Error-driven changes in hippocampal representations accompany flexible re-learning

P. Dylan Rich[1✉], Stephan Y. Thiberge[1,2], Nathaniel D. Daw[1,3]*, David W. Tank[1,2,]*

✉corresponding author
*co-senior authors
[1]Princeton Neuroscience Institute, Princeton University, Princeton, NJ, USA
[2]Bezos Center for Neural Circuit Dynamics, Princeton University, Princeton, NJ, USA
[3]Department of Psychology, Princeton University, Princeton, NJ, USA

Abstract

Flexible behavior requires both the learning of new associations, and the suppression of previous ones, but how neural circuits achieve this balance remains unclear. Here we show that continuous changes in hippocampal representations, known as drift, may facilitate this process. We used voluntary head-fixation and calcium imaging to record from CA1 in rats during an odor-guided navigation task that required frequent re-learning. We found systematic representational changes over the course of the multi-hour sessions that were increased following errors. A simple neural network model revealed that such error-driven drift can enable flexible re-learning by allowing new associations to form from new neural patterns. A consequence of this is that previous associations are maintained in latent synaptic weights. These findings reconcile the apparent tension between representational drift and stable memory storage, demonstrating how dynamic neural codes could support both flexible behavior and lasting memories.

Introduction

Humans and animals must learn what to do in new situations. Such learning is thought to rely on the formation and strengthening of associations between neural representations of current states or situations and corresponding downstream actions that lead to a desired outcome (Brown and Sharp, 1995; McNaughton and Morris, 1987). The hippocampus is believed to play a central role in this process by providing a neural representation of states or situations; cells in the hippocampus code for locations (O'Keefe and Dostrovsky, 1971), collectively discriminate between different contexts (Leutgeb et al., 2004; McKenzie et al., 2014; Muller and Kubie, 1987), task states (Aronov and Tank, 2014; Pastalkova et al., 2008; Sun et al., 2020) and abstract quantities (Nieh et al., 2021). Moreover, the reactivation of hippocampal cells is sufficient to elicit the recall of previously learned behaviors (Liu et al., 2012; Robinson et al., 2020).

Although initially considered to be stable (Thompson and Best, 1990), spatial representations in the hippocampus change over time, even with the same conditions are revisited (Kentros et al., 2004; Manns et al., 2007; Ziv et al., 2013). This phenomenon, termed representational drift, provides a challenge for a state representation view of the hippocampus' function, which tends to presuppose a static neural representation that can allow the retrieval of learned information. How the dynamic nature of hippocampal representation relates to its role in learning and memory remains to be explored.

One view is that representational drift is undesirable, a by-product of synaptic turn-over (Mongillo et al., 2017), and a problem for stable behavior (Driscoll et al., 2022), with mechanisms proposed for how downstream circuits can compensate for drift (Rule and O'Leary, 2022). An alternative view is that drift is useful, and that changes in representations are required for ongoing learning (Driscoll et al., 2022; Micou and O'Leary, 2023) and dynamic memory encoding (Mau et al., 2020). Computational models of memory function such as the temporal context model (Howard and Kahana, 2002) posit that an ongoing, changing neural representation links events that occur close in time (Cai et al., 2016; Rubin et

45   al., 2015). A dynamic representation has also been proposed to support the adaptive learning rates that
46   can be useful in volatile environments (Razmi and Nassar, 2022).
47   Although the idea that representational changes exist to accommodate changes in an
48   environment and enable learning is intuitive, evaluating this hypothesis rigorously has been hampered by
49   limitations in existing studies where drift has been observed. In order to effectively map place cells,
50   animals are typically engaged in simple tasks, such as foraging, that do not have an explicit learning
51   component that is under experimental control. Behavioral factors, such as running speed, that are known
52   to modulate hippocampal responses (McNaughton et al., 1983) are often uncontrolled and may explain
53   aspects of drift (Liberti et al., 2022; Sadeh and Clopath, 2022). Finally, drift is often assessed across days
54   or sessions (Geva et al., 2023; Khatib et al., 2023), masking how the short timescale elements of the
55   change (Bittner et al., 2017; Priestley et al., 2022) may relate to specific learning events.
56   To shed light on these issues, we recorded hippocampal activity while animals were engaged in a
57   task where learning and re-learning could be examined multiple times within a session. We used
58   voluntary head-fixation (Rich et al., 2024) and two-photon calcium imaging to record the activity of
59   neurons in hippocampal region CA1 during the stimulus presentation epoch of an odor-cued navigation
60   task. During this period animals' heads were in the same position, allowing unambiguous attribution of
61   neural firing differences to cognitive variables (Igarashi et al., 2014; MacDonald et al., 2013). Crucially,
62   the spatial configuration of the maze animals navigated was changed multiple times per session,
63   requiring animals to periodically adapt to new contingencies.
64   We found the task representation in CA1 was driven by the cue stimulus, changed throughout
65   the session, and that these changes were increased following trials when the animal made a navigational
66   error. To understand the consequences of drift in CA1 to a downstream circuit, we simulated a simple
67   feed-forward network that had to perform the same task and adapt to changing odor-action associations.
68   The model illustrated the advantage of the evolving task state representation for both learning new
69   associations and suppression of previous, irrelevant ones (Razmi and Nassar, 2022) . However, previous
70   rules were not entirely lost, but, as a consequence of the ongoing drift, were preserved in latent synaptic
71   weights. For this reason, such a model suggests that representational drift could serve as a common
72   mechanism bridging two processes that have typically been thought to have separate neural and
73   computational substrates (Yu et al., 2021): both incremental adaptive learning (Razmi and Nassar, 2022)
74   and latent state retrieval .

Results

76   We trained animals to perform a novel olfactory-cued navigation task (Fig. 1a-c). Animals had to navigate
77   to cued goals within an environment that was reconfigured multiple times within the same session;
78   animals had to learn and re-learn discrete spatial contingencies in blocks of trials. Each trial in the task
79   began with the animal sampling an odor at the central start location. The particular odor the animal was
80   presented with indicated that there was a reward available at one of four remote goal locations. Animals
81   learned this odor goal association over the shaping period, and were able to go to the correct goal with
82   high accuracy (Fig. 1f), indicating a high confidence with the basic task structure. In blocks of trials within
83   a single session, the location of a barrier within the maze was moved to one of three discrete locations
84   (Fig. 1a). The location of the barrier determined the correct route to each goal (Fig. 1b, Fig. 2b).
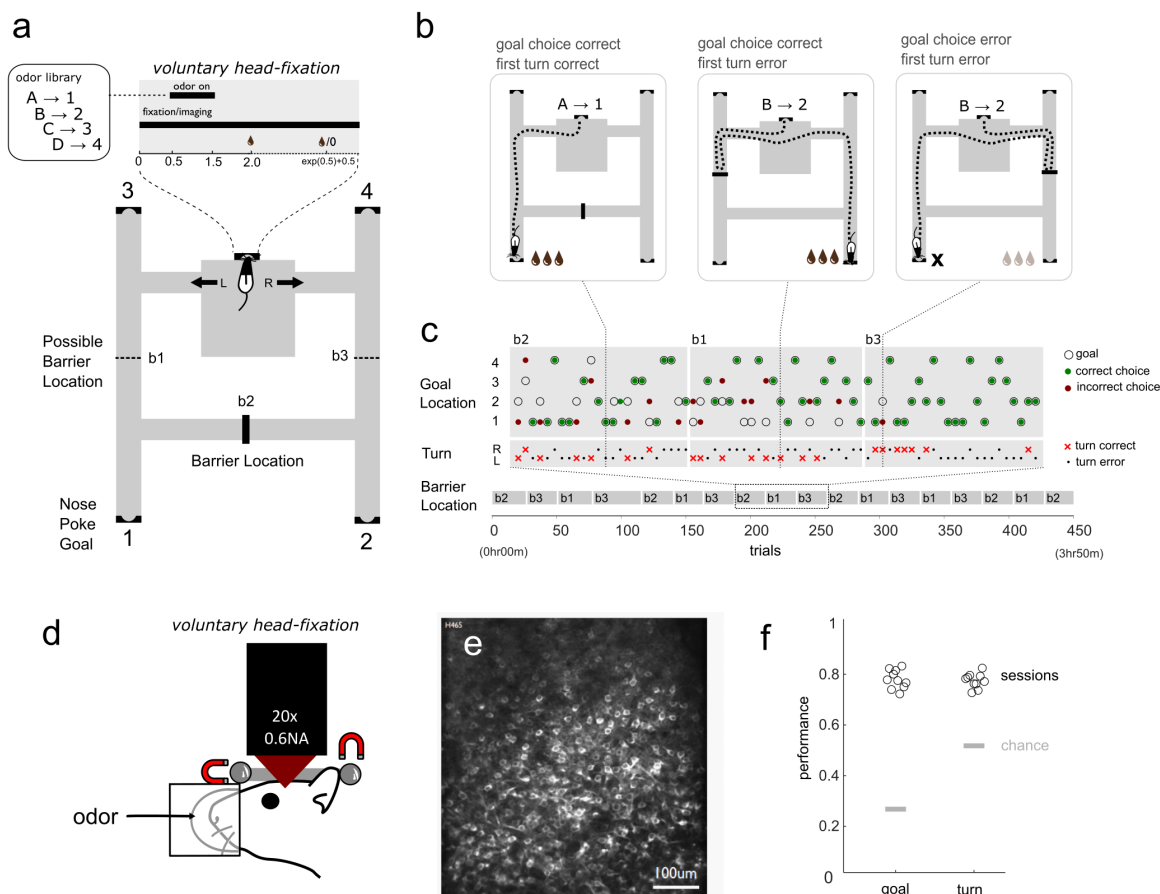85

Figure 1 Rats perform a flexible navigation task.

 a – Animals perform the flexible navigation task in a maze. In the central chamber, animals are presented with an odor during a voluntary head-fixation period. Following the head-fixation animals choose to turn left or right to proceed to one of four remote goal locations. The presented odor indicates which goal location has a reward available. Animals have to make an immediate left or right turn decision to navigate to the correct goal. Within the maze, there is a single barrier in one of three locations which is moved in blocks of trials within the session. **b** – Showing three example trials from a session. Animals can make a correct or incorrect goal choice and the first turn the animal makes can either be correct or an error. **c** – The barrier location is changed within blocks of trials within a single session, the trials for three example blocks are shown. **d** – Voluntary head-fixation during odor presentation, the animal is trained to maintain its head in precisely the same position, aided by a light magnetic force and bearing system. Within trial stability and across trial repeatability allows 2-photon calcium imaging of neurons in hippocampal region CA1. **e** – A subset of cells in CA1 express the calcium sensor GCaMP6f, which allows the monitoring of the activity of a large number of neurons. **f** – Over all sessions, performance is generally high for both the goal choice (chance is 0.25) and the first turn (chance is 0.5).

