

The Same Hippocampal CA1 Population Simultaneously Codes Temporal Information over Multiple Timescales

Highlights

- Time cell sequences recorded with Ca^{2+} imaging change over minutes and days
- Second-resolution temporal information is retained despite population drift
- Time cell ensembles carry information about time over seconds, minutes, and days
- This population could organize temporal encoding over these timescales

Authors

William Mau, David W. Sullivan,
Nathaniel R. Kinsky,
Michael E. Hasselmo, Marc W. Howard,
Howard Eichenbaum

Correspondence

wmau@bu.edu

In Brief

Episodic memories span timescales of seconds, minutes, and days. Mau et al. use calcium imaging to longitudinally monitor cell sequences in hippocampal CA1. Bayesian decoder analyses show that the same population of neurons carries information about time across all three scales.



The Same Hippocampal CA1 Population Simultaneously Codes Temporal Information over Multiple Timescales

William Mau,^{1,2,4,*} David W. Sullivan,¹ Nathaniel R. Kinsky,^{1,2} Michael E. Hasselmo,¹ Marc W. Howard,¹ and Howard Eichenbaum^{1,3}

¹Center for Memory and Brain, Boston University, Commonwealth Avenue, Boston, MA 02215, USA

²Graduate Program for Neuroscience, Boston University, Commonwealth Avenue, Boston, MA 02215, USA

³We dedicate this paper to Howard Eichenbaum, who passed away July 2017

⁴Lead Contact

*Correspondence: wmau@bu.edu

<https://doi.org/10.1016/j.cub.2018.03.051>

SUMMARY

It has long been hypothesized that a primary function of the hippocampus is to discover and exploit temporal relationships between events. Previously, it has been reported that sequences of “time cells” in the hippocampus extend for tens of seconds. Other studies have shown that neuronal firing in the hippocampus fluctuates over hours and days. Both of these mechanisms could enable temporal encoding of events over very different timescales. However, thus far, these two classes of phenomena have never been observed simultaneously, which is necessary to ascribe broad-range temporal coding to the hippocampus. Using *in vivo* calcium imaging in unrestrained mice, we observed sequences of hippocampal neurons that bridged a 10 s delay. Similar sequences were observed over multiple days, but the set of neurons participating in those sequences changed gradually. Thus, the same population of neurons that encodes temporal information over seconds can also be used to distinguish periods of time over much longer timescales. These results unify two previously separate paradigms of temporal processing in the hippocampus that support episodic memory.

INTRODUCTION

The mammalian hippocampus is critical for linking spatiotemporally defined events to form episodic memories [1, 2]. Numerous experiments in both rodents and humans have shown that representations of temporal delays or temporal order are generated in the hippocampus [3–8] (for reviews, see [9–11]). In a particularly striking example, CA1 pyramidal cells reliably spike in sequence during defined temporal intervals within experimentally imposed delays of up to 20 s [4–6, 8, 9]. Sequences of this nature had been predicted in computational models of hippocampal function [12, 13], suggesting that the “time cells” that comprise these sequences provide temporal information about successive

events at a behavioral timescale (i.e., “microtime”) [14]. In support of this, time cell sequences differentiate goal locations [6], odors [15], tones, and behavioral decisions [16]. After learning, time cell sequences are required for appropriate memory for past events, supported by evidence that interruption of these sequences impairs performance in memory tasks [17, 18]. Furthermore, CA1 temporal structure is compromised in the time periods before erroneous decisions [6, 15, 16, 19], reflecting their importance in maintaining task-relevant information about the past.

Time cell sequences span seconds, making them well suited to encode temporal information in microtime, but for timescales exceeding seconds, the hippocampus appears to utilize a different mechanism. That is, neural representations of memories occurring minutes to days apart (“macrotime”) diverge in order to support accurate retrieval. For instance, in a recent human functional imaging study, the neural similarity of the activation in the anterior hippocampus evoked by remembered events tracked objective distance in time over the scale of hours, days, and weeks [20]. In analogous animal studies, the hippocampus exhibits population “drift” whereby neuronal outputs gradually and continuously change. For example, the spatial code in CA1 has been found to progressively differ with increasing temporal distance under constant conditions [21–23]. One purported role for population drift is the timestamping of mnemonic representations via indexing within neuronal engrams that continuously turn over [22, 24, 25]. That is, memories of events within a certain temporal window are allocated to subpopulations of cells, with memories that occur proximally in time residing in overlapping populations [26–29]. Those populations ultimately underlie memory representations during subsequent retrieval [29, 30]. This so-called temporal context model has powerful implications for how neural circuits distinguish between events occurring far apart in time [31, 32] and would require macroscopic-timescale basal dynamics in the brain. Indeed, population drift is required to support memory tasks with a temporal demand [19, 33].

Theories of hippocampal function have hypothesized its role in binding events into a spatiotemporal configuration for memory storage and retrieval [1, 10, 11]. Hippocampal time cell sequences, in conjunction with spatial responses (i.e., place cells) [34], are thought to fulfill this role and represent events for informing future behavior [9]. Thus, it is imperative to examine the long-term activity of cell sequences representing temporal order



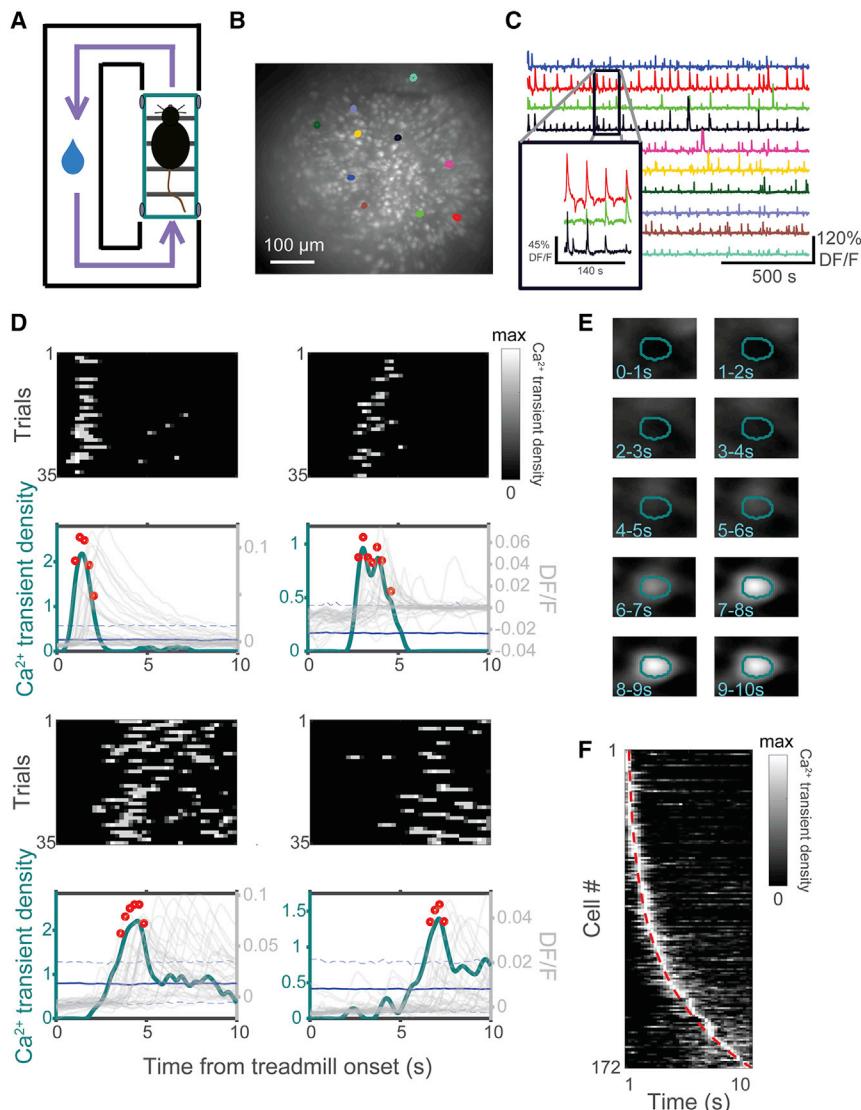


Figure 1. Sequentially Activated Time Cells Were Observed Using Calcium Imaging

(A) Task schematic. Mice run for 10 s on a motorized treadmill then turn left to retrieve a sucrose water reward at a well.

