (Alme et al. 2014; McClelland et al., 1995; Hasselmo and Wyble 1997). These hypotheses could be tested with computational models looking at task dimensionality coding in connected areas drifting at different rates, and in future experiments employing simultaneous recordings in HPC and MEC during the DNMP task.

Our results here show that the stability of hippocampal representations is heterogeneous, displaying different rates of change in task-relevant activity across activity levels, cognitive demands, maze locations, and levels of analysis. These changes are largely attributable to changes in the allocation of newly reliably active cells among task-modulated activity types,

as well as individual cells' transitioning from coding both task dimensions to just coding for one. Together, the results suggest that reorganization of memory representations actively reshapes hippocampal memories among single neurons but not at the population level. Future studies should seek to clarify the behavioral parameters which predict the rate of cell replacement, the allocation of newly reliably active cells, and the cellular and network mechanisms which mediate them.

Software and Data availability

Software used in our analysis is freely available on GitHub. TENASPIS is available at https://github.com/SharpWave/TENASPIS, and all other analysis software is available at https://github.com/samjlevy/CaImageRelated. Data can be made available from the authors upon reasonable request.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments:

This work was supported by the following grants: U.S. Office of Naval Research MURI N00014-16-1-2832 and N00014-19-1-2571, National Institutes of Health R01 MH052090, R01 MH 051570, and MH060013, and MH120073. We would like to thank Howard Eichenbaum, who initially conceived of the experimental questions and study design and secured funding to perform them, but passed away in July of 2017 before the analysis had progressed significantly. We additionally thank Ian Davison for providing extensive consultation during data analysis and manuscript preparation, Jay Bladon, Andy Alexander, Dan Sheehan, and Catherine Mikkelsen for input on results, analysis and interpretation, Elin Anwar and Alex Terzibachian for help with data processing, and Dan Orlin, Wing Ning, Denise Parisi, Shelley Russek, and Sandra-Jean Grasso for administrative support. We would like to acknowledge the GENIE Program, specifically Vivek Jayaraman, PhD, Douglas S. Kim, PhD, Loren L. Looger, PhD, Karel Svoboda, PhD from the GENIE Project, Janelia Research Campus, Howard Hughes Medical Institute, for providing the GCaMP6f virus. Finally, we would like to acknowledge Inscopix, Inc. for making single-photon calcium imaging miniscopes widely available, and specifically Lara Cardy and Vardhan Dani for all their technical support throughout the experiment.

References

- Alme CB, Miao C, Jezek K, Treves A, Moser EI, & Moser M-B (2014). Place cells in the hippocampus: Eleven maps for eleven rooms. Proceedings of the National Academy of Sciences, 111(52), 18428–18435. 10.1073/pnas.1421056111
- Attardo A, Fitzgerald JE, & Schnitzer MJ (2015). Impermanence of dendritic spines in live adult CA1 hippocampus. Nature, 523(7562), 592–596. 10.1038/nature14467 [PubMed: 26098371]
- Barnes CA, Suster MS, Shen J, & Mcnaughton BL (1997). Multistability of cognitive maps in the hippocampus of old rats. Nature, 388(August 1996), 272–275. [PubMed: 9230435]
- Brown TI, Ross RS, Keller JB, Hasselmo ME, & Stern CE (2010). Which way was I going? Contextual retrieval supports the disambiguation of well learned overlapping navigational routes. Journal of Neuroscience, 30(21), 7414–7422. 10.1523/JNEUROSCI.6021-09.2010 [PubMed: 20505108]
- Brown TI, & Stern CE (2014). Contributions of medial temporal lobe and striatal memory systems to learning and retrieving overlapping spatial memories. Cerebral Cortex, 24(7), 1906–1922. 10.1093/ cercor/bht041 [PubMed: 23448868]
- Chambers AR, & Rumpel S (2017). A stable brain from unstable components: Emerging concepts and implications for neural computation. Neuroscience, 357, 172–184. 10.1016/ j.neuroscience.2017.06.005 [PubMed: 28602920]