Animals learn new routes in a continual spatial learning task

For two of the goal locations, the correct turn response must be *changed* after certain block switches; for the other two goal locations, the turn response is always *fixed* regardless of what position the barrier is in (Fig. 2b). We looked at the profile of performance for the first turn following a block transition; performance improved as a function of distance from the block transition for changed-turn goals (Fig. 2c) indicating that animals learned the new correct first turn that corresponded to a new route for a given

109   goal. In contrast for fixed-turn goals, performance remained high, indicating animals were able to
110   maintain a fixed policy for some conditions despite the environmental changes going on.
111   In principle, animals could use a contextual switching strategy to solve this task; the same
112   conditions are revisited multiple times every session, so the animal would have to retrieve the correct set
113   of actions previously learned for each of the three conditions or contexts. Such a strategy would predict
114   fast switching of behavior following a block transition, for instance using state inference (Gershman et al.,
115   2014; Wilson et al., 2014), but this expectation was at odds with the observed gradual learning curves
116   (Fig. 2d,e).
117   Such gradual or incremental learning is more typically associated with trial-and-error learning, a
118   strategy where an agent learns about the correct action separately for each stimulus (in this case, the
119   odor-goal pair) based on the direct outcomes of its action (Daw et al., 2011). In its simplest form, this
120   strategy would predict learning rates that depended only on the direct experience animals had with each
121   odor, which is the same for all types of blocks. However, animals showed different profiles of learning in
122   different block transition types, in deviation from this simple view (Fig. 2d,e). Specifically, we observed
123   faster learning in block transitions where more than one odor/goal had a changed route (Fig. 2e),
124   indicating that learning was in some way adaptively modulated based on changes in the environment.
125   We quantified these effects using a logistic regression to analyze animals' turn choice following a block
126   switch for changed goals (Fig. 2e). We found the number of trials following a block switch was a
127   significant predictor of the correct turn ($\beta_{ntrials}$ = 2.60, p = 2.11x10^{-10}), confirming that animals were
128   learning the correct action, and that this learning significantly interacted with the block type (one vs. two
129   goals changed routes, $\beta_{ntrial:blockType}$ = 2.04, p = 2.08x10^{-6}), indicating that animals are sensitive to broader
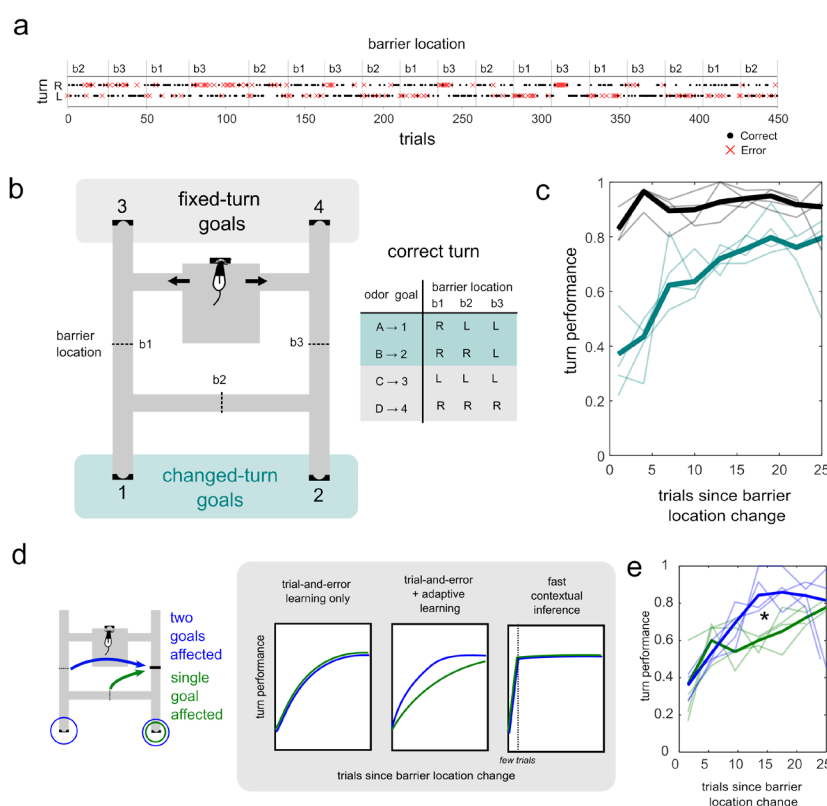130   environmental variability to adjust their learning strategy.
131

Figure 2 Rats learn new routes in a changing environment

a – A full session (same as shown in Fig1) showing turn decisions and the blocks of trials of the different barrier location. b – The goal locations in the maze could be characterized as *changed-turn goals* when the correct turn response was changed in different blocks of barrier location, and *fixed-turn goals* when the correct turn response was constant across all barrier locations. c – Learning curves for the two goal types, in bold shows the average learning curve and thin lines show individual animals. d – Diagram of two block transition types and predicted learning curves. In the diagram the colored arrows represent the movement of the barrier from one location to another. For the blue transition, the correct turn for two goal locations (circled) are affected. For the green transition one a single goal location (circled) is affected. The panels illustrate the predicted learning curves for three types of behavioral strategy. e – Learning curves for the two of two block transition types in bold shows the average learning curve and thin lines show individual animals. * – significant difference between the two block transition types ($\beta_{ntrial:blockType}$ = 2.04, p = 2.08x10^{-6}).

132
133

Neural responses in the hippocampus encode odor but not upcoming turn or barrier location

During the odor presentation, animals had been trained to voluntarily head-fixed beneath a two-photon microscope (Fig. 1), which, combined with an optical cannula and transgenic expression of GCaMP6f in pyramidal neurons, allowed us to monitor the population dynamics of a large number of dorsal CA1 neurons (n = 1143 over 10 sessions, session length 3.9 ± 0.8hrs mean ± std). The voluntary head-fixation system consists of a magnetic bearing system that aligns the head to the same position with micron-scale repeatability from trial-to-trial and animals can engage and disengage entirely at will (Rich et al., 2024). Voluntary head-fixation provides stability to perform population calcium imaging and allows us to unambiguously attribute any differences in neural signal in the hippocampus to cognitive rather than positional determinants (Liberti et al., 2022).

The critical decision-making during this task is sequestered within the epoch when the odor is presented and the animal has to make an immediate navigational decision (turn left or turn right, Fig. 1a). This decision is based on the integration of two distinct sources of information: the odor that is presented and the location of the barrier in the maze. We performed a linear decoding analysis to detect which of three behavioral variables: odor, turn, and barrier location were present in the ensemble activity. As expected (Igarashi et al., 2014), odor identity was decodable from the activity of cells in each trial shortly after, but not before odor presentation (all sessions $p < 0.01$, Fig. 3b). There was no pattern of significant decoding of the upcoming turn the animal was about to make (almost all time points for all sessions $p > 0.01$ Fig. 3b).

To ask whether the barrier location was represented in the ensemble, we used a decoder that was trained and tested over different blocks within the session (this was to ensure that we would be detecting consistent representation of barrier location - see methods). We were unable to decode the current barrier location when a classifier was trained on other blocks (almost all time points for all sessions $p > 0.01$, Fig. 3b), indicating that there was no reflection of the barrier location context in the hippocampal code during the decision-making period. More sophisticated classifiers that could account for a structured drift of context, that for instance preserved some consistent geometry between conditions (Keinath et al., 2022) also did not reveal any contextual coding (Fig. S1).

We next investigated whether the representation was stable over the multi-hour sessions. We hypothesized that a dynamic representation may facilitate some aspects of the behavior we saw; specifically we wanted to test whether ongoing learning might be facilitated by representational drift (Rule et al., 2019), which would imply an ongoing change in neural activity through the session. We trained odor classifiers on single blocks of trials, and evaluated them for other blocks (Fig. 3c). Looking at the average performance as a function of the lag between the training and the test blocks, we found that there was a gradual decline in the decoder performance for more distant lags. Additionally, the block identity itself (rather than the barrier location) could be decoded, indicating a robust implicit encoding of time within the session (Yang et al., 2024).
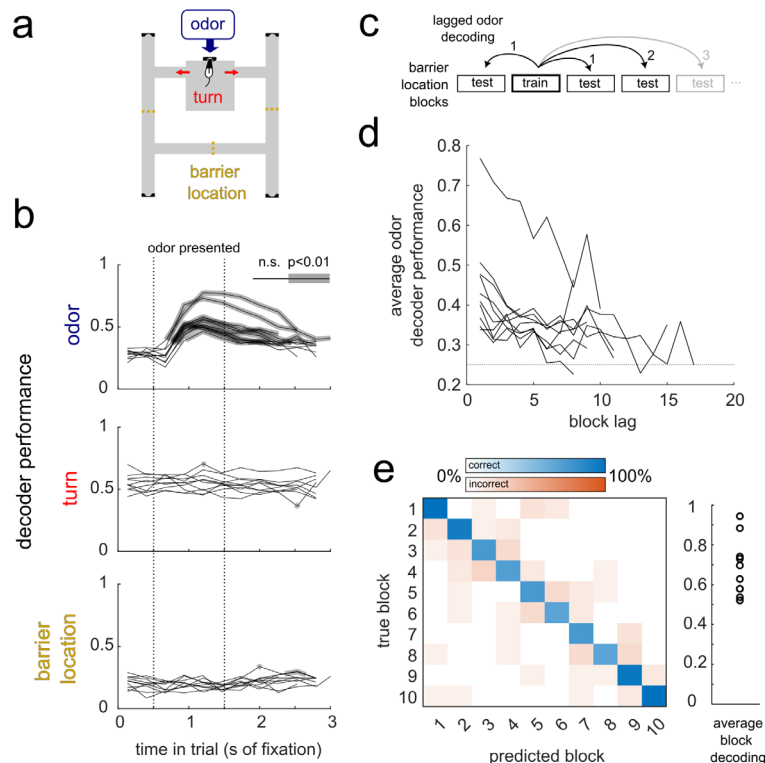
Figure 3 Decoding of odor, but not other task variables.

a – Illustration of the three factors in the task: The odor presented to the animal, the turn the animal takes and the location of the Barrier. **b** – The time courses of decoding performance during the head-fixation (recording) period in the task. Each line is an individual session. The lines are thickened where the decoding performance exceeded the p = 0.01 proportion of the null distribution, indicating above chance decoding performance. **c** – illustration of the lagged block odor decoder cross-validation scheme. **d** - The average performance of a whole-trial odor classifier as a function of the lag between training and testing. The lines show individual sessions; the nominal chance decoding is 0.25. **e** – Block decoding confusion matric for a single session. Plot on the right shows the average block decoding across different sessions.