(B) Example of the field of view through an implanted lens aimed at CA1, depicted as the maximum temporal projection of fluorescence activity. Ten neuron ROIs are highlighted.

(C) Fluorescence traces of highlighted neurons in (B). Inset, zoom.

(D) Activity patterns of four representative time cells. Top plots are Ca^{2+} transient density maps, aligned to treadmill onset. Bottom plots are temporal receptive fields, averaged across treadmill runs (teal). Also shown are receptive fields of time-shuffled data (blue, solid) with 95% confidence intervals (blue, dashed), regions where empirical data are statistically significant from time-shuffled data (red dots), and fluorescence traces from individual treadmill runs (gray).

(E) Trial-averaged time lapse images of last cell in (D).

(F) Receptive fields (grayscale) of all classified time cells in one mouse during one session, sorted by location of field peaks (red line).

See also Figures S1 and S2 and Video S1.

in particular, as this would elucidate how the brain encodes time along multiple scales. One popular prediction is that the hippocampus must distinguish between events occurring minutes or hours apart via population drift [21, 24], but this has not yet been explicitly observed in hippocampal time cell sequences. Fortunately, recent advances in imaging technology permit longitudinal recording of brain regions at cellular resolution [35, 36], allowing us to track long-term evolution of these sequences. In this study spanning four days, we demonstrated that on each day, time cell sequences retain a semblance of the previous day's structure while also systematically varying, producing temporal signals over multiple timescales.

RESULTS

Behavioral Task and Epifluorescence Imaging of Calcium Transients

Due to the strong place selectivity of hippocampal neurons and the direct relationship between space and time, it is necessary

to control for spatial variables when observing temporal sequences. To do this, we utilized forced treadmill running as a method for clamping position while measuring neural activity as a function of time, as done in previous studies [5, 18]. Mice were trained to traverse a rectangular track followed by running in place on a motorized treadmill for 10 s at a constant velocity to receive sucrose water reward after traversing an additional part of the track (Figure 1A). Thus, a "trial" in this study refers to a 10 s treadmill-

running interval. We virally transfected dorsal CA1 neurons with the calcium indicator GCaMP6f using a synapsin promoter (Figures S1A and S1B) and used *in vivo* one-photon microscopy to image calcium (Ca^{2+}) transient activity and reliably capture the activity patterns of hundreds of cells simultaneously [23, 35, 37] in each of four adult mice during laps around the track and treadmill running (Figures 1B and 1C). Optical recording began after training ensured that mice reliably ran ~30 laps per day at a constant velocity on the treadmill. To extract fluorescence traces and infer Ca^{2+} transient events, we utilized an image segmentation algorithm designed to minimize optical crosstalk between overlapping neurons (see STAR Methods; D.W. Sullivan et al., 2017, Soc. Neurosci., abstract). To identify sequentially active time cells, we aligned Ca^{2+} transient traces to the treadmill's onset and averaged across trials to characterize putative temporal receptive fields. Fields for each cell were then compared to surrogate receptive fields generated by shuffling Ca^{2+} transient time stamps along the 10 s delay for each run. We identified a large population of cells whose receptive fields were statistically

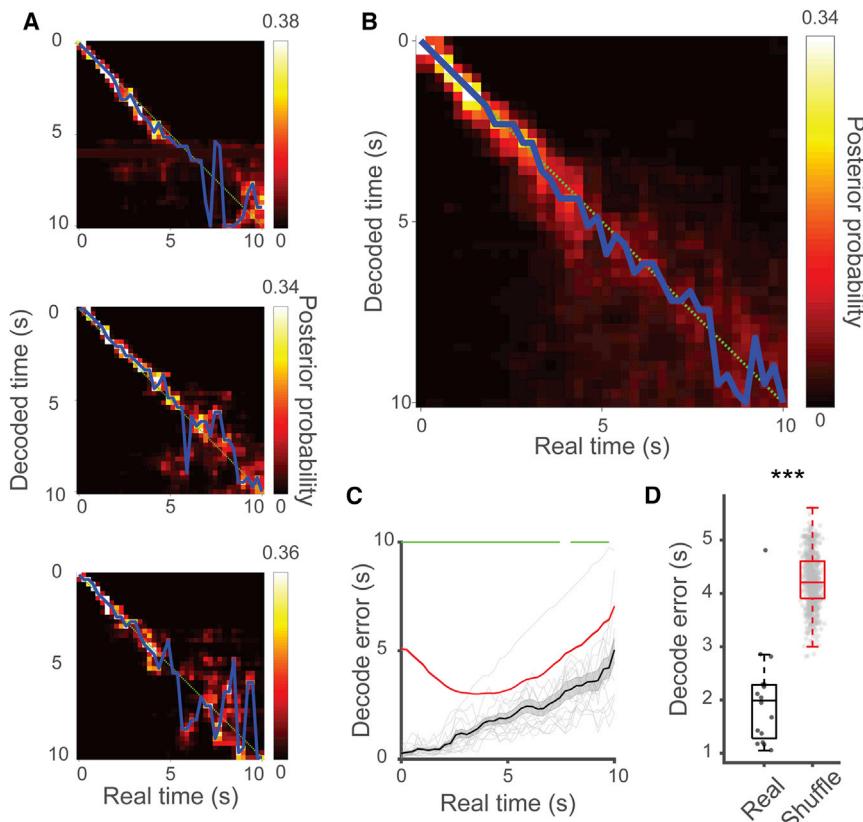


Figure 2. Time Cell Sequences Contained Information about Relative Time on the Scale of Seconds

(A) Decoding results of individual treadmill runs. Color bar indicates posterior probabilities, and blue lines denote decoder's most confident estimation. Green lines signify hypothetical perfect decoding.

(B) Decoding results of all sessions, averaged.

(C) Average decoding error as a function of elapsed time. Chance (red) was calculated by shuffling cell identity. Decoder performs better than chance for the majority of the temporal delay (green, $p < 0.05$). Data are represented as means \pm SEM.

(D) Average decoding error for each mouse and session compared to chance (Mann-Whitney U test; $p < 7.5 \times 10^{-10}$).

See also Figure S3.

significant compared to the randomized fields ($p < 0.01$, $n = 1,111$ time cells/10,315 neurons recorded over four days, 10.8% of total population of cells that fired at least 1 Ca^{2+} transient, Figures 1D–1F; see also Video S1 and Figure S1C). As described in previous literature, these cells reliably fired during specific moments relative to the start of the treadmill run (Figures 1D and 1E) and spanned the entirety of the 10 s delay (Figure 1F). Similar to observations in previous studies [4–6], the distribution of temporal receptive fields along the delay was skewed toward the beginning of the delay onset, possibly reflecting the relative salience of the treadmill turning on and scalar representation of time [31]. A sizeable percentage of time cells also exhibited place fields ($n = 164/1,111$ time cells, 17.0% of the time cell population, example shown in Figure S2). The overlap in time and place cells was not significantly different from chance, calculated via random sampling (Mann-Whitney U test, $p = 0.15$), suggesting that temporal and spatial information are interchangeably encoded by the neuronal population [14].