- Chanales AJH, Oza A, Favila SE, & Kuhl BA (2017). Overlap among Spatial Memories Triggers Repulsion of Hippocampal Representations. Current Biology, 27(15), 2307–2317.e5. 10.1016/ j.cub.2017.06.057 [PubMed: 28736170]
- De Paola V, Holtmaat A, Knott G, Song S, Wilbrecht L, Caroni P, & Svoboda K (2006). Cell typespecific structural plasticity of axonal branches and boutons in the adult neocortex. Neuron, 49(6), 861–875. 10.1016/j.neuron.2006.02.017 [PubMed: 16543134]
- Deadwyler S. a, Bunn T, & Hampson RE (1996). Hippocampal ensemble activity during spatial delayed-nonmatch-to-sample performance in rats. The Journal of Neuroscience : The Official Journal of the Society for Neuroscience, 16(1), 354–372. [PubMed: 8613802]
- Dudchenko PA, Wood ER, & Eichenbaum H (2000). Neurotoxic hippocampal lesions have no effect on odor span and little effect on odor recognition memory but produce significant impairments on spatial span, recognition, and alternation. The Journal of Neuroscience : The Official Journal of the Society for Neuroscience, 20(8), 2964–2977. Retrieved from http://www.ncbi.nlm.nih.gov/ pubmed/10751449 [PubMed: 10751449]
- Ferbinteanu J, & Shapiro ML (2003). Prospective and retrospective memory coding in the hippocampus. Neuron, 40(6), 1227–1239. 10.1016/S0896-6273(03)00752-9 [PubMed: 14687555]
- Frank LM, Brown EN, & Wilson M (2000). Trajectory encoding in the hippocampus and entorhinal cortex. Neuron, 27(1), 169–178. 10.1016/S0896-6273(00)00018-0 [PubMed: 10939340]
- Ghosh VE, & Gilboa A (2013). What is a memory schema? A historical perspective on current neuroscience literature. Neuropsychologia, 53, 104–114. 10.1016/j.neuropsychologia.2013.11.010 [PubMed: 24280650]
- Grieves RM, Wood ER, & Dudchenko PA (2016). Place cells on a maze encode routes rather than destinations. ELife, 1–24. 10.7554/eLife.15986
- Griffin AL, Eichenbaum H, & Hasselmo ME (2007). Spatial Representations of Hippocampal CA1 Neurons Are Modulated by Behavioral Context in a Hippocampus-Dependent Memory Task. Journal of Neuroscience, 27(9), 2416–2423. 10.1523/jneurosci.4083-06.2007 [PubMed: 17329440]
- Grutzendler J, Kasthuri N, & Gan W (2002). Long-term spine dendritic spine stability in the adult cortex. Letters to Nature, 420(December), 812–816. 10.1038/nature01151.1.
- Hampson RE, Jarrard LE, & Deadwyler SA (1999). Effects of Ibotenate Hippocampal and Extrahippocampal Destruction on Delayed-Match and -Nonmatch-to-Sample Behavior in Rats. The Journal of Neuroscience, 19(4), 1492–1507. 10.1523/jneurosci.19-04-01492.1999 [PubMed: 9952425]
- Han J, Kushner S. a, Yiu AP, Cole CJ, Matynia A, Brown R. a, ... Josselyn S. a. (2007). During Memory Formation. Science, 316(April), 457–460. 10.1126/science.1139438 [PubMed: 17446403]
- Hasselmo ME (2005). The role of hippocampal regions CA3 and CA1 in matching entorhinal input with retrieval of associations between objects and context: Theoretical comment on Lee et al. (2005). Behavioral Neuroscience, 119(1), 342–345. 10.1037/0735-7044.119.1.342 [PubMed: 15727540]
- Hasselmo ME (2008). Grid cell mechanisms and function: Contributions of entorhinal persistent spiking and phase resetting. Hippocampus, 18(12), 1213–1229. 10.1002/hipo.20512.Grid [PubMed: 19021258]
- Hasselmo ME, & Eichenbaum Howard. (2005). Hippocampal mechanisms for the context-dependent retrieval of episodes. Neural Networks, 18(9), 1172–1190. 10.1016/j.neunet.2005.08.007 [PubMed: 16263240]
- Hasselmo ME, & Schnell E (1994). Laminar selectivity of the cholinergic suppression of synaptic transmission in rat hippocampal region CA1: Computational modeling and brain slice physiology. Journal of Neuroscience, 14(6), 3898–3914. 10.1523/jneurosci.14-06-03898.1994 [PubMed: 8207494]
- Hasselmo ME, & Wyble BP (1997). Free recall and recognition in a network model of the hippocampus: simulating effects of scopolamine on human memory function. Behavioural Brain Research, 89, 1–34. 10.1080/00268977500102411 [PubMed: 9475612]