194

Representation change is apparent at the individual and ensemble level and is continual throughout the session.

Inspection of individual units revealed clear non-stationarities in the responses of cells across the recording sessions. Neurons showed a long-tail distribution of the number of trials activity, and each trial showed a sparse recruitment of ~5% of neurons (Fig. S2). We observed a large variety of individual cell responses, with cells both turning on and off during a session, and with both gradual and abrupt time courses (Fig. 4a-c).

As well as strong responses to single odors, we noticed that cells also responded to multiple odors at varying levels, while being modulated by the overall time within a session. We used an encoding model for the individual neurons that explained the average activity of a cell within a trial as a multiplication of a parametric temporal response kernel over a session and a characteristic odor response vector (Fig. 4d). This approach allowed us to test whether a cell was modulated by time during the session, by the presented odor, or by a multiplicative combination of the two (Fig. 4d). We found cells in all categories, with a large proportion of cells showing this conjunctive session and odor coding. We also found a distribution of the odor responsiveness of cells, with a large proportion of cells showing a strong response to a single odor (Fig. 4e). Inspection of the significant session modulation kernels revealed a broad distribution of shape profiles for cells, with some cells showing small firing fields over the session's timeline, and others broader activation (Fig. 4f); these responses tiled the session (Fig. 4c, f). There was no correlation between the strength of the session modulation and the isolation quality of the ROI (Fig. S3) indicating that the modulation we saw was not a result of movement of the imaging plane (Fig. 4a). Finally, we noticed that many of the cells had a large transient early, or on the very first trial of their transition from silent to active (Fig 4a, Fig. S4). These large transients are likely plateau potentials that are the initiation events for behavioral timescale plasticity (Bittner et al., 2017).

218

$$y = s(x_{trials}) \times bx_{odor} + \varepsilon$$

Figure 4 –Neurons respond to odor and modulated within the session.

a – An example neuron showing the individual trial fluorescence activity within the head-fixation period. The first plot is sorted in order of the trials within a behavioral session, and the second plot is the same data, sorted first by the odor that was presented to the animal that trial. The periods of time that the odor was presented within the trial is shown as a black bar on the top of the plot. The three subpanels on the far right show the averaged image centered around that neuron for three selected trials throughout the session, with the dotted lines following the trials back to the session ordered plot. **b** – Neurons from the same session showing examples of the types of odor and session modulation that was observed. Each column is a single neuron, with the two orderings for each neuron as in A. The peak dF/F values are shown at the lower right corner of each plot. **c** – The average activity, following odor presentation, for each trial for all neurons that passed the minimum activity threshold within a single session. Each column is a single neuron, and the color scale corresponds to the average dF/F on each trial, following odor presentation. Before averaging, transients were detected and any non-transients in individual trials are set to zero. The neurons are sorted in three blocks, session modulated neurons, odor only modulated neurons and mean response only neurons. Session modulated cells are sorted within the category according to the center of mass of the session response kernel. **d** – The statistical model used to capture the average responses of individual cells (y) as a product of a parametric session response kernel and a characteristic odor response vector. To the right is the proportion of neurons in each category, showing a significant contribution of one or both of the two factors. **e** – The distribution of the normalized odor responses for all neurons. Neurons are sorted by the lifetime sparseness with respect to the estimated odor vector, and the odors are ordered for each neuron according to their responsiveness. **f** – Showing all session modulation kernels for all neurons that showed a significant modulation across the session. Each column is an individual neuron. The modulation kernels are aligned to the first recorded trial in the session, and are only shown for the number of trials within each session. The neurons are sorted first into total activity terciles according to cumulative sum of the session modulation kernel. Within each tercile neurons are sorted by the center of mass of the kernels.

Representational change increases with learning

We next asked whether the representational changes we observed were modulated by any behavioral features related to learning. One possibility is that representational change is entirely dependent on time, encoding the elapsed time within a session at a constant rate. Alternatively, there may be factors in the experiment that are related to learning that show correspondence with the representational change. To begin, we calculated a pairwise population vector correlation matrix to visualize the change in the ensemble response over a session (representational similarity matrix Fig. 5a) and plotted the average

249 neural similarity as a function of trial lag (Fig. 5b). In line with our univariate description and previous
250 work (Manns et al., 2007; Yang et al., 2024), we saw that time within a session was a strong predictor of
251 neural similarity, evidenced by a gradual fall off of neural similarity with trial lag. As expected from the
252 decoding analysis and univariate encoding analysis, pairs of trials with the same odor were more similar
253 than those with different odors (Fig. 5b). Breaking the category of same odor trial-pairs down further, we
254 found that the average population vector distance was systematically increased for changed-turn
255 odor/goals compared to fixed-turn odor/goals (Fig. 5b,c; see Fig 2b), suggesting that the learning
256 demands of the task increased the rate of change of the population.
257   To investigate the fine time-scale changes in neural representations and to test them statistically,
258 we examined neural representations between adjacent trials (consecutive trials in the experimental
259 sequence). This approach allowed us to investigate how representational changes might be related to
260 learning on a trial-by-trial level. We found that there was a significant effect of the fixed-turn category on
261 the neural representational distance between adjacent trials (β_{fixed} = -0.09, p = 0.016) even accounting for
262 the absolute time difference between trials ($\beta_{timeDiff}$ = 0.00, p = 0.169). Looking at adjacent trials also
263 allowed us to exclude the possibility that the difference between fixed-turn and changed-turn odor/goal
264 trials was driven by exposure frequency (fixed-turn odor/goal trials were presented less frequently than
265 changed-turn trials).
266   In many computational models of learning, learning is driven by errors, which can represent a
267 discrepancy between the predicted and actual outcome. We asked if the neural representational
268 similarity between adjacent trials for the same odor for changed-turn odor/goals was modulated by turn
269 errors. We found that neural distance was increased between adjacent trials of the same (changed-turn)
270 odor if there was an intervening turn error (i.e. an error on the first trial, Fig. 5d, β_{error} = 0.07 p =0.005),
271 this difference was not explained by absolute time difference between trials ($\beta_{timeDiff}$ = 0.00, p = 0.46). This
272 strongly indicates that the process of representational drift serves a purpose in the ongoing learning
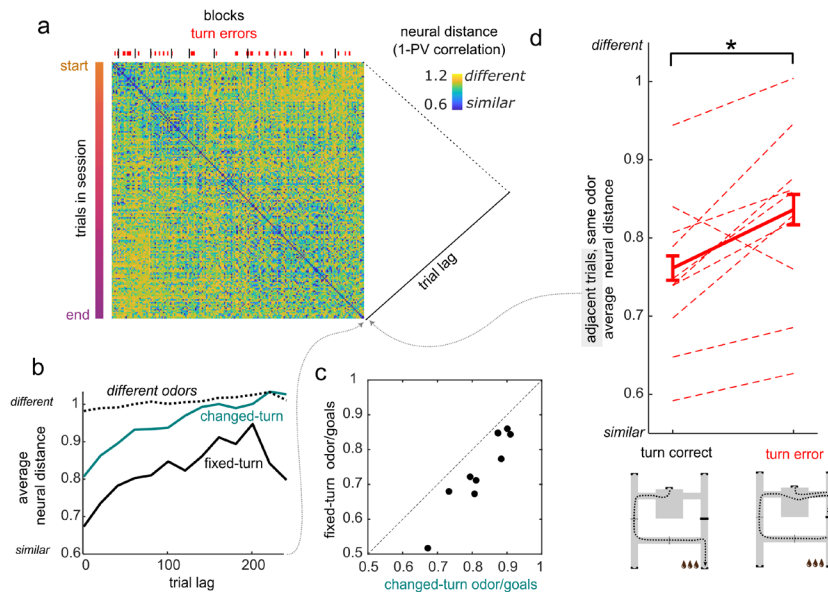273 process.
274

275

Figure 5 Neural change is modulated by factors of learning

a – Representational similarity matrix for a single session, each element is the trial-to-trial similarity, the neural distance is 1 minus the population vector correlation. Along the top edge shows the blocks and the turn errors within a session; along the side edge shows the trials in the session (matrix is sorted by time) **b** – the representational similarity matrix is collapsed along the diagonal to generate the similarity with respect to trial lag for a single session. The three conditions correspond to three trial-to-trial comparison types. Solid lines are between trials with the same odor, in the two categories (Fig. 2), dotted line is between trials of different odors (bin width 20 trials). **c** – The average distance between pairs of trials of the same odor for fixed vs changed turn conditions for the first trial lag bin in b (1-20 trials), shown for each session separately. **d** – The average neural distance between adjacent trials of the same odor, which is the collapse of the representational similarity matrix along the off diagonal, separated as to whether the animal made a turn error on the first trial. Dotted line represents the average values for individual sessions. * – p = 0.005 for main effect of error in mixed-effect model.

Representational change enables flexible re-learning without forgetting

In order to understand the mechanism of how dynamic, non-stationary neural data could give rise to the patterns of learning we observed in the behavior, we constructed a simple feedforward neural network model. The input layer corresponds to the CA1 response to a given odor, and the output units model the decision of whether to turn left or right, for instance in a downstream target (Fig. 6a). A weight matrix representing synaptic weights is updated based on a supervised Hebbian learning rule. Such learning rules typically rely on decay, normalization or depression of weights to achieve re-learning by erasing old weights; for instance, if the contingencies for a certain stimulus are reversed. However, the erasure of previous associations following re-learning is at odds with experimental evidence that associations can be recalled later following reactivation of tagged engram cells (Liu et al., 2012). To account for this, we included a constraint that a synapse could only *increase* in weight and would not decay or be depressed, allowing the possibility for associations that could persist in time. Under this constraint, a static input representation of the odor stimuli prohibits any re-learning since the previous association would compete with any new one (Fig. S5). We reasoned that drift in the input representation would allow re-learning by 1) moving activated inputs away from the previous rules in the weights and 2) providing a fresh, un-associated set of weights on which to make new associations.

We imposed representational change in the input layer in our network to mirror the neural data that we recorded in CA1 and investigated two conditions for how the drift/change was driven (Fig. 6b). In the first condition, the representational change occurred as a fixed amount per trial (time-driven), and in the second change only occurred in the input layer when the agent made a turn error (error-driven) - more closely mirroring our observation from the neural data. In both cases, the network was able to both re-learn new rules following block transitions (Fig. 6d). The drift of the input vector opened up a fresh, unused set of weights on which new association could be made, and the previously learned set of weights for changed units would not contribute to the action since the input units were now silent (Fig. S5).