Reconstructing Temporal Information from Ordered Neural Firing

Though these cells exhibited temporal firing fields, a separate question concerns whether temporal sequence information was embedded in the ensemble activity. To test this, we investigated our ability to reconstruct elapsed time from Ca^{2+} transient population vectors derived from the time cell ensemble. For each session, we trained a naive Bayesian classifier with all the time cells' Ca^{2+} transient activity on that day from a random 50% of the trials and used the trained classifier to predict elapsed time

relative to the treadmill onset from the activity on the other 50% of the session. We found that we were able to accurately decode elapsed time on individual trials (Figure 2A) and over all sessions (Figure 2B). Interestingly, the classifier's error increased as a function of time relative to the treadmill onset (one-way ANOVA, $F_{15,639} = 16.79$, $p < 3.4 \times 10^{-37}$; Figure 2C), reflecting the uncertainty associated with scalar representation of time [38].

To assess the performance of our classifier compared to chance, we trained it on a dataset with shuffled cell identities and found that this dramatically increased the classifier's error (Mann-Whitney U test, $p < 7.5 \times 10^{-10}$), reinforcing the idea that the order of neuronal firing is necessary to generate accurate representation of temporal information (Figures 2C and 2D). This effect was extremely robust, and even a small percentage of cells contributed to encoding temporal information (Figure S3).

Evolution of Time Cell Sequences on the Scale of Minutes

Although numerous studies have primarily characterized population changes across days [21–23], it is also informative to observe these changes at a finer temporal resolution. We hypothesized that if population differences are apparent at time points hours apart, they might also be visible at time points minutes apart. Therefore, we explored how the time cell ensemble evolved over minutes by tracking its activity profile throughout the course of a recording session. To quantify when cells were actively encoding temporal information during the session, we identified the trial numbers on which a cell fired in its receptive field and computed the average of those trial numbers and then normalized by the number of trials in that session. Using this method, cells that fired on every trial would receive a “within-session trial bias score” of 0.5, whereas cells that fire only early in the session would have lower scores and cells that fire only later in the session would have higher scores. Scores were centered around 0.5 (0.53 ± 0.0029), but extremes in either

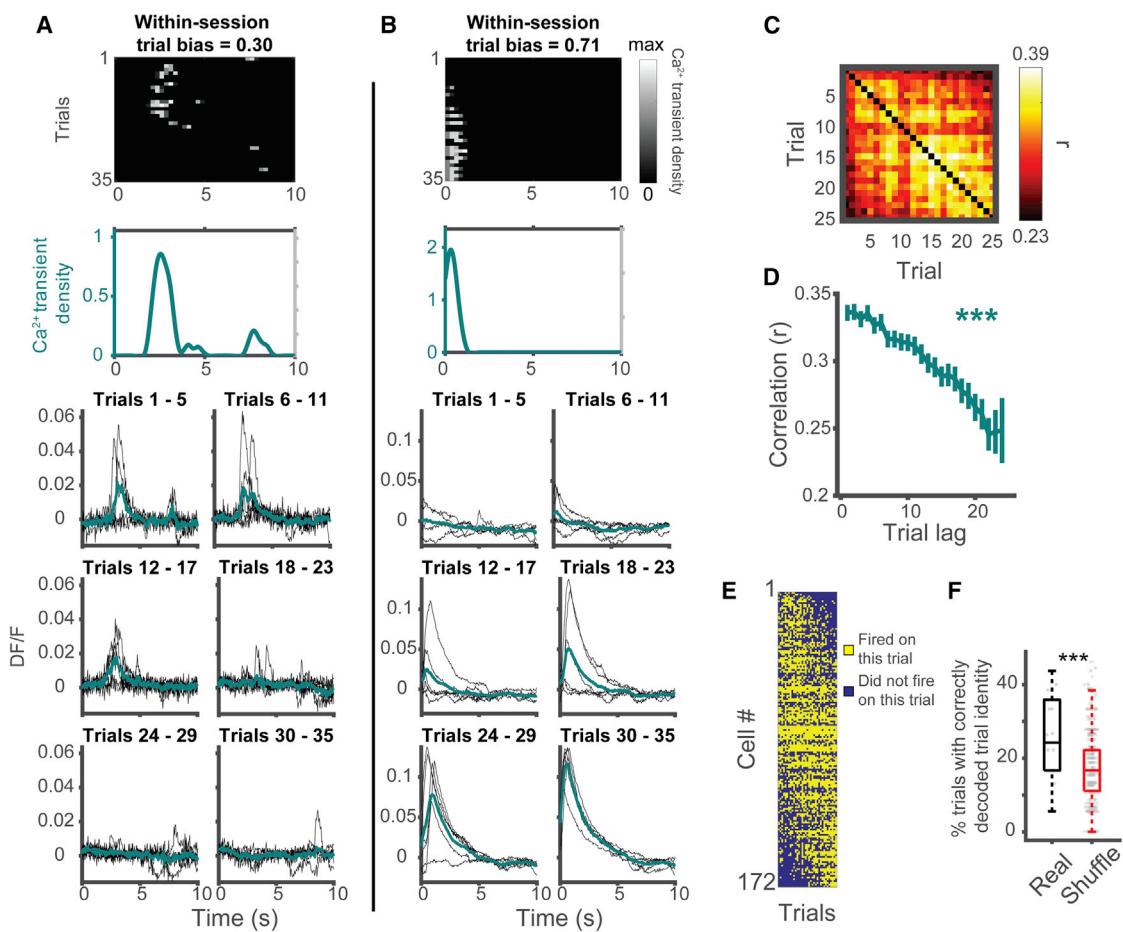


Figure 3. Time Cells Encoded Information about Elapsed Time on the Scale of Minutes

(A) Activity profile of example time cell active early in the session. Top plots are Ca^{2+} transient density maps and trial-averaged receptive field. Bottom plots are fluorescence traces from individual runs (black), divided into treadmill run blocks and averaged within a block (teal).
 (B) Same as (A) but for a time cell active late in the session.
 (C) Trial-by-trial correlation matrix of fluorescence traces.
 (D) Correlation as a function of trial lag, averaged from off-diagonals of matrix in (C). Data are represented as means \pm SEM.
 (E) Trial-by-trial activity of time cells during one session. Yellow indicates trials where that cell fired in its receptive field. Blue indicates trials where it did not. Data are sorted by within-session trial bias scores.
 (F) Treadmill run block decoder performance compared to chance (shuffling run identity; Mann-Whitney U test; $p < 9.0 \times 10^{-6}$).
 See also Figure S4.

direction were also prevalent, as reflected in scores that were lower and higher than expected by chance compared to a distribution derived from randomized trial numbers (Figure S4C). We identified cells that were disproportionately active earlier in the session (permutation tests, $p < 0.05$, $n = 109$, 9.8% of the time cell population; e.g., Figure 3A) or later in the session ($n = 167$, 15.0% of the time cell population, e.g., Figure 3B). The proportion of early-active cells versus late-active cells was not significantly different (Mann-Whitney U test $p > 0.92$). This trial-modulated activity was not a result of shifting the plane of focus on the microscope because Ca^{2+} transients were still detected in early cells at later time points on the track, but not during the delay (Figure S4A) and same for late cells at early time points (Figure S4B). This indicates that, despite the treadmill running task being highly familiar, the hippocampus nonetheless showed changes in its activity patterns, with cells

forming and losing temporal receptive fields throughout each session.