- Kempadoo KA, Mosharov EV, Choi SJ, Sulzer D, & Kandel ER (2016). Dopamine release from the locus coeruleus to the dorsal hippocampus promotes spatial learning and memory. Proceedings of the National Academy of Sciences of the United States of America, 113(51), 14835–14840. 10.1073/pnas.1616515114 [PubMed: 27930324]
- Kentros CG, Agnihotri NT, Streater S, Hawkins RD, & Kandel ER (2004). Increased Attention to Spatial Context Increases Both Place Field Stability and Spatial Memory. Neuron, 42(2), 283–295. 10.1016/S0896-6273(04)00192-8 [PubMed: 15091343]
- Kinsky NR, Mau W, Sullivan DW, Levy SJ, Ruesch EA, & Hasselmo ME (2020). Trajectorymodulated hippocampal neurons persist throughout memory-guided navigation. Nature Communications, 11(1), 2443. 10.1038/s41467-020-16226-4
- Kinsky NR, Sullivan DW, Mau W, Hasselmo ME, & Eichenbaum HB (2018). Hippocampal Place Fields Maintain a Coherent and Flexible Map across Long Timescales. Current Biology, 28(22), 3578–3588. 10.1016/j.cub.2018.09.037 [PubMed: 30393037]
- Kobayashi T, Tran AH, Nishijo H, Ono T, & Matsumoto G (2003). Contribution of hippocampal place cell activity to learning and formation of goal-directed navigation in rats. Neuroscience, 117(4), 1025–1035. 10.1016/S0306-4522(02)00700-5 [PubMed: 12654354]
- Komorowski RW, Manns JR, & Eichenbaum H (2009). Robust conjunctive item-place coding by hippocampal neurons parallels learning what happens where. The Journal of Neuroscience : The Official Journal of the Society for Neuroscience, 29(31), 9918–9929. 10.1523/ JNEUROSCI.1378-09.2009 [PubMed: 19657042]
- Koster R, Chadwick MJ, Chen Y, Berron D, Banino A, Düzel E, ... Kumaran D (2018). Big-Loop Recurrence within the Hippocampal System Supports Integration of Information across Episodes. Neuron, 99(6), 1342–1354.e6. 10.1016/j.neuron.2018.08.009 [PubMed: 30236285]
- Kumaran D, & McClelland JL (2012). Generalization through the recurrent interaction of episodic memories: A model of the hippocampal system. Psychological Review, 119(3), 573–616. 10.1037/ a0028681 [PubMed: 22775499]
- Law LM, Bulkin DA, & Smith DM (2016). Slow stabilization of concurrently acquired hippocampal context representations. Hippocampus, 26(12), 1560–1569. 10.1002/hipo.22656 [PubMed: 27650572]
- Lee I, Griffin AL, Zilli EA, Eichenbaum H, & Hasselmo ME (2006). Gradual Translocation of Spatial Correlates of Neuronal Firing in the Hippocampus toward Prospective Reward Locations, 639– 650. 10.1016/j.neuron.2006.06.033
- Lever C, Wills T, Cacucci F, Burgess N, & Keefe JO (2002). Long-term plasticity in hippocampal place-cell representation of environmental geometry. Letters to Nature, 416(March), 236–238. 10.1038/416090a
- Lewis PA, Knoblich G, & Poe G (2018). How Memory Replay in Sleep Boosts Creative Problem-Solving. Trends in Cognitive Sciences, 22(6), 491–503. 10.1016/j.tics.2018.03.009 [PubMed: 29776467]
- Lipton PA, White JA, & Eichenbaum H (2007). Disambiguation of Overlapping Experiences by Neurons in the Medial Entorhinal Cortex, 27(21), 5787–5795. 10.1523/ JNEUROSCI.1063-07.2007
- Lisman JE, & Otmakhova NA (2001). Storage, recall, and novelty detection of sequences by the hippocampus: Elaborating on the SOCRATIC model to account for normal and aberrant effects of dopamine. Hippocampus, 11(5), 551–568. 10.1002/hipo.1071 [PubMed: 11732708]
- Low RJ, Lewallen S, Aronov D, Nevers R, & Tank DW (2018). Probing variability in a cognitive map using manifold inference from neural dynamics. BioRxiv 10.1101/418939
- Mankin EA, Diehl GW, Sparks FT, Leutgeb S, & Leutgeb JK (2015). Hippocampal CA2 Activity Patterns Change over Time to a Larger Extent than between Spatial Contexts. Neuron, 85(1), 190– 201. 10.1016/j.neuron.2014.12.001 [PubMed: 25569350]
- Mankin EA, Sparks FT, Slayyeh B, Robert J, Leutgeb S, & Leutgeb JK (2012). Correction for Mankin et al., Neuronal code for extended time in the hippocampus: Fig. 5.. Proceedings of the National Academy of Sciences, 102(47), 19462–19467. 10.1007/s007660070019
- Mau W, Sullivan DW, Kinsky NR, Hasselmo ME, Howard MW, & Eichenbaum H (2018). The Same Hippocampal CA1 Population Simultaneously Codes Temporal Information over Multiple