328   Only the error-driven change model exhibited the learning rate difference between the different
329   categories of block transition that was seen in the behavioral data (Fig. 6d *c.f.* Fig. 2e). In the time-driven
330   drift, the rate of change was constant despite ongoing changes in performance. However, when the drift
331   was driven by turn errors, this caused an increased level of drift in blocks with more errors (blue, two
332   goal changed condition). Since the drift affects the representation generally, for both odors, this is then
333   seen as an increase in the overall learning rate. After an error for one odor there is a weakening of the
334   rule for the other odor; when both odors need to be relearned, this is beneficial, as it weakens the
335   previous, now incorrect association for the other odor. In the case where only one odor is relearned,
336   there are fewer errors and the generalized unlearning is balanced by the ongoing associative
337   strengthening that occurs on correct trials.
338   Finally, at the end of a simulated session, following learning, we simulated a hypothetical engram
339   reactivation experiment; we reactivated the patterns of input from pervious trials in the session and
340   presented them to the final end-of-session synaptic weights. The network was able to retrieve the
341   correct rule better than chance, indicating that a record of the previously appropriate actions was
342   successfully stored in the final weight matrix (Fig. 6d, Fig. S5). Although we used the full distribution of
343   time scales as estimated from the neural data, investigation of the effective weight matrices (Fig S –
344   effective weights) indicated that only the fast-modulated simulated neurons were contributing to this
345   recall.

Figure 6 Error-driven change enables adaptive re-learning

a – Feedforward neural network architecture. The Input layer is activated by the presentation of two odor, which project to an output layer with a turn left and turn right unit. The network had to perform a simplified version of the behavioral task, using only the two changed turn odors. The coarse temporal/session modulation of simulated cells is determined by resampling the session response kernels estimated from the neural data. **b** – The fine time-scale of the session modulation was tested in two conditions i) time driven which advanced the representation on every trial or ii) error-driven which advanced the representation following errors **c** - Activity of simulation neurons across the session. Neurons are sorted in the same manner as Fig. 4f, first into 3 terciles based on total activity and then, within each quartile by the center of mass within the session. **d** – The average learning curves of the network as a function of the number of trials since the block change (barrier changes location). The categories are the same as for Fig. 2b,c: block transitions with two goals affected and block transitions with only a single goal affected. The last data point is the performance of the network during an engram reactivation test, where activations from previous points in the session are sent through the final weight matrix that has been learned at the end of the session.
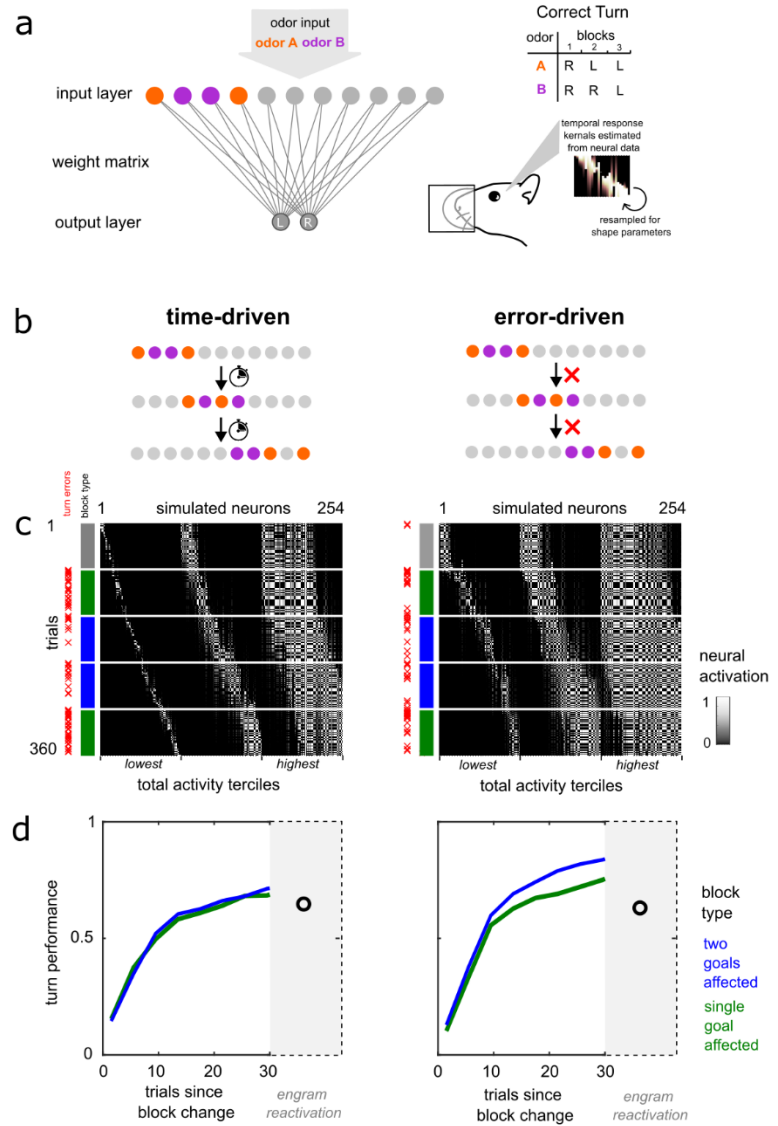
Discussion

We present a study to examine the neural bases of how animals adapt and learn in changing environments. We designed a novel learning paradigm that challenged animals to learn and re-learn, revisiting repeated conditions, throughout long, multi-hour sessions. We used magnetic voluntary head-fixation during the crucial decision-making epoch in the task, allowing us to monitor large numbers of pyramidal neurons in the CA1 of dorsal hippocampus. We did not observe a representation of the maze configuration in the neural data, nor behaviorally the sort of fast rule switching that such a representation could support. Rather, the neural task representation, which primarily encoded the instructional odor stimuli, changed steadily throughout the session. On a trial-by-trial level the representational change was increased following errors, indicated the role of the change for ongoing learning in the task. Although often considered a problem for the stable readout of activity in downstream circuits (Rule et al., 2019), we show how this representational change may in fact serve a useful purpose by allowing flexible behavior, enabling the learning of new rules and the suppression or extinction of previous ones. This architecture resolves a tension between stable memory storage and representational drift, as previous associations are stored and may be recalled from latent synaptic weights. Finally, the modulation of the rate of drift by errors emerges as a mechanism by which learning rates can be altered in environmental conditions with varying volatility, indicating a role for drift beyond simple temporal encoding.

Drift in Hippocampal representations

Activity in the hippocampus delineates spatial (Leutgeb et al., 2005; Muller and Kubie, 1987) latent (Pastalkova et al., 2008) and abstract (Nieh et al., 2021) contexts. Reactivation studies have shown that this activity confers the contextual specificity of long-term memories (Liu et al., 2012). However, hippocampal activity is increasingly appreciated to change over time, showing different responses even when identical environmental conditions are experienced (Bladon et al., 2019; Kentros et al., 2004; Mankin et al., 2012; Manns et al., 2007; Ziv et al., 2013). In our work, we make a number of contributions to the understanding of representational drift and representations in the hippocampus.

Much of the previous information about hippocampal drift comes from looking at representations across sessions or days (Mankin et al., 2012; Ziv et al., 2013), although some studies have looked at fast time-scale changes (Bittner et al., 2017; Monaco et al., 2014). The long sessions we used provided the temporal sampling required to estimate fast and slow changes in coding. We found a diversity of time courses of modulation of individual cells,

The hippocampus also encodes non-spatial elements of the environment (Wood et al., 1999; Manns et al., 2007; Igarashi et al., 2014), we demonstrate that representational changes are not exclusive to the spatial representations in the hippocampus, which have been the focus of previous investigations of drift (Kentros et al., 2004; Geva et al., 2023; Khatib et al., 2023). Head-fixation allows an unambiguous physiological demonstration of drift since potential positional determinants of neural coding are eliminated (Liberti et al., 2022); two-photon calcium imaging is ideally suited for the study of representational drift (Driscoll et al., 2017) by eliminating any concerns that change is due to electrode movement.

State representations and mechanisms of learning

Learning can be envisaged as the process of forming synaptic links from the representation of the current state to the appropriate action (Brown and Sharp, 1995). When the environment changes, the appropriate actions may also change. However, the encoding of new information can lead to catastrophic interference (McCloskey and Cohen, 1989) a well-known problem facing a neural network when existing information is overwritten by the new. We have demonstrated how a dynamic representation can ameliorate this concern and provides a scalable solution for continual learning (Driscoll et al., 2022).

431   In our model of drift and learning, the dynamic state representation serves two roles when
432 contingencies change. The first is to suppress the expression of previous learned information, which is
433 stored in the downstream weights of now silent neurons, which represents an orthogonal neural
434 population. The second is to provide a new set of active neurons with fresh downstream weights for new
435 information to be stored. This idea has been proposed as a mechanism for memory encoding (Mau et al.,
436 2020; Driscoll et al., 2022; Antony et al., 2024) that would involve a subset of neurons that are recruited
437 to be part of the memory trace (Zhou et al., 2009).
438   Perhaps the closest previous theoretical model to ours is Razmi and Nassar's (2022) proposal of a
439 dynamic state mechanism underlying adaptive learning. Our data provide direct cellular support for this
440 concept and connect it to phenomena of hippocampal representational drift. The prominent temporal
441 context model (Howard and Kahana, 2002) proposes that a gradual drift of a contextual memory
442 representation explains contiguity effects in memory recall. This line of work emphasizes the role of
443 experiential effects (stimulus encounters rather than the passage of time per se) in driving such change,
444 and evidence from both animals (Geva et al., 2023; Khatib et al., 2023) and humans (DuBrow et al., 2017)
445 has suggested that the rate of drift in the hippocampus is not solely temporally driven, but is influenced
446 by experience and behavior.
447

No upcoming turn activity seen in the hippocampus
449 We did not see any decoding of the turn the animal made in the neural activity in CA1. This was
450 somewhat surprising given the existence of prospective trajectory modulation of place activity in the
451 hippocampus (Frank et al., 2000; Wood et al., 2000); however such "splitter cell" activity is not always
452 observed (Berke et al., 2009; Duvelle et al., 2023). Furthermore, most studies observe prospective activity
453 during locomotion either on a track or a running wheel (Duvelle et al., 2023), indeed in a brain-wide
454 survey during a head-fixed visual decision task, the hippocampus was one of the few areas that did not
455 show upcoming choice coding (Steinmetz et al., 2019). Many studies that observe trajectory-modulated
456 activity have the structure that memory of the prior trial uniquely determines the upcoming response,
457 whereas in our task the decision to turn left or right is more complex, requiring the integration of
458 information about both the current odor stimulus and the previous history.
459