Our single-cell observations prompted us to investigate whether the hippocampal population as a whole exhibited global changes in temporal coding over the course of a session. To do this, we performed pairwise correlations between fluorescence traces on single trials for each cell and then averaged across all cells to find the global population similarity for each pair of treadmill runs (Figure 3C). Trials that occurred farther apart in time became gradually more decorrelated, revealing continuous population drift over the entire recording session (one-way ANOVA $F_{23,4799} = 8.77$, $p < 8.8 \times 10^{-30}$; Figures 3D and 3E). The time cells participating in the sequence also changed as a function of time (Figure S4D). To test the possibility that temporal information on the scale of minutes could be encoded in this systematic variance, we designed another naive Bayes classifier to

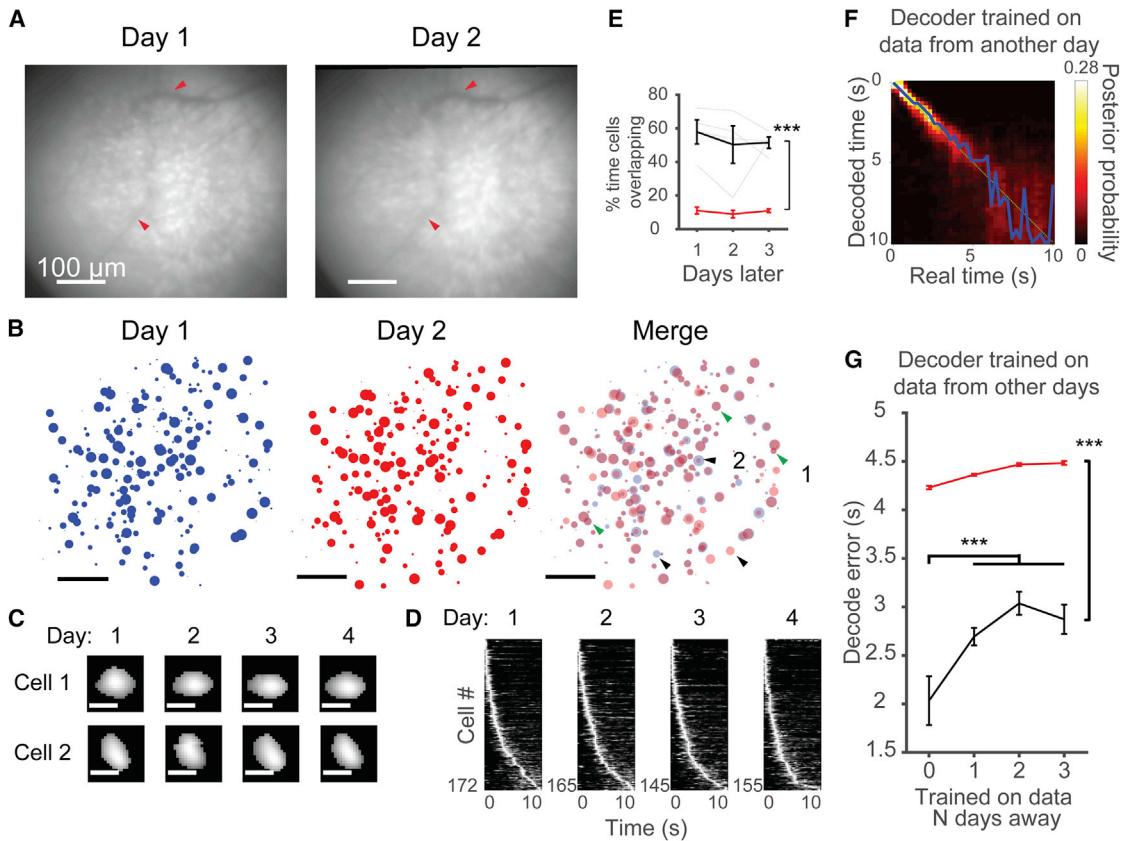


Figure 4. Time Cell Sequences Were Stably Recorded over Days

(A) Example fields of view for image alignment from the same mouse, on different days. Visible vasculature is indicated by red arrows.

(B) Time cell ensembles on two consecutive days (left and middle), with individual ROIs sized by temporal position in time cell sequence (larger, later in sequence) and overlaid (right). Green arrows indicate cells with similar temporal tuning curves across the two days; black arrows indicate otherwise. The scale bars represent 100 μm.

(C) Cell masks of neurons marked in (B) over all four days. The scale bars represent 10 μm.

(D) Ensemble plots of time cell ensembles, filtered day by day. Rows in each panel represent different neurons. For rows representing the same neuron, see Figure 5B.

(E) Ensemble overlap (black) as a function of temporal distance compared to chance (red; two-way ANOVA $F_{1,1211} = 611.88$, $p < 0.001$; post hoc Tukey's HSD test, $p < 0.001$). Gray lines indicate separate mice. Data are represented as means ± SEM.

(F) Decoded output of Bayesian classifier trained and tested on different days. Plotting conventions are the same as those in Figure 2.

(G) Seconds-level decoder performance from training decoder on data from a day different from the test set. Decoder error (black) is significantly below chance (red) for all temporal distances here (two-way ANOVA $F_{1,2039} = 483.19$, $p < 0.001$; post hoc Tukey's HSD tests, $p < 0.001$). Decoder performs better when trained on data from the same day (post hoc Tukey's HSD test; $p < 0.001$). Data are represented as means ± SEM.

See also Figure S5.

decode approximate trial number (trial blocks with each session split into six blocks, ~5 trials per block) from the collective Ca^{2+} transient activity of the time cell ensemble on each treadmill run. Again, we trained this classifier on 50% of each session's trials and then tested on the other 50% and asked it to predict which blocks those trials belonged to. The classifier was able to predict trial blocks significantly better than chance (calculated by shuffling trial blocks; Mann-Whitney U test, $p < 9.0 \times 10^{-6}$; Figure 3F). Different-sized trial blocks were tested with similar results (see STAR Methods). This demonstrates that temporal information on the order of minutes (across trials) can be extracted from population drift occurring over the course of a recording session in the same subset of neurons that also encode sequential structure within each trial.

Longitudinal Tracking of Time Cell Sequences

Next, we sought to define how the time cell ensemble developed over macroscopic time on the order of days. We exploited the ability of *in vivo* calcium imaging to track neurons over long time-scales and examined time cell sequences longitudinally. Across-day cell registration was performed by first aligning the minimum projection of the field of view for each pair of days via rigid translations and rotations, utilizing vasculature as landmarks (Figure 4A; see also Figures S5A–S5D). Then neuronal regions of interest (ROIs) on one day were matched to the closest ROI on another day based on distance between ROI centroids (all matches $<3.3 \mu\text{m}$ away; Figure 4B). After determining which cells were the same across days, we were able to visualize time cell ensembles over the duration of the experiment (Figure 4C).