Timescales. Current Biology, 28(10), 1499–1508.e4. 10.1016/j.cub.2018.03.051 [PubMed: 29706516]

- McClelland JL, McNaughton BL, & O'Reilly RC (1995). Why There Are Complementary Learning Systems in the Hippocampus and Neocortex: Insights From the Successes and Failures of Connectionist Models of Learning and Memory. Psychological Review, 102(3), 419–457. 10.1037/0033-295X.102.3.419 [PubMed: 7624455]
- McKenzie S, Robinson NTM, Herrera L, Churchill JC, & Eichenbaum H (2013). Learning Causes Reorganization of Neuronal Firing Patterns to Represent Related Experiences within a Hippocampal Schema. Journal of Neuroscience, 33(25), 10243–10256. 10.1523/ JNEUROSCI.0879-13.2013 [PubMed: 23785140]
- McNaughton BL, & Morris RGM (1987). Hippocampal synaptic enhancement and information storage within a distributed memory system. Trends, 10(10), 408–415.
- Mehta MR, Quirk MC, & Wilson MA (2000). Experience-dependent asymmetric shape of hippocampal receptive fields. Neuron, 25(3), 707–715. 10.1016/S0896-6273(00)81072-7 [PubMed: 10774737]
- Norman KA, & O'Reilly RC (2003). Modeling Hippocampal and Neocortical Contributions to Recognition Memory: A Complementary-Learning-Systems Approach. Psychological Review, 110(4), 611–646. 10.1037/0033-295X.110.4.611 [PubMed: 14599236]
- Pfeiffer T, Poll S, Bancelin S, Angibaud J, Inavalli VK, Keppler K, ... Nägerl UV (2018). Chronic 2P-STED imaging reveals high turnover of dendritic spines in the hippocampus in vivo. ELife, 7, 1– 17. 10.7554/elife.34700
- Poe GR, Nitz DA, McNaughton BL, & Barnes CA (2000). Experience-dependent phase-reversal of hippocampal neuron firing during REM sleep. Brain Research, 855(1), 176–180. 10.1016/ S0006-8993(99)02310-0 [PubMed: 10650147]
- Redish AD, & Touretztky DS (1998). The Role of the Hippocampus in the Morris Water Maze. Computational Neuroscience, 10(1), 73–111. 10.1007/978-1-4615-4831-7_17
- Rubin A, Geva N, Sheintuch L, & Ziv Y (2015). Hippocampal ensemble dynamics timestamp events in long-term memory. ELife, 4(DECEMBER2015), 1–16. 10.7554/eLife.12247
- Rule ME, O'Leary T, & Harvey CD (2019). Causes and consequences of representational drift. Current Opinion in Neurobiology, 58, 141–147. 10.1016/j.conb.2019.08.005 [PubMed: 31569062]
- Schapiro AC, Turk-Browne NB, Botvinick MM, & Norman KA (2017). Complementary learning systems within the hippocampus: A neural network modelling approach to reconciling episodic memory with statistical learning. Philosophical Transactions of the Royal Society B: Biological Sciences, 372(1711). 10.1098/rstb.2016.0049
- Sun C, Yang W, Martin J, & Tonegawa S (2019). CA1 pyramidal cells organize an episode by segmented and ordered events. BioRxiv Preprint, 565689. 10.1101/565689
- Takeuchi T, Duszkiewicz AJ, Sonneborn A, Spooner PA, Yamasaki M, Watanabe M, ... Morris RGM (2016). Locus coeruleus and dopaminergic consolidation of everyday memory. Nature, 537(7620), 357–362. 10.1038/nature19325 [PubMed: 27602521]