Absence of latent context representation
461 We did not observe any reflection of the latent context, the closed barrier location, in our neural data;
462 our analysis would have detected invariant stationary code as well as a drifting code with consistent
463 geometry or similarity for contexts (Keinath et al., 2022). Although well-trained, the animals in this study
464 did not immediately deploy the correct rules in a given context block despite having experienced it many
465 times before. Latent context models (Gershman and Niv, 2010; Wilson et al., 2014) suggest that
466 previously learned rules can be accessed through inference of a latent context representation, providing
467 behavioral benefits such as fast learning or generalization. The lack of contextual reactivation in the
468 neural data during our task is therefore, entirely consistent with the lack of fast contextual switching in
469 the behavior.
470   A caveat to our findings is that we are not recording from the hippocampus throughout the
471 whole task. Although we designed the task to sequester the crucial decision-making period in the head-
472 fixation imaging epoch, it may be that a latent context representation would be activated at other
473 timepoints, for instance when the animal is actively navigating. However, even with continuous
474 recording, it may prove difficult to definitely interpret any putative contextual signal since there would be
475 covariate positional changes such as slowing down before encountering a closed barrier.
476   Contextual coding in the hippocampus is frequently observed when context is overt, either
477 spatial context (Muller and Kubie, 1987,Kubie et al., 2020) or stimulus-defined contexts (Markus et al.,
478 1995; McKenzie et al., 2014). However, other studies with a similar, block-based switching task to ours
479 did not any behavioral context switching in rats or any reflective contextual neural representation in CA1

480   (Bladon et al., 2019), suggesting that the reliance on incremental learning over contextual inference may
481   be a common strategy adopted by animals in certain conditions. Simple incremental learning algorithms
482   fail to account for learning phenomena such as savings over repeated reversals and reinstatement
483   following extinction (Redish et al., 2009) where animals show patterns of learning that deviate from what
484   would be expected for a direct incremental learner. In our study, we saw other deviations from simple
485   incremental learning (Fig. 2e) that can be accounted for by the effective learning rate being modulated by
486   errors.
487

488   **Incremental learning and state inference**
489   Trial-and-error incremental learning and latent state inference have been viewed as distinct cognitive
490   processes with specialized neural mechanisms that can both contribute to behavior (Bouchacourt et al.,
491   2022). In our account, these two seemingly disparate processes can in fact be understood as separate
492   operations of the same system (Yu et al., 2021). A dynamic state representation provides a fresh neural
493   substrate for continual learning of new information and previously learned information is stored and can
494   be recalled via state inference. Mechanistic questions that have risen from these separate accounts can
495   be reinterpreted in terms of the control of hippocampal representational dynamics. The rate of change of
496   the state representation may be used to modulate learning rates (Razmi and Nassar, 2022), and the
497   error-driven change we observe is a simple implementation of what might otherwise be a complex
498   estimation of environment volatility (Piray and Daw, 2024).
499   What factors might influence reliance on contextual or trial-and-error incremental learning
500   strategy? The frontal cortex is thought to be critical for latent state inference (Wilson et al., 2014) which
501   could act as a controller to reactivate previously active patterns in the hippocampus and recall the
502   information. While a static context representation allows rapid access to a fixed set of rules, such a
503   system may take time to learn and potentially less useful in dynamic environments, especially where
504   there is uncertainty about the statistics of change.
505   Although in our experiment we don't see any instantiation of the latent context, our model
506   predicts that reactivation of patterns of previous activity from a previous block, for instance cell-specific
507   optogenetic stimulation (Rickgauer et al., 2014; Robinson et al., 2020), would cause the animal to retrieve
508   the appropriate rule from that block, rather than simply providing the percept of the stimulus.
509   Furthermore, we would predict that in conditions where fast contextual switching was seen in the
510   behavior, a concordant contextual representation would be seen in the hippocampus.
511   The distribution of modulation rates across the population we observe in our data also suggests a
512   reservoir of temporal scales that would be useful for learning associations across different time scales.
513   We found that only the fast-modulated neurons would be contributing to recall of the block specific
514   information since it was varying at a similar timescale. In the same way that place cell field propensity is
515   broadly distributed across the population to enable efficient coding of large and small environments
516   (Eliav et al., 2021; Rich et al., 2014), a multi-scale distribution of temporal modulation would allow a
517   similar advantage for making associations over various times-scales.
518

519   ***Outlook and Future Directions***
520   Representational drift has also been observed in other brain regions (Deitch et al., 2021; Driscoll et al.,
521   2017; Karlsson et al., 2012; Schoonover et al., 2021; Tsao et al., 2018) indicating that it is a general
522   feature of many neural systems. Understanding the purpose and mechanisms of dynamic, non-stationary
523   neural circuits generally is a pressing problem in contemporary neuroscience (Driscoll et al., 2022; Rule et
524   al., 2019) and advances in longitudinal neural recording capabilities (Rich et al., 2024) have created an
525   opportunity for the continued investigation of this phenomenon.
526

References

Antony J, Liu XL, Zheng Y, Ranganath C, O'Reilly RC. 2024. Memory out of context: Spacing effects and decontextualization in a computational model of the medial temporal lobe. *Psychological Review*. doi:10.1037/rev0000488

Aronov D, Tank DW. 2014. Engagement of neural circuits underlying 2D spatial navigation in a rodent virtual reality system. *Neuron* **84**:442–456. doi:10.1016/j.neuron.2014.08.042

Berke JD, Breck JT, Eichenbaum H. 2009. Striatal versus hippocampal representations during win-stay maze performance. *J Neurophysiol* **101**:1575–1587. doi:10.1152/jn.91106.2008

Bittner KC, Milstein AD, Grienberger C, Romani S, Magee JC. 2017. Behavioral time scale synaptic plasticity underlies CA1 place fields. *Science* **357**:1033–1036. doi:10.1126/science.aan3846

Bladon JH, Sheehan DJ, Freitas CSD, Howard MW. 2019. In a Temporally Segmented Experience Hippocampal Neurons Represent Temporally Drifting Context But Not Discrete Segments. *J Neurosci* **39**:6936–6952. doi:10.1523/JNEUROSCI.1420-18.2019

Bouchacourt F, Tafazoli S, Mattar MG, Buschman TJ, Daw ND. 2022. Fast rule switching and slow rule updating in a perceptual categorization task. *eLife* **11**:e82531. doi:10.7554/eLife.82531

Brown MA, Sharp PE. 1995. Simulation of spatial learning in the Morris water maze by a neural network model of the hippocampal formation and nucleus accumbens. *Hippocampus* **5**:171–88. doi:10.1002/hipo.450050304

Cai DJ, Aharoni D, Shuman T, Shobe J, Biane J, Song W, Wei B, Veshkini M, La-Vu M, Lou J, Flores SE, Kim I, Sano Y, Zhou M, Baumgaertel K, Lavi A, Kamata M, Tuszynski M, Mayford M, Golshani P, Silva AJ. 2016. A shared neural ensemble links distinct contextual memories encoded close in time. *Nature* **534**:115–118. doi:10.1038/nature17955

Daw ND, Gershman SJ, Seymour B, Dayan P, Dolan RJ. 2011. Model-Based Influences on Humans' Choices and Striatal Prediction Errors. *Neuron* **69**:1204–1215. doi:10.1016/j.neuron.2011.02.027

Deitch D, Rubin A, Ziv Y. 2021. Representational drift in the mouse visual cortex. *Current Biology* **31**:4327-4339.e6. doi:10.1016/j.cub.2021.07.062

Driscoll LN, Duncker L, Harvey CD. 2022. Representational drift: Emerging theories for continual learning and experimental future directions. *Current Opinion in Neurobiology* **76**:102609. doi:10.1016/j.conb.2022.102609

Driscoll LN, Pettit NL, Minderer M, Chettih SN, Harvey CD. 2017. Dynamic Reorganization of Neuronal Activity Patterns in Parietal Cortex. *Cell* **170**:986-999.e16. doi:10.1016/j.cell.2017.07.021

DuBrow S, Rouhani N, Niv Y, Norman KA. 2017. Does mental context drift or shift? *Current Opinion in Behavioral Sciences*, Memory in time and space **17**:141–146. doi:10.1016/j.cobeha.2017.08.003

Duvelle É, Grieves RM, van der Meer MA. 2023. Temporal context and latent state inference in the hippocampal splitter signal. *eLife* **12**:e82357. doi:10.7554/eLife.82357

Eliav T, Maimon SR, Aljadeff J, Tsodyks M, Ginosar G, Las L, Ulanovsky N. 2021. Multiscale representation of very large environments in the hippocampus of flying bats. *Science* **372**:eabg4020. doi:10.1126/science.abg4020

Frank LM, Brown EN, Wilson M. 2000. Trajectory encoding in the hippocampus and entorhinal cortex. *Neuron* **27**:169–178.

Gershman SJ, Niv Y. 2010. Learning latent structure: carving nature at its joints. *Curr Opin Neurobiol* **20**:251–256. doi:10.1016/j.conb.2010.02.008

Gershman SJ, Radulescu A, Norman KA, Niv Y. 2014. Statistical Computations Underlying the Dynamics of Memory Updating. *PLOS Computational Biology* **10**:e1003939. doi:10.1371/journal.pcbi.1003939

Geva N, Deitch D, Rubin A, Ziv Y. 2023. Time and experience differentially affect distinct aspects of hippocampal representational drift. *Neuron* **111**:2357-2366.e5. doi:10.1016/j.neuron.2023.05.005