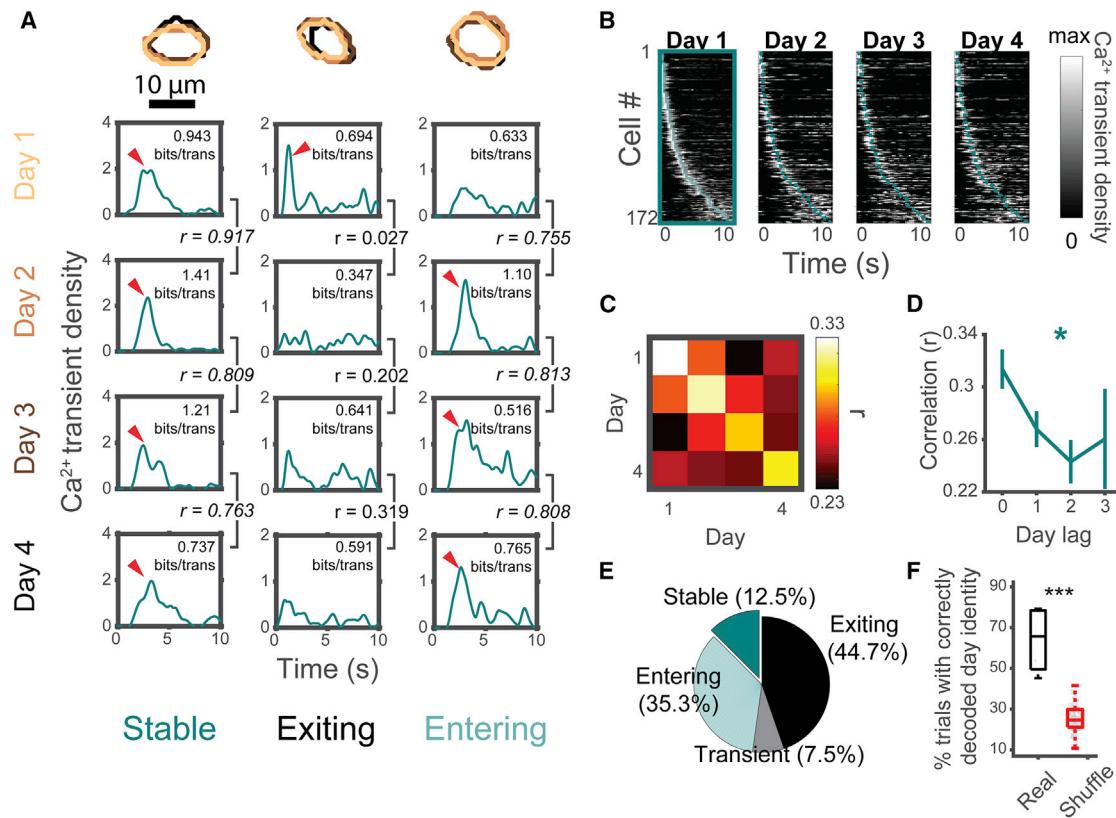


Figure 5. Time Cell Sequences Carried Information about Relative Time on the Scale of Days

(A) Receptive fields of three example cells exhibiting different across-days dynamics with accompanying ROI masks across days (top). Red arrows denote significant temporal receptive field peaks. Also shown are temporal mutual information (bits per transient for each cell) and tuning curve Pearson correlations. Italicized coefficients indicate statistically significant correlations.

(B) Time cell ensemble on day 1 of one mouse across four days. Teal line outlines the peaks on day 1 across all successive days.

(C) Correlation matrix of population similarity for all day pairs. Each value in the matrix represents the grand average of population correlations between all trials in that day pair for all animals.

(D) Correlation as a function of day lag, data from (C). Data are represented as means \pm SEM.

(E) Proportion of time cells exhibiting stability characteristics described in (A); $n = 486$ unique time cells.

(F) Performance of Bayesian decoder trained to decode day compared to chance from shuffling days (Mann-Whitney U test; $p < 7.0 \times 10^{-4}$).

See also Figure S5.

Although a different subset of the CA1 population encoded relative time on the scale of 10 s each day (Figure 4D), there was substantial overlap in time cells on one day compared to time cells up to three days later (Figure 4E). This overlap was significantly different from chance as calculated by ensemble overlap when random cells were drawn from the population instead (two-way ANOVA $F_{1,1211} = 611.88$, $p < 0.001$; post hoc Tukey's honest significant difference [HSD] test, $p < 0.001$). We next inquired whether this partial overlap was sufficient to preserve temporal information. Indeed, training the Bayesian classifier on activity rate vectors from one day allowed us to accurately decode elapsed time within the 10 s delay interval one day later (Figure 4F). This was true even when the classifier was trained using data from three days prior (chance calculated by performance trained on data with shuffled cell identity; two-way ANOVA $F_{1,1223} = 206.35$, $p < 0.001$; post hoc Tukey's HSD test, $p < 0.001$; Figure 4G). Collectively, this evidence suggests that despite our observations that a different subpopulation of neurons participate in the time cell ensemble each day, the activity

of the sequence is preserved to allow for extraction of meaningful temporal information.

Evolution of Time Cell Sequences on the Scale of Days

After establishing that the time cell sequence remains sufficiently intact over days, we focused on the content of all cells that were classified as a time cell at any point during the experiment. Out of that pool of cells ($n = 486$ time cells), we characterized longitudinal sequence coding by correlating tuning curves relative to the treadmill run. Time cells that were consistently correlated across all sessions and had statistically significant temporal receptive fields (Pearson correlation $p < 0.01$, Bonferroni corrected) were considered stable (Figure 5A, left), whereas cells that lost or gained temporal firing fields between a pair of sessions were designated exiting (i.e., had a temporal receptive field on one day, but not the next, as defined by permutation tests described in Figure 1; Figure 5A, center) or entering (i.e., had no significant temporal receptive field on one day but gained one on the next; Figure 5A, right), respectively. A modest percentage

($12.5\% \pm 2.5\%$) of time cells were stable throughout the entire 4-day experiment, whereas most either entered ($35.3\% \pm 4.5\%$) or exited the ensemble ($44.7\% \pm 2.2\%$; **Figure 5E**). A minority of cells ($7.5\% \pm 2.0\%$) both lost and gained temporal tuning at least once over the experiment, and as such, their activity was considered “transient”. These observations were unlikely to be due to shifts in the focal plane due to consistently high spatial correlation of cell masks (**Figure S5C**) and virtually no change in ROI movement or orientation across days (**Figure S5D**). Furthermore, to address the possibility that entering and exiting cells might reflect erroneous across-days cell registration, we computed the ROI displacements of those cells across days and compared them to stable cells. ROI displacements of entering and exiting cells were indistinguishable from those of stable cells, discrediting the possibility that our registration threshold allowed inaccurate matching of different cells across days (**Figure S5E**). Interestingly, we also found time cells with higher temporal information were more likely to be stable over two consecutive days (**Figure S5F**), which parallels previous findings that reliability of dendritic branch spiking predicts place field stability [39]. Turnover of the activity of single cells contributed to the evolution of the time cell ensemble day by day, gradually introducing variance to the system (**Figure 5B**), while simultaneously, a reliable time signal persisted (**Figure 4G**). This population drift was quantified by correlating fluorescence traces in the same manner as in **Figure 3C** and then averaging across trials to calculate the overall level of difference between pairs of days (**Figure 5C**). Similar to our results across trials during single sessions, we found a significant and systematic decorrelation in ensemble activity across days (one-way ANOVA $F_{3,39} = 9.43$, $p < 0.025$; **Figure 5D**). Analogous to our minutes-timescale analysis, to determine whether macroscopic temporal information on the order of days was present in this population drift, we built a third type of classifier to predict on which day a sample of population activity occurred. Just as the minutes-timescale classifier was able to correctly identify trial blocks, our day-scale classifier could accurately distinguish between recording days (days 1–4) based on population activity better than chance (Mann-Whitney U test, $p < 7.0 \times 10^{-4}$; **Figure 5F**). These results imply that temporal information on the order of days can be derived from macrotimescale drift from the same population that encodes sequence order on the order of seconds.

DISCUSSION

We showed that time cell sequences spanning 10 s durations occur in CA1 neurons of mice running in place during a simple goal-seeking task (**Figure 1**). There was sufficient temporal information contained in those sequences for a Bayesian classifier to faithfully decode elapsed time (**Figure 2**). Interestingly, notwithstanding the lack of salient cues during our temporal interval, this information is preserved over multiple days, as we can successfully train and test a classifier to decode elapsed time with data collected on different days (**Figure 4**). Despite stability in its overall temporal structure, the time cell ensemble systematically varied with the passage of minutes (**Figure 3**) and days (**Figure 5**). This time-dependent variance similarly contained temporal information, in parallel with and on larger scales compared to the content of time cell sequences. Thus, minute-scale and day-

scale Bayesian classifiers were capable of accurately inferring temporal position within and between imaging sessions. Collectively, these data demonstrate that the hippocampus has the capacity to encode temporal information along multiple timescales in support of episodic memory.