- Taxidis J, Pnevmatikakis E, Mylavarapu AL, Arora JS, Samadian KD, Hoffberg EA, & Golshani P (2018). Emergence of stable sensory and dynamic temporal representations in the hippocampus during working memory. BioRxiv, 474510. 10.1101/474510
- Thompson LT, & Best PJ (1990). Long-term stability of the place-field activity of single units recorded from the dorsal hippocampus of freely behaving rats. Brain Research, 509(2), 299–308. 10.1016/0006-8993(90)90555-P [PubMed: 2322825]
- Treves A, & Rolls ET (1994). Computational Analysis of the role of hippocampus in memor. Hippocampus, 4(3), 374–391. [PubMed: 7842058]
- Winocur G, Moscovitch M, & Bontempi B (2010). Memory formation and long-term retention in humans and animals: Convergence towards a transformation account of hippocampal-neocortical interactions. Neuropsychologia, 48(8), 2339–2356. 10.1016/j.neuropsychologia.2010.04.016 [PubMed: 20430044]
- Wood ER, Dudchenko PA, Robitsek RJ, & Eichenbaum HB (2000). Hippocampal Neurons Encode Information about Different Types of Memory Episodes Occurring in the Same Location. Neuron, 27, 623–633. 10.1016/s0896-6273(00)00071-4 [PubMed: 11055443]

- Yiu AP, Mercaldo V, Yan C, Richards B, Rashid AJ, Hsiang HLL, ... Josselyn SA (2014). Neurons Are Recruited to a Memory Trace Based on Relative Neuronal Excitability Immediately before Training. Neuron, 83(3), 722–735. 10.1016/j.neuron.2014.07.017 [PubMed: 25102562]
- Zhou Y, Won J, Karlsson MG, Zhou M, Rogerson T, Balaji J, ... Silva AJ (2009). CREB regulates excitability and the allocation of memory to subsets of neurons in the amygdala. Nature Neuroscience, 12(11), 1438–1443. 10.1038/nn.2405 [PubMed: 19783993]
- Ziv Y, Burns LD, Cocker ED, Hamel EO, Ghosh KK, Kitch LJ, ... Schnitzer MJ (2013). Long-term dynamics of CA1 hippocampal place codes. Nature Neuroscience, 16(3), 264–266. 10.1038/ nn.3329 [PubMed: 23396101]

Significance statement:

Single photon calcium imaging using head-mounted miniature microscopes in freely moving animals has enabled researchers to measure the long term stability of hippocampal pyramidal cells during repeated behaviors. Previous studies have demonstrated instability of neural circuit components including dendritic spines and axonal boutons. It is now known that single units in the neuronal population exhibiting behaviorally relevant activity eventually become inactive and that previously silent neurons can quickly acquire task-relevant activity. The function of such population dynamics is unknown. We show here that population dynamics differ for cells coding distinct task dimensions, suggesting such dynamics are part of a mechanism for latent memory reorganization. These results add to a growing body of work showing that maintenance of episodic memory is an ongoing and dynamic process.