575   Howard MW, Kahana MJ. 2002. A Distributed Representation of Temporal Context. *Journal of*
576   *Mathematical Psychology* **46**:269–299. doi:10.1006/jmps.2001.1388
577   Igarashi KM, Lu L, Colgin LL, Moser M-B, Moser EI. 2014. Coordination of entorhinal–hippocampal
578   ensemble activity during associative learning. *Nature* **510**:143–147. doi:10.1038/nature13162
579   Karlsson MP, Tervo DGR, Karpova AY. 2012. Network Resets in Medial Prefrontal Cortex Mark the Onset
580   of Behavioral Uncertainty. *Science* **338**:135–139. doi:10.1126/science.1226518
581   Keinath AT, Mosser C-A, Brandon MP. 2022. The representation of context in mouse hippocampus is
582   preserved despite neural drift. *Nat Commun* **13**:2415. doi:10.1038/s41467-022-30198-7
583   Kentros CG, Agnihotri NT, Streater S, Hawkins RD, Kandel ER. 2004. Increased Attention to Spatial Context
584   Increases Both Place Field Stability and Spatial Memory. *Neuron* **42**:283–295. doi:10.1016/S0896-
585   6273(04)00192-8
586   Khatib D, Ratzon A, Sellevoll M, Barak O, Morris G, Derdikman D. 2023. Active experience, not time,
587   determines within-day representational drift in dorsal CA1. *Neuron*.
588   doi:10.1016/j.neuron.2023.05.014
589   Kubie JL, Levy ERJ, Fenton AA. 2020. Is hippocampal remapping the physiological basis for context?
590   *Hippocampus* **30**:851–864. doi:10.1002/hipo.23160
591   Leutgeb S, Leutgeb JK, Barnes CA, Moser EI, McNaughton BL, Moser M-B. 2005. Independent codes for
592   spatial and episodic memory in hippocampal neuronal ensembles. *Science* **309**:619–623.
593   doi:10.1126/science.1114037
594   Leutgeb S, Leutgeb JK, Treves A, Moser M-B, Moser EI. 2004. Distinct ensemble codes in hippocampal
595   areas CA3 and CA1. *Science* **305**:1295–1298. doi:10.1126/science.1100265
596   Liberti WA, Schmid TA, Forli A, Snyder M, Yartsev MM. 2022. A stable hippocampal code in freely flying
597   bats. *Nature* **604**:98–103. doi:10.1038/s41586-022-04560-0
598   Liu X, Ramirez S, Pang PT, Puryear CB, Govindarajan A, Deisseroth K, Tonegawa S. 2012. Optogenetic
599   stimulation of a hippocampal engram activates fear memory recall. *Nature* **484**:381–385.
600   doi:10.1038/nature11028
601   MacDonald CJ, Carrow S, Place R, Eichenbaum H. 2013. Distinct hippocampal time cell sequences
602   represent odor memories in immobilized rats. *J Neurosci* **33**:14607–14616.
603   doi:10.1523/JNEUROSCI.1537-13.2013
604   Mankin EA, Sparks FT, Slayyeh B, Sutherland RJ, Leutgeb S, Leutgeb JK. 2012. Neuronal code for extended
605   time in the hippocampus. *Proc Natl Acad Sci USA* **109**:19462–19467.
606   doi:10.1073/pnas.1214107109
607   Manns JR, Howard MW, Eichenbaum H. 2007. Gradual Changes in Hippocampal Activity Support
608   Remembering the Order of Events. *Neuron* **56**:530–540. doi:10.1016/j.neuron.2007.08.017
609   Markus EJ, Qin YL, Leonard B, Skaggs WE, McNaughton BL, Barnes CA. 1995. Interactions between
610   location and task affect the spatial and directional firing of hippocampal neurons. *J Neurosci*
611   **15**:7079–7094.
612   Mau W, Hasselmo ME, Cai DJ. 2020. The brain in motion: How ensemble fluidity drives memory-updating
613   and flexibility. *eLife* **9**:e63550. doi:10.7554/eLife.63550
614   McCloskey M, Cohen NJ. 1989. Catastrophic Interference in Connectionist Networks: The Sequential
615   Learning Problem In: Bower GH, editor. Psychology of Learning and Motivation. Academic Press.
616   pp. 109–165. doi:10.1016/S0079-7421(08)60536-8
617   McKenzie S, Frank AJ, Kinsky NR, Porter B, Rivière PD, Eichenbaum H. 2014. Hippocampal Representation
618   of Related and Opposing Memories Develop within Distinct, Hierarchically Organized Neural
619   Schemas. *Neuron* **83**:202–215. doi:10.1016/j.neuron.2014.05.019
620   McNaughton B, Morris R. 1987. Hippocampal synaptic enhancement and information storage within a
621   distributed memory system. *Trends Neurosci* **10**:408–415.
622   McNaughton BL, Barnes CA, O'Keefe J. 1983. The contributions of position, direction, and velocity to
623   single unit activity in the hippocampus of freely-moving rats. *Exp Brain Res* **52**:41–49.

624   Micou C, O'Leary T. 2023. Representational drift as a window into neural and behavioural plasticity.
625   *Current Opinion in Neurobiology* **81**:102746. doi:10.1016/j.conb.2023.102746
626   Monaco JD, Rao G, Roth ED, Knierim JJ. 2014. Attentive scanning behavior drives one-trial potentiation of
627   hippocampal place fields. *Nat Neurosci* **17**:725–731. doi:10.1038/nn.3687
628   Mongillo G, Rumpel S, Loewenstein Y. 2017. Intrinsic volatility of synaptic connections — a challenge to
629   the synaptic trace theory of memory. *Current Opinion in Neurobiology*, Computational
630   Neuroscience **46**:7–13. doi:10.1016/j.conb.2017.06.006
631   Muller RU, Kubie JL. 1987. The effects of changes in the environment on the spatial firing of hippocampal
632   complex-spike cells. *J Neurosci* **7**:1951–1968.
633   Nieh EH, Schottdorf M, Freeman NW, Low RJ, Lewallen S, Koay SA, Pinto L, Gauthier JL, Brody CD, Tank
634   DW. 2021. Geometry of abstract learned knowledge in the hippocampus. *Nature* 1–5.
635   doi:10.1038/s41586-021-03652-7
636   O'Keefe J, Dostrovsky J. 1971. The hippocampus as a spatial map. Preliminary evidence from unit activity
637   in the freely-moving rat. *Brain Res* **34**:171–175.
638   Pastalkova E, Itskov V, Amarasingham A, Buzsáki G. 2008. Internally generated cell assembly sequences in
639   the rat hippocampus. *Science* **321**:1322–1327. doi:10.1126/science.1159775
640   Piray P, Daw ND. 2024. Computational processes of simultaneous learning of stochasticity and volatility in
641   humans. *Nat Commun* **15**:9073. doi:10.1038/s41467-024-53459-z
642   Priestley JB, Bowler JC, Rolotti SV, Fusi S, Losonczy A. 2022. Signatures of rapid plasticity in hippocampal
643   CA1 representations during novel experiences. *Neuron* **110**:1978-1992.e6.
644   doi:10.1016/j.neuron.2022.03.026
645   Razmi N, Nassar MR. 2022. Adaptive Learning through Temporal Dynamics of State Representation. *J*
646   *Neurosci* **42**:2524–2538. doi:10.1523/JNEUROSCI.0387-21.2022
647   Redish D, Jensen S, Johnson A, Kurth-Nelson Z. 2009. Reconciling Reinforcement Learning Models With
648   Behavioral Extinction and Renewal: Implications for Addiction, Relapse, and Problem Gambling.
649   *Psychological review* **116**:518. doi:10.1037/a0016243
650   Rich PD, Liaw H-P, Lee AK. 2014. Large environments reveal the statistical structure governing
651   hippocampal representations. *Science* **345**:814–817. doi:10.1126/science.1255635
652   Rich PD, Thiberge SY, Scott BB, Guo C, Tervo DGR, Brody CD, Karpova AY, Daw ND, Tank DW. 2024.
653   Magnetic voluntary head-fixation in transgenic rats enables lifespan imaging of hippocampal
654   neurons. *Nat Commun* **15**:4154. doi:10.1038/s41467-024-48505-9
655   Rickgauer JP, Deisseroth K, Tank DW. 2014. Simultaneous cellular-resolution optical perturbation and
656   imaging of place cell firing fields. *Nat Neurosci* **17**:1816–1824. doi:10.1038/nn.3866
657   Robinson NTM, Descamps LAL, Russell LE, Buchholz MO, Bicknell BA, Antonov GK, Lau JYN, Nutbrown R,
658   Schmidt-Hieber C, Häusser M. 2020. Targeted Activation of Hippocampal Place Cells Drives
659   Memory-Guided Spatial Behavior. *Cell*. doi:10.1016/j.cell.2020.09.061
660   Rubin A, Geva N, Sheintuch L, Ziv Y. 2015. Hippocampal ensemble dynamics timestamp events in long-
661   term memory. *eLife*. doi:10.7554/eLife.12247
662   Rule ME, O'Leary T. 2022. Self-healing codes: How stable neural populations can track continually
663   reconfiguring neural representations. *Proc Natl Acad Sci U S A* **119**:e2106692119.
664   doi:10.1073/pnas.2106692119
665   Rule ME, O'Leary T, Harvey CD. 2019. Causes and consequences of representational drift. *Current Opinion*
666   *in Neurobiology*, Computational Neuroscience **58**:141–147. doi:10.1016/j.conb.2019.08.005
667   Sadeh S, Clopath C. 2022. Contribution of behavioural variability to representational drift. *eLife*
668   **11**:e77907. doi:10.7554/eLife.77907
669   Schoonover CE, Ohashi SN, Axel R, Fink AJP. 2021. Representational drift in primary olfactory cortex.
670   *Nature* **594**:541–546. doi:10.1038/s41586-021-03628-7
671   Steinmetz NA, Zatka-Haas P, Carandini M, Harris KD. 2019. Distributed coding of choice, action and
672   engagement across the mouse brain. *Nature* **576**:266–273. doi:10.1038/s41586-019-1787-x

673 Sun C, Yang W, Martin J, Tonegawa S. 2020. Hippocampal neurons represent events as transferable units
674   of experience. *Nat Neurosci* 1–13. doi:10.1038/s41593-020-0614-x
675 Thompson LT, Best PJ. 1990. Long-term stability of the place-field activity of single units recorded from
676   the dorsal hippocampus of freely behaving rats. *Brain Res* **509**:299–308.
677 Tsao A, Sugar J, Lu L, Wang C, Knierim JJ, Moser M-B, Moser EI. 2018. Integrating time from experience in
678   the lateral entorhinal cortex. *Nature* **561**:57–62. doi:10.1038/s41586-018-0459-6
679 Wilson RC, Takahashi YK, Schoenbaum G, Niv Y. 2014. Orbitofrontal Cortex as a Cognitive Map of Task
680   Space. *Neuron* **81**:267–279. doi:10.1016/j.neuron.2013.11.005
681 Wood ER, Dudchenko PA, Eichenbaum H. 1999. The global record of memory in hippocampal neuronal
682   activity. *Nature* **397**:613–616.
683 Wood ER, Dudchenko PA, Robitsek RJ, Eichenbaum H. 2000. Hippocampal neurons encode information
684   about different types of memory episodes occurring in the same location. *Neuron* **27**:623–633.
685 Yang W, Sun C, Huszár R, Hainmueller T, Kiselev K, Buzsáki G. 2024. Selection of experience for memory
686   by hippocampal sharp wave ripples. *Science* **383**:1478–1483. doi:10.1126/science.adk8261
687 Yu LQ, Wilson RC, Nassar MR. 2021. Adaptive learning is structure learning in time. *Neuroscience &*
688   *Biobehavioral Reviews* **128**:270–281. doi:10.1016/j.neubiorev.2021.06.024
689 Zhou Y, Won J, Karlsson MG, Zhou M, Rogerson T, Balaji J, Neve R, Poirazi P, Silva AJ. 2009. CREB
690   regulates excitability and the allocation of memory to subsets of neurons in the amygdala. *Nat*
691   *Neurosci* **12**:1438–1443. doi:10.1038/nn.2405
692 Ziv Y, Burns LD, Cocker ED, Hamel EO, Ghosh KK, Kitch LJ, El Gamal A, Schnitzer MJ. 2013. Long-term
693   dynamics of CA1 hippocampal place codes. *Nat Neurosci* **16**:264–266. doi:10.1038/nn.3329
694
695