Robustness and Fluidity of Sequential Firing over Days

A key finding of this study was that of a heterogeneous population of CA1 neurons that displaying diverse stability patterns over days. Although some cells retained stable temporal receptive fields throughout the experiment, others gradually acquired and lost firing fields. Our observations that a time cell can change its temporal receptive field while not completely disrupting the downstream sequence reveals flexibility in the CA1 temporal coding regime, ruling out a simple synfire chain model for time cell sequence generation [9]. Instead, the evidence presented here suggests a more complex system where CA1 continuously reassigns encoding responsibilities to distributed ensembles of cells during population drift. This mechanism is manifested in time cells that lost their temporal receptive fields. Such a strategy, termed “dropout” by neural network researchers, is utilized in artificial intelligence by randomly removing single units during encoding to prevent the network from becoming disproportionately dependent on particular neurons. Dropout had been previously explored and successfully implemented to prevent overfitting in artificial neural networks by the deep learning community [40]. Taking inspiration from this model, in a biological neural network, population drift might serve multiple different purposes—(1) to facilitate the formation of neural associations in a diverse and resilient population of neurons, (2) to timestamp neural events along an extended timeline, and (3) to assimilate experiences across multiple timescales.

Advantages of Neural Instability in an Unstable World: Drift as a Mechanism for Timestamping Events

Why might a system continuously vary its activity patterns at the population level? Intuition leads one to believe that stability would be preferable in order for neural circuits to preserve and promote consistent outputs at the behavioral level. Following this line of thought, early experiments focused disproportionately on how hippocampal ensembles remained stable over long time periods [41, 42]. However, recent studies have begun to appreciate the potential benefits and advantages of an “unstable” system [21, 22, 24, 25]. Such a strategy might enable dynamic allocation of memories to distributed networks of neurons for mentally separating experiences in time [21, 22, 27, 29] while simultaneously providing a mechanism for avoiding interference during retrieval [43]. From the results presented here, we concluded that this mechanism is prominent in a population of neurons that encode sequential information. Through Bayesian decoding analyses, we found that this was true for minute- to day-level timescales, opening up the prospect that sequences of events occurring at different time points could be disambiguated based on the active population. This disambiguation might be performed by a downstream reader interpreting the network state to infer time.

A question triggered by our findings is the mechanism by which population drift manifests and how it might subserve memory encoding. One conspicuous possibility is plasticity

via synaptic changes in the cellular network. Despite the reputation of the hippocampus for being a long-term memory storage unit, it is not uncommon to observe ample synaptic turnover. A recent imaging experiment has estimated the lifetime of CA1 dendritic spines in live mice to be only 1 or 2 weeks [44]. The impermanence of hippocampal synapses may be due to a perpetual rotation of cells constantly forming and eliminating potentiated connections [26]. These dynamics at the synaptic level seem to be modulated by expression of key plasticity proteins, such as cyclic AMP (cAMP) response element-binding protein (CREB), where cells with high levels of CREB are more likely to be potentiated and recruited into a memory trace [26]. CREB also modulates cellular excitability [45], and endogenous cycling of CREB in cells could explain emergence and decay of time cell activity over both minute and day timescales through their impact on time cell excitability. The time course of CREB phosphorylation, which occurs over minutes [46], is consistent with our observations of changes in time cell responses over similar timescales, and *stability* of time cell responses is most likely the byproduct of sustained synaptic potentiation mediated by CREB [26]. On the other hand, *changes* in time cell sequences over days could reflect competitive processes [29], where cells with ramping CREB dominate over existing time cells.

CREB-induced excitability may underlie the emergence of firing fields in the hippocampus [47] as well as other structures [45]. Therefore, it is conceivable that the hippocampus routinely recruits neurons into sequential patterns to establish associative connections [48] based on CREB expression. Furthermore, recent findings of synaptic plasticity windows in CA1 that occur on the behavioral timescale [49] lend credence to the idea that constituents of time cell ensembles, which activate over seconds, could be linked in this manner, yielding a neural storage unit for sequential information across long timescales, as observed in our study.

Another possible mechanism for hippocampal time cell generation and the population's neural drift could originate from mathematical model utilizing a two-layer feedforward network [31, 50]. According to this model, one set of cells responds to a salient event and then decays exponentially, with different cells decaying with a range of time constants. A second set of cells receives and filters input from the exponentially decaying ensemble to generate sequentially activated time cells. The sequence extends over a range of times controlled by the range of time constants in the exponentially decaying population. Previous modeling work has suggested that a calcium-activated non-specific cation (CAN) current dependent on muscarinic receptor activation may be sufficient to generate the observation of drifting time cells presented in this study [51].

A Unified Framework of Event Sequence Coding in Hippocampus over Long Timescales

Although it has been shown that population drift serves to timestamp place cells in CA1 [21, 22], until now, no studies have shown that population drift also applies to sequence coding in the hippocampus. This finding is a novel demonstration of a unified representation of temporal order along many scales, which is critical for episodic memory. Here, we observed time cell ensembles that fired in sequence but also rode on top of a basal

and continuous population-level dynamic that changed with the passage of minutes and days. Drifting time cell ensembles synthesize different regimes of temporal coding in the hippocampus by describing a population of neurons that simultaneously reflects temporal information about microtime within a 10 s delay interval and much longer timescales of minutes and even days (macrotime). This framework could potentially allow events occurring in sequence (including episodic features beyond that of spatial features) to be encoded while simultaneously providing a signal for distinguishing broad temporal context within a common subpopulation of neurons [31]. In addition, sequential firing could enable these neurons to reactivate, generating predictions of the future to inform behavioral decisions [52].

Formation of Schemata via Integration of Experiences across Macrotome

The paradigm of continuous neural drift might also support the integration of novel information during learning. Accumulation of knowledge occurs as a function of time as organisms continuously sample their environment. In a psychological context, this evidence accumulation is harnessed for the assimilation of concepts into a pre-existing mental "schema" [53]. The biological basis of assimilation might rest in the merging of neural representations, most likely manifested in neural sequences, such as time cell assemblies. In support of this, new neurons become incorporated into established sequences during learning and sleep [48, 54]. These "incoming" neurons may be primed by the continuously shifting hippocampal network to encode potentially useful new data. Furthermore, by amalgamating neurons into a sequence network, this places the brain in an advantageous position to make associative links to pre-existing memories and thus form cognitive schemata [55, 56]. Indeed, memories may be linked physically by the overlap in ensembles encoding them [27, 29, 54]. Our study recorded longitudinally from time cells, which by their very nature are critical for representing temporally separated events [4]. In the resulting analyses, we presented a key piece of evidence for this mechanism of schema formation by demonstrating that individual time cells insert and remove themselves from existing sequences from previous days.