Figure 1.

A, Task outline: each trial has a Study and Test Phase, separated by a 20-second delay. Each trial is followed by a ~15s inter-trial interval in the mouse's home cage, adjacent to the alternation maze (not shown). B, Performance of individual mice (separate colors) over all days of recording. Only sessions with performance above 70% were included, excluded sessions are marked in red. C, Example viral expression and lens placement in dorsal CA1. Green is GCaMP6f-EYFP, blue is DAPI. D, Map showing regions for activity analysis. Purple indicates Central Stem, Orange indicates return arms. E, Top: Activity maps for one

cell (a Turn Splitter Neuron; see Figure 2) over five days of recording. Each plot represents the average activity map for one set of task conditions, ordered clockwise from top-left: Study-Left, Study-Right, Test-Right, Test-Left. In each plot, the black trace is the animal's recorded position, and colored dots indicate frames where the cell was active. Dots are colored based on the local transient likelihood, normalized by local occupancy, where red is the highest transient likelihood within that day and blue is the lowest. Bottom: Cell ROI masks for that recording day. Cell of interest is colored in green, and indicated with red arrow on first day shown. Masks were aligned across days based on relative positions of cells and cells were aligned based on distance between cell centers and correlation of masks. F, Same as E but for a place cell.





Figure 2.

A, Example activity maps for each type of splitter neuron on the central stem. B, Percentage of cells in each recording that had at least one calcium event on the central stem. Colored lines indicate individual animals, black line is best fit regression. Statistic: Spearman rank correlation. C, Percentage of splitter cells out of the active cell population on each day for all animals. Box shows inter-quartile range and middle line shows median. Statistic: Wilcoxon signed-rank test. D, Percentage of splitter neurons in individual animals (unique colors) and group regression (black) over the course of the experiment. Color of box indicates splitter

type as described by y-axis label. Significance calculated with Spearman rank correlation between percentage of splitters and recording day number for all included sessions (n=38). E, Same as B but for percentage of cells that exceed the activity threshold (see Methods). F, Same as C but for reliably active cells. G, Distribution of centers-of-mass of event activity for Turn and Phase splitter neurons. Statistic: Mann-Whitney U-test. H, same as D but for reliably active cells.

* p<0.05, ** p<0.01, ***p<0.001

Levy et al.



Figure 3.

A, Method for making population vector correlations. Calcium event likelihoods for one day, from one trial type, and from one spatial bin are correlated against event likelihoods from another (or the same) day and trial type but in the same spatial bin. Calcium event likelihoods are included for all cells found on both recording days of interest. B, Population vector correlations between trials of the same turn direction and task phase (gray), different turn directions (red) and different task phases(blue). Correlations in this panel B are generated from trials that occur on the same day. Shaded patch indicates 95% of points for

the indicated correlation type in that spatial bin, trend line indicates mean. Statistic: Wilcoxon rank-sum test on all points for these groups. C,D, Mean correlation for pairs of spatial bins over the course of recordings. Thin lines indicate individual animals' correlations, bold lines are best fit regression. Statistic: Spearman rank correlation on points from all recording days. E,F, Correlations between trials on separate recording days for indicated pairs of spatial bins. See text and supplementary data tables for statistics. * p<0.05, ** p<0.01, ***p<0.001

Levy et al.



Figure 4.

A, Percentage of cells that are still present at increasing day lags. Statistic: Wilcoxon signedrank test. B, Percentage of each splitter type by what that cell was on the prior day of recording. C, Percentage of active cells that were inactive on the prior day of recording. Colored lines are individual animals, black line is best fit regression. Statistic: Spearman rank correlation. D, Percentage of each splitting phenotype among each recording day's set of previously inactive cells (from second recording day forward). Statistic: Wilcoxon signedrank test. E, Changes in the distribution of splitting phenotypes among previously inactive over the course of recordings. Colored lines are individual animals, black line is best fit regression. Color of box indicates cell type as described by y-axis label. Statistic is indicated at right: Spearman rank correlation.

* p<0.05, ** p<0.01, ***p<0.001