696 **Acknowledgements**

701

702 **Data and Code Availability Statement**
703 Upon publication, data will be made available on a publicly accessible data sharing site. Code for
704 analyzing the data will be available on GitHub

705

706 **Declarations of Interests**
707 The authors declare no competing interests

708

709 **Inclusion and Ethics Statement**
710 All animal procedures were performed in accordance with the Princeton Universty Institutional Animal
711 Care and Use Committee (IACUC) guidelines and approved under protocol #1837
712

713 # METHODS

714 **Animals and Surgery description**
715 All procedures performed in this study were approved by the Institutional Animal Care and Use Committee at Princeton
716 University (IACUC protocol number 1837) and were performed in line with the Guide for the Care and Use of Laboratory Animals
717 (National Research Council, 2011). We used 4 male transgenic rats that expresses GCaMP6f under the thy-1 promotor aged
718 between 3 and 9 months at the time of surgery, these rats had were among the animals reported in a previous report (Rich et al.,
719 2024).

720 Rats underwent surgery to obtain optical access to the dorsal hippocampus. Surgery was performed under aseptic conditions,
721 and animals were anesthetized with isoflurane (3.5% induction, 0.5 -1% for maintenance). Animals received pre-operative
722 buprenorphine (0.02mg/kg) and either post-operative meloxicam (1mg/kg) or buprenorphine (0.02mg/kg) for analgesia. Body
723 temperature was maintained with a heating pad (Harvard Apparatus). The skull was exposed and the periosteum was retracted.
724 A craniotomy was performed on the right hemisphere over the dorsal CA1 region of the hippocampus (4.2AP, 3.0ML) The
725 overlying cortex was aspirated to exposure the fibers of the external capsule, which was carefully retracted to exposure the
726 fibers of the alveus. A thin layer of Kwik-Sil (WPI) was squeezed on the exposed area and the custom conical canula (5/4 mm
727 top/bottom diameter 2.6mm height) was quickly lowered and cemented in place with adhesive cement (C&B Metabond,
728 Parkell). A titanium base head-plate was mounted parallel to the stereotaxic frame and cemented to the skull. The canula was
729 held in place at its rim by three magnets (1/16" diameter, D12-N52, K&J magnets) and was implanted centrally to the head-plate
730 aperture with the cover-glass parallel to the stereotaxic frame. These steps resulted in the cover-glass being parallel and
731 concentric to the head plate for imaging. Following surgery animals were left for ~1 month before the imaging window was
732 assessed under anesthesia. At this point the kinematic head-plate was installed into the base head-plate.

733 At the conclusion of the study, animals were anesthetized with an overdose of ketamine xyline and euthanized with a
734 transcardial perfusion of PBS followed by 4% paraformaldehyde. Brains were extracted, and left in fixative for 12 h and then
735 transferred to 30% sucrose in PBS until the brain had sunk. Brains were sectioned on a freezing microtome (Leica) in 50μm
736 sections and floated onto slides. Slides were inspected on a epifluorescence microscope to confirm that the field of view of the
737 cannula was in the CA1 region of the hippocampus as assessed by gross anatomy.

738 **Imaging system**
739 We performed imaging with a custom two-photon microscope controlled with the Scan Image software (Vidrio) running in
740 MATLAB (Mathworks). Laser illumination was provided by a Ti:Sapphire laser (Chameleon Vision II, Coherent) operating at 920
741 nm. We used a long working distance air immersion objective lens (20x/0.6NA/13mm WD, Edmunds optics) and a GaAsP
742 photomultiplier tubes (H10770PA-40, Hamamatsu). The beam power was modulated by a Pockels cell (350-80 LA BIC -02
743 Conoptics) and the power used for imaging, measured at the front of the objective, was 150-200mW. Horizontal scans of the
744 laser were achieved using a resonant galvanometer (Thorlabs). Typical fields of view measured approximately 600 × 600μm and
745 data were acquired at 30 Hz.

746 **Magnetic voluntary head-fixation**
747 Animals were first trained to voluntarily clip into a magnetic head-fixation system (Rich et al., 2024) that provided micron-scale
748 registration of the head required for trial-to-trial calcium imaging of the hippocampus. Animals were initially trained to do this as
749 previously described (Rich et al., 2024). Briefly, animals were first trained to poke at the stainless-steel nose poke beneath the
750 head fixation system, to receive a chocolate milk drop reward (1 drop, Ensure milk chocolate flavor nutritional shake, Abbott). A
751 variable period was introduced while the animal had to maintain the hold to receive the reward. If the animal made contact with
752 the bearing surfaces, it was given an immediate larger reward (2 drops). The various further steps of contact with the kinematic
753 clamp were shaped using this strategy hold small reward, next step no hold large reward, until the animals were making routine
754 fixations with the magnetic bearing system. At the start of all sessions animals typically were given 10 warm-up trials with
755 shorter fixations.

756 **Odor-guided navigation task**
757 Once animals had learned to reliably clip-in to the head fixation system they were trained to perform an odor guided navigation
758 task. We delivered odors (acetic acid 20%, 2-propanol, propyl acetate, anethole, trans-cinnamaldehyde, nutmeg, cardamom,
759 clove, rosemary, star anise) to animals using a custom built olfactometer. Air was directed to flow through vials containing
760 odorants using PTFE valves (Neptune Research) at a flow rate of 0.5L/min. At the start of head fixation, blank air was delivered
761 for 0.5s, followed by 1s of odorized air, followed by a scavenger vacuum through the system to evacuate any residual odorant on
762 that trial and for the next trial. Animals were trained that different odors corresponded to different reward location in a maze
763 constructed of enclosed linear track elements constructed of plastic U-channel and acrylic wall and ceiling. Different arms of the

764 maze were made texturally distinct by adding metal elements to the track floor such as small ramps or protruding screw heads.
765 Animals preformed the fixation and odor sampling in a central chamber. Once the animal had made the fixation for the required
766 duration (2-3s), breaking the fixation would open two choice doors (metal duct blast gates; McMaster-Carr) on the left and right
767 side of the animal, actuated with a servo moto. The animal then chose to turn left or right, and to proceed to the remote goal
768 location in the maze. The maze was configured with a single barrier location (metal duct blast gates; McMaster-Carr), that meant
769 that there was a single correct route to each goal, and so a single correct left right initial turn to make after fixation. During
770 navigation the maze room was lit with a dim purple/near UV LED lamp with additional smaller UV LEDs mounted on one wall to
771 aim as an orientation cue. These lights were turned off during imaging to prevent light contamination. The barrier location was
772 changed manually by the experimenter between each block of trials. If the animal nose poked to the correct goal location given
773 the odor presented it received a large reward (5-8 drops). A nose poked at an incorrect location would elicit a short white noise
774 sound and a 30s timeout. During the timeout the choice-doors were closed, preventing re-access to the main chamber. Following
775 a correct goal choice or the end of the timeout, the animal was free to re-enter the main chamber, where poke a rewarded poke
776 at the reset location at the rear of the main chamber, caused the choice doors to close, and the apparatus to be reset for the
777 next trial. Odors were presented pseudo-randomly in block of 12 trials, with 2 trials for each fixed and 4 trials for each changed
778 odor/goal trial types. Blocks for each barrier location were 24 -36 trials long. The task apparatus was controlled using Arduino
779 microcontroller and custom MATLAB software running of a laptop computer.

780 **Behavioral analysis**
781 We assessed the learning of correct turn using a logistic regression analysis. We only considered trials for goal locations that had
782 changed a route and therefore the correct turn in a new block. The probability of a correct turn was modeled function of the
783 number of trials since the blocks transition and the categorical block transition type. The block transition type was one of two
784 categories that depended on the previous and current location of the barrier. A single goal turn change was a move of the barrier
785 that only affected the correct turn for one goal location and a multiple goal turn change was a move of the barrier that affected
786 two (both) of the (learned) goal locations. The logistic regression included an intercept: to account for the baseline on the first
787 trial of the block; a main effect of the number of trials since the barrier location change, to capture the animals' improvement
788 over each block; and an interaction term for the number of trials since the barrier location change and the block transition type,
789 to account for differences in the slope the learning curves. We did not include a main effect of the block transition alone, since it
790 would have modeled a difference in performance before the animal knew what block transition type it was in; which was not
791 supported empirically by the data, or in principle since the block transition type order was random. We fit the model across the
792 dataset of animals and sessions using a mixed effects model in MATLAB, including the random effect of animal ID.

793 **Data pre-processing and neuron identification**
794 To correct for non-rigid brain motion in imaging data, we employed custom MATLAB code based on the NoRMCorre algorithm.
795 (Nieh et al., 2021, Pnevmatikakis and Giovannucci 2017). This method involves dividing the image into overlapping patches and
796 estimating a rigid translation for each patch and frame by aligning it with a template. The resulting translations are then
797 interpolated to form a motion field, which is applied to smaller overlapping patches. The displacement for the x and y axes were
798 determined by the maximum values observed across all patches in each frame.

799 Fluorescence traces for individual cells were extracted from the motion-corrected images using the Constrained Non-negative
800 Matrix Factorization (CNMF) algorithm (Pnevmatikakis et al., 2016) The initialization of spatial components for CNMF and the
801 classification of these components into cell-like and non-cell-like categories followed previously published methods (Koay et al.,
802 2020) All components were manually reviewed and reclassified if necessary. The ΔF/F for each cell was calculated using the
803 modal fluorescence value from a 3-minute windows as the baseline. Since CNMF only identifies cells with calcium activity during
804 imaging, cells that did not exhibit activity were not included in this analysis. Shape correlation was used as a metric for cell
805 isolation quality. The shape correlation is the correlation of activity in each pixel with the average, reference image for that cell.
806 A change in the shape correlation time course over a session could reflect some change in the measurement, such as a shift in
807 the focal plane. Any units with a clear temporal trend in the shape correlation time course were excluded during manual
808 curation.