Outstanding Questions in Long-Term Sequence Representations

Persistence and variance of temporal information across long timescales support the idea that the hippocampus stores and modifies firing patterns to support memory. However, several questions remain. The constant flux of excitable neurons situates the hippocampus in an ideal position for integrating new information into existing schemata [48, 54], but this has yet to be explicitly demonstrated. Promising recent advances in imaging technology have permitted other groups to longitudinally track network states of various brain regions across macroscopic time [36], but few have addressed the evolution of hippocampal cell assembly sequences during learning. Given that sequence generation appears to be the default activity of the hippocampal network [57, 58], future investigation into this domain promises fruitful gains in knowledge about how learning is imprinted onto the neural substrate.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- **KEY RESOURCES TABLE**
- **CONTACT FOR REAGENT AND RESOURCE SHARING**
- **EXPERIMENTAL MODEL AND SUBJECT DETAILS**
 - Animal Subjects
 - Viral Constructs
- **METHOD DETAILS**
 - Stereotactic Surgeries
 - Treadmill Running Behavior
 - Freely-Moving Calcium Imaging and Mouse Tracking
 - Histology and Epifluorescent Microscopy
- **QUANTIFICATION AND STATISTICAL ANALYSIS**
 - Time Cell Selection
 - Within-Session Trial Bias Score
 - Population Correlations
 - Naive Bayes Classifiers
 - Statistical Tests
- **DATA AND SOFTWARE AVAILABILITY**

SUPPLEMENTAL INFORMATION

Supplemental Information includes five figures and one video and can be found with this article online at <https://doi.org/10.1016/j.cub.2018.03.051>.

ACKNOWLEDGMENTS

This work is dedicated to Howard Eichenbaum, who tragically passed away before this manuscript was completed. We thank Steve Ramirez, Jon Rueckemann, William Liberti III, Adam Miller, John Bladon, Ryan Place, Daniel Sheehan, Sam Levy, Catherine Mikkelsen, and Daniel Orlin for comments and discussion. We thank Daniel Salz and Columbus Instruments for help with the treadmill setup. We thank Lara Cardy and Vardhan Dani at Inscopix for their technical support. We thank Daniel Aharoni and the UCLA miniscope project for experimental and analytical advice. We thank UNC Vector Core for providing GCaMP6f virus and Inscopix for making the microendoscopy technique commercially available. We thank Daniel Orlin, Wing Ning, Denise Parisi, Sandra Jean Grasso, and Shelley Russek for their administrative support. We also thank Helen Fawcett, David Boas, and the NSF Neurophotonics Research Traineeship Program. This work was funded by NIMH R01-MH112169, R01-MH051570, and R01-MH095297.

AUTHOR CONTRIBUTIONS

Conceptualization, Methodology, and Investigation, W.M. and H.E.; Software, W.M., D.W.S., and N.R.K.; Formal Analysis, W.M.; Writing – Original Draft, W.M., M.W.H., and M.E.H.; Writing – Reviewing & Editing, W.M., D.W.S., N.R.K., M.W.H., and M.E.H.; Supervision, M.W.H., M.E.H., and H.E.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: February 18, 2018

Revised: March 16, 2018

Accepted: March 22, 2018

Published: April 26, 2018

REFERENCES

2. Scoville, W.B., and Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. *J. Neurol. Neurosurg. Psychiatry* **20**, 103–113.
3. Fortin, N.J., Agster, K.L., and Eichenbaum, H.B. (2002). Critical role of the hippocampus in memory for sequences of events. *Nat. Neurosci.* **5**, 458–462.
4. MacDonald, C.J., Lepage, K.Q., Eden, U.T., and Eichenbaum, H. (2011). Hippocampal “time cells” bridge the gap in memory for discontiguous events. *Neuron* **71**, 737–749.
5. Kraus, B.J., Robinson, R.J., 2nd, White, J.A., Eichenbaum, H., and Hasselmo, M.E. (2013). Hippocampal “time cells”: time versus path integration. *Neuron* **78**, 1090–1101.
6. Pastalkova, E., Itskov, V., Amarasingham, A., and Buzsaki, G. (2008). Internally generated cell assembly sequences in the rat hippocampus. *Science* **321**, 1322–1327.
7. Ezzyat, Y., and Davachi, L. (2014). Similarity breeds proximity: pattern similarity within and across contexts is related to later mnemonic judgments of temporal proximity. *Neuron* **81**, 1179–1189.
8. Modi, M.N., Dhawale, A.K., and Bhalla, U.S. (2014). CA1 cell activity sequences emerge after reorganization of network correlation structure during associative learning. *eLife* **3**, e01982.
9. Eichenbaum, H. (2014). Time cells in the hippocampus: a new dimension for mapping memories. *Nat. Rev. Neurosci.* **15**, 732–744.
10. Ranganath, C., and Hsieh, L.T. (2016). The hippocampus: a special place for time. *Ann. N.Y. Acad. Sci.* **1369**, 93–110.
11. Davachi, L., and DuBrow, S. (2015). How the hippocampus preserves order: the role of prediction and context. *Trends Cogn. Sci.* **19**, 92–99.
12. Levy, W.B. (1996). A sequence predicting CA3 is a flexible associator that learns and uses context to solve hippocampal-like tasks. *Hippocampus* **6**, 579–590.
13. Wallenstein, G.V., Eichenbaum, H., and Hasselmo, M.E. (1998). The hippocampus as an associator of discontiguous events. *Trends Neurosci.* **21**, 317–323.
14. Eichenbaum, H. (2017). On the Integration of Space, Time, and Memory. *Neuron* **95**, 1007–1018.
15. MacDonald, C.J., Carrow, S., Place, R., and Eichenbaum, H. (2013). Distinct hippocampal time cell sequences represent odor memories in immobilized rats. *J. Neurosci.* **33**, 14607–14616.
16. Terada, S., Sakurai, Y., Nakahara, H., and Fujisawa, S. (2017). Temporal and Rate Coding for Discrete Event Sequences in the Hippocampus. *Neuron* **94**, 1248–1262.e4.
17. Wang, Y., Romani, S., Lustig, B., Leonardo, A., and Pastalkova, E. (2015). Theta sequences are essential for internally generated hippocampal firing fields. *Nat. Neurosci.* **18**, 282–288.
18. Robinson, N.T.M., Priestley, J.B., Rueckemann, J.W., Garcia, A.D., Smeglin, V.A., Marino, F.A., and Eichenbaum, H. (2017). Medial Entorhinal Cortex Selectively Supports Temporal Coding by Hippocampal Neurons. *Neuron* **94**, 677–688.e6.
19. Manns, J.R., Howard, M.W., and Eichenbaum, H. (2007). Gradual changes in hippocampal activity support remembering the order of events. *Neuron* **56**, 530–540.
20. Nielson, D.M., Smith, T.A., Sreekumar, V., Dennis, S., and Sederberg, P.B. (2015). Human hippocampus represents space and time during retrieval of real-world memories. *Proc. Natl. Acad. Sci. USA* **112**, 11078–11083.
21. Mankin, E.A., Sparks, F.T., Slayyeh, B., Sutherland, R.J., Leutgeb, S., and Leutgeb, J.K. (2012). Neuronal code for extended time in the hippocampus. *Proc. Natl. Acad. Sci. USA* **109**, 19462–19467.
22. Rubin, A., Geva, N., Sheintuch, L., and Ziv, Y. (2015). Hippocampal ensemble dynamics timestamp events in long-term memory. *eLife* **4**, e12247.
23. Ziv, Y., Burns, L.D., Cocker, E.D., Hamel, E.O., Ghosh, K.K., Kitch, L.J., El Gamal, A., and Schnitzer, M.J. (2013). Long-term dynamics of CA1 hippocampal place codes. *Nat. Neurosci.* **16**, 264–266.

24. Chambers, A.R., and Rumpel, S. (2017). A stable brain from unstable components: emerging concepts and implications for neural computation. *Neuroscience* 357, 172–184.
25. Clopath, C., Bonhoeffer, T., Hübener, M., and Rose, T. (2017). Variance and invariance of neuronal long-term representations. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 372, 20160161.
26. Rogerson, T., Cai, D.J., Frank, A., Sano, Y., Shobe, J., Lopez-Aranda, M.F., and Silva, A.J. (2014). Synaptic tagging during memory allocation. *Nat. Rev. Neurosci.* 15, 157–169.
27. Cai, D.J., Aharoni, D., Shuman, T., Shobe, J., Biane, J., Song, W., Wei, B., Veshkini, M., La-Vu, M., Lou, J., et al. (2016). A shared neural ensemble links distinct contextual memories encoded close in time. *Nature* 534, 115–118.
28. Yokose, J., Okubo-Suzuki, R., Nomoto, M., Ohkawa, N., Nishizono, H., Suzuki, A., Matsuo, M., Tsujimura, S., Takahashi, Y., Nagase, M., et al. (2017). Overlapping memory trace indispensable for linking, but not recalling, individual memories. *Science* 355, 398–403.
29. Rashid, A.J., Yan, C., Mercaldo, V., Hsiang, H.-L.L., Park, S., Cole, C.J., De Cristofaro, A., Yu, J., Ramakrishnan, C., Lee, S.Y., et al. (2016). Competition between engrams influences fear memory formation and recall. *Science* 353, 383–387.
30. Liu, X., Ramirez, S., Pang, P.T., Puryear, C.B., Govindarajan, A., Deisseroth, K., and Tonegawa, S. (2012). Optogenetic stimulation of a hippocampal engram activates fear memory recall. *Nature* 484, 381–385.
31. Howard, M.W., and Eichenbaum, H. (2013). The hippocampus, time, and memory across scales. *J. Exp. Psychol. Gen.* 142, 1211–1230.
32. Howard, M.W., Fotedar, M.S., Datey, A.V., and Hasselmo, M.E. (2005). The temporal context model in spatial navigation and relational learning: toward a common explanation of medial temporal lobe function across domains. *Psychol. Rev.* 112, 75–116.
33. Jenkins, L.J., and Ranganath, C. (2010). Prefrontal and medial temporal lobe activity at encoding predicts temporal context memory. *J. Neurosci.* 30, 15558–15565.
34. O'Keefe, J., and Burgess, N. (2011). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res.* 34, 171–175.
35. Ghosh, K.K., Burns, L.D., Cocker, E.D., Nimmerjahn, A., Ziv, Y., Gamal, A.E., and Schnitzer, M.J. (2011). Miniaturized integration of a fluorescence microscope. *Nat. Methods* 8, 871–878.
36. Hamel, E.J.O., Grawe, B.F., Parker, J.G., and Schnitzer, M.J. (2015). Cellular level brain imaging in behaving mammals: an engineering approach. *Neuron* 86, 140–159.
37. Resendez, S.L., Jennings, J.H., Ung, R.L., Namboodiri, V.M.K., Zhou, Z.C., Otis, J.M., Nomura, H., McHenry, J.A., Kosyk, O., and Stuber, G.D. (2016). Visualization of cortical, subcortical and deep brain neural circuit dynamics during naturalistic mammalian behavior with head-mounted microscopes and chronically implanted lenses. *Nat. Protoc.* 11, 566–597.
38. Howard, M.W., Shankar, K.H., Aue, W.R., and Criss, A.H. (2015). A distributed representation of internal time. *Psychol. Rev.* 122, 24–53.
39. Sheffield, M.E.J., and Dombeck, D.A. (2015). Calcium transient prevalence across the dendritic arbor predicts place field properties. *Nature* 517, 200–204.
40. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. (2014). Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* 15, 1929–1958.
41. Thompson, L.T., and Best, P.J. (1990). Long-term stability of the place-field activity of single units recorded from the dorsal hippocampus of freely behaving rats. *Brain Res.* 509, 299–308.
42. Kentros, C., Hargreaves, E., Hawkins, R.D., Kandel, E.R., Shapiro, M., and Muller, R.V. (1998). Abolition of long-term stability of new hippocampal place cell maps by NMDA receptor blockade. *Science* 280, 2121–2126.
43. Colgin, L.L., Moser, E.I., and Moser, M.-B. (2008). Understanding memory through hippocampal remapping. *Trends Neurosci.* 31, 469–477.
44. Attardo, A., Fitzgerald, J.E., and Schnitzer, M.J. (2015). Impermanence of dendritic spines in live adult CA1 hippocampus. *Nature* 523, 592–596.
45. Zhou, Y., Won, J., Karlsson, M.G., Zhou, M., Rogerson, T., Balaji, J., Neve, R., Poirazi, P., and Silva, A.J. (2009). CREB regulates excitability and the allocation of memory to subsets of neurons in the amygdala. *Nat. Neurosci.* 12, 1438–1443.
46. Bito, H., Deisseroth, K., and Tsien, R.W. (1996). CREB phosphorylation and dephosphorylation: a Ca(2+)- and stimulus duration-dependent switch for hippocampal gene expression. *Cell* 87, 1203–1214.
47. Sheffield, M.E.J., Adoff, M.D., and Dombeck, D.A. (2017). Increased Prevalence of Calcium Transients across the Dendritic Arbor during Place Field Formation. *Neuron* 96, 490–504.e5.
48. Grosmark, A.D., and Buzsáki, G. (2016). Diversity in neural firing dynamics supports both rigid and learned hippocampal sequences. *Science* 351, 1440–1443.
49. Bittner, K.C., Milstein, A.D., Grienberger, C., Romani, S., and Magee, J.C. (2017). Behavioral time scale synaptic plasticity underlies CA1 place fields. *Science* 357, 1033–1036.
50. Howard, M.W., MacDonald, C.J., Tiganj, Z., Shankar, K.H., Du, Q., Hasselmo, M.E., and Eichenbaum, H. (2014). A unified mathematical framework for coding time, space, and sequences in the hippocampal region. *J. Neurosci.* 34, 4692–4707.
51. Tiganj, Z., Hasselmo, M.E., and Howard, M.W. (2015). A simple biophysically plausible model for long time constants in single neurons. *Hippocampus* 25, 27–37.
52. Lisman, J., and Redish, A.D. (2009). Prediction, sequences and the hippocampus. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 364, 1193–1201.
53. Piaget, J. (1952). *The Origins of Intelligence in Children* (W W Norton & Co).
54. Lewis, P.A., and Durrant, S.J. (2011). Overlapping memory replay during sleep builds cognitive schemata. *Trends Cogn. Sci.* 15, 343–351.
55. Tse, D., Langston, R.F., Kakeyama, M., Bethus, I., Spooner, P.A., Wood, E.R., Witter, M.P., and Morris, R.G. (2007). Schemas and memory consolidation. *Science* 316, 76–82.
56. McKenzie, S., Frank, A.J., Kinsky, N.R., Porter, B., Rivière, P.D., and Eichenbaum, H. (2014). Hippocampal representation of related and opposing memories develop within distinct, hierarchically organized neural schemas. *Neuron* 83, 202–215.
57. Villette, V., Malvache, A., Tressard, T., Dupuy, N., and Cossart, R. (2015). Internally Recurring Hippocampal Sequences as a Population Template of Spatiotemporal Information. *Neuron* 88, 357–366.
58. Buzsáki, G. (2006). *Rhythms of the Brain* (Oxford University Press).

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Bacterial and Virus Strains		
AAV9.Syn.GCaMP6f.WPRE.SV40	Penn Vector Core	Cat#AV9-PV2822
Deposited Data		
Raw imaging and behavioral videos	Google Drive	https://drive.google.com/open?id=1r4lpk_Z6rttc4btQAlml-S0uSmv0pTPW
Processed data	Mendeley Data	https://dx.doi.org/10.17632/f9fmrj98n3.1
Experimental Models: Organisms/Strains		
C57BL/6J mice	Jackson Laboratories	https://www.jax.org/ ; Cat#000664
Software and Algorithms		
MATLAB 2016b	MathWorks	https://www.mathworks.com/ ; RRID: SCR_001622
nVista	Inscopix	https://www.inscopix.com/
Mosaic 1.2	Inscopix	https://www.inscopix.com/
Tenaspis	