809 We detected transients as periods where the ΔF/F signal exceeded 2 standard deviations in the positive direction for at least 4
810 frames (133ms). For all analysis we zeroed out all non-transient portions of the individual time-traces for each neuron. Our
811 expression of the GCaMP6f was sparse compared to other systems when contamination of signals from neighboring neurons was
812 not a concern, so false attribution of multiple cells contributing to the same ROI was unlikely. For all analysis, we excluded cells
813 that did have at least one transient in at least 5% of trials within a session.

814

Univariate encoding model.

The encoding model was formulated as a multiplicative session modulation by odor modulation model. The average calcium response on each trial, following odor delivery, y was modeled by a temporal kernel s across the trials in a session x_{trials} multiplied with a characteristic odor response vector b which has a different value depending on the odor presented that trial x_{odor} .

$$y = s(x_{trials}) \cdot b x_{odor}$$

The session kernel s was a three-parameter compound gaussian with a center position, and separate rise and fall width parameters. The rise and fall parameters allowed the shape to smoothly transition from a step function to fully flat. The characteristic odor response vector b was a positive valued vector. The multiplication of these two terms would give the underlying latent excitability value for each cell per trial. We noticed that residuals of this model fitted to the trial-average calcium data showed strong heteroskedasticity, the variability of the average calcium response increased with the local mean; this is expected since the underlying spiking process is roughly Poisson. We applied the Anscombe transform to the data to stabilize the residual variance, which could then be appropriately modelled with a gaussian distribution for likelihood fitting purposes.

We used the MATLAB function *fminunc* to minimize the negative log likelihood. In order to classify cells, we estimated 5 models for each cell; (1) a mean only model, (2) an odor coding only model, (3) session-time only, (4) the full session-time x odor model and (5) a separate session-time kernel for each odor model.

We used lowest value of the Bayesian information criterion to select the best fitting model for each neurons. We excluded fitted models that showed narrow session kernel(s) width (less than a cumulative sum of the kernel of 6 trial units, which corresponds to a symmetrical gaussian kernel with ~6 trials s.d.); this constraint corresponded to overfitting of single isolated transients for sparsely firing neurons. Since very few cells (2%) were best explained by the separate model, and those which were, showed a high degree of overlap of the separate kernels or low rates of firing, we excluded this category from further analysis and chose the model with the next lowest BIC score for these neurons.

We defined the time modulation index as the minimum of the two estimated normalized shape parameters for each cell, with a low value near zero indicating a strong temporal modulation, and a value of one being no temporal modulation. We calculated the degree of odor modulation by computing the lifetime sparseness (Vinje and Gallant 2000).

$$Sparseness = \frac{1 - \frac{1}{N}\frac{(\sum_i b_i)^2}{\sum_i b_i^2}}{1 - \frac{1}{N}}$$

Where b_i is the estimated odor response vector and N is the number of odors ($N = 4$).

Decoding model

For the windowed decoder, the three task factors we decoded where, the *odor* presented to the animal, the location of the *barrier* in the current block, and the upcoming *turn* decision of the animal. We used a linear discriminant decoder for these three factors. For each cell and trial, we took the average ΔF/F value in 27ms wide, non-overlapping temporal windows aligned to the start of the fixation. Odor and turn decoders were cross-validated using a n-fold random selection, where n is the number of blocks within a session. Barrier location was cross-validated using temporal partition scheme, where for each fold, a single block was selected as a test and the remaining blocks as training. This partition scheme was specifically designed to ensure the decoder was capturing barrier location signals that generalized across experimental blocks rather than block-specific characteristics of the data (Harris 2021).

To compare the performance of classifiers to chance, we constructed 400 null shuffles for each factor. For the odor and the turn factors, the trial identity was randomly permuted within the current block of trials, this ensured that the nulls reflected the local statistics of the real data. For the barrier location factor, we created a shuffled sequence of barrier location blocks that respected the original sequence (i.e. same number of each condition, and no repeats of the same barrier location on adjacent blocks). For the turn decoder we only considered trials from the two *changed turn goals,* we excluded the *fixed turn* goals because the animal's behavior was almost entirely the correct turn for each odor, and so any odor coding could be mistaken for turn coding.

860 For the block lagged odor decoder, we trained a separate linear decoder for each block for the odor identity and evaluated each
861 decoder on the test data from all other blocks; we averaged the decoder performance as a function of the block lags between
862 training and test.

863 For the block identity decoder, we considered the average active of active cells in the whole trial time window, and used a
864 multiclass categorical linear discriminant decoder. The decoder was cross validated with a 5-fold partition, where every trial in
865 each block was randomly assigned to one of 5 equal sized partitions. This ensured the that every training and test partition had
866 some trials from every block.

867

868 **Representational similarity analysis**
869 For each cell and every trial within a session we calculated a population vector of the average ΔF/F response for each neuron for
870 the whole trial following the presentation of the odor. A representational similarity/distance matrix was then constructed by
871 taking the population vector (normalized to zero to one for each cell) Pearson's correlation between every pair of trials. The
872 average neural distance as a function of trial lag was calculated in bins of 20 trials. The comparison across animals for the fixed vs
873 changed action odor/goals was made based on the first trial lag bin. Since the values in the representational similarity matrix are
874 not independent observations, we only considered the similarity of adjacent trials for statistical testing. We took the population
875 vector correlation between two adjacent trials and modeled it as a multiple, mixed effects regression model across the whole
876 dataset. For the fixed vs change analysis, we selected only adjacent trials of the same odor; the fixed effects were the category of
877 odor/goal (fixed or changed) and the absolute time difference between trials; the random effects structure included random
878 intercepts for animal identity and session. For the turn error analysis, we selected only adjacent trials of the same odor for
879 changed odor/goals (since animals made few errors for fixed odor/goals); the fixed effects we examined were: if the animal
880 made a turn error on the first trial and the absolute time difference between the trials. The random effects structure included
881 random intercepts for animal identity and session.

882 For the barrier location decoding analysis, we used the neural representational similarity matrix and the category similarity
883 matrix for the closed barrier location for each trial (zero if same, one if different) We then shuffled the closed barrier location in
884 the same manner as the linear decoding analysis, shuffling the closed barrier location ID for each block, but respecting the
885 structure of the task and calculated a series of null category similarity matrices. Each of these was correlated to the neural
886 representational similarity matrix to generate a distribution of correlation scores under the null that the true value could be
887 compared to. In this way we could test for representations of closed barrier that might drift, but showed some above average
888 similarity when comparing same and different closed barrier location blocks.

889 **Neural network modeling**
890 We modelled the learning of the correct action for a given stimuli following a block change with a single layer feedforward
891 network. The input layer represents a state input, reflective of the state representation in CA1. The input layer x projects to two
892 output units y through a weight matrix w. The activity in the output units are the weighted sum of the activity of the input layer.

893

894
$$y_j = \sum_i x_i w_{ij}$$

895

896 The values of the action units are passed through a softmax function, that generates a probability of taking a left or right action.

897

898
$$P(action_j) = \frac{e^{(y_j / T)}}{\sum_n e^{(y_n / T)}}$$

899

900

901

902 The weights were adjusted with a supervised Hebbian rule

903

904

$$w_{ij} \leftarrow w_{ij} + \alpha x_i y_j$$

905

906

$$j = \begin{cases} action\ taken, & correct\ response \\ action\ not\ taken, & incorrect\ response \end{cases}$$

907

908   w_{ij} is the synaptic weight between the ith input unit and the the jth output unit, α is the learning rate. This rule corresponds to a
909   strengthening of the synapse for the correct response regardless of the outcome. Such a rule could be implemented by eligibility
910   traces established during action selection. After the action has been taken and the outcome evaluated, If the action taken was
911   correct, the action taken index is potentiated, if the action was incorrect then the action not taken index is potentiated
912   (counterfactual learning). The weights of synapses were restricted to grow no larger than 1 (the maximum value of the input
913   units), to ensure that weights did not continue to increase indefinitely. There is no decay term nor is there any re-normalization.
914   The was no reduction of any weights for the incorrect action whether it was taken or not.

915   The activity of the input layer was determined by a generative model that was based on the dynamics of the real neural data. For
916   each simulated cell we selected a session response profile by resampling the from the temporal kernel shapes from the encoding
917   model for real neurons and randomly generated a center point for the response within that session. This gave a predetermined
918   response kernel which spanned the length of the simulated session and was indexed by the number of total trials in a session.

919   In the **time-driven** condition, the state index n of the response kernels advances by 1 on every trial:

920

$$n \leftarrow n + 1$$

921   In the **error-driven** condition is only advanced following an incorrect action

922

$$n \leftarrow n + \gamma$$

923

$$\gamma = \begin{cases} 0, & if\ correct\ response \\ s, & if\ incorrect\ response \end{cases}$$

924   Since the total number of errors is not known in advance, the step size s must be adjusted dynamically to ensure that the full
925   response profile is covered by the end of the session. We used the following procedure to determine s; s is initialized with a
926   value taking into account the number of errors in real animal data to cover the full response profiles, we first run an estimate the
927   expected number of errors by running. If the estimated step size leads to premature completion or incomplete coverage, it is
928   iteratively refined. This ensures that the average, long run response profiles of simulated neurons in the error-driven condition
929   match the **time-driven** condition and are representative of the neural data in both cases.

930   We wanted to isolate the contributions of representational drift to the learning profiles, so we restricted the simulation to cells
931   with single odor representations (which are present in many actual cells recorded); cells were simulated with a binary one-hot
932   response to either odor, which was then multiplied with the temporal kernel previously selected. This ensured that behavioral
933   effects we saw in the simulation would be the result of the modulation of the drift rather than any correlation of the response
934   vectors for each odor.

935   We ran the simulation with a simplified version of the task that only featured two learned odors (since the correct action for
936   fixed odors would simply be learned initially and be maintained throughout the session), with the three closed door positions.
937   For each session we presented 360 trials in blocks of 72 trials, this was to ensure that the network more completely leaned the
938   rule for the current block, before having to switch to a new block. We simulated 200 sessions to generate the learning curves.

939   The softmax inverse temperature ($T = 5$) and learning rate ($\alpha = 0.4$) were adjusted so there was a suitable match to the actual
940   animal behavioral data. For the simulated engram reactivation condition, we re-instantiated the activity on each trial of a session
941   through the final weight matrix w that was learned after training and recorded if the network generated the appropriate action
942   for that given trial. The final engram reactivation score was the proportion of these trials that were correct.

943   For the Hebbian plus depression condition, we allowed the weights to adjust downwards with the same learning rate for the
944   incorrect action. The softmax inverse temperature ($T = 10$) and learning rate ($\alpha = 0.1$) were adjusted for this condition to give a
945   similar learning profile following block changes.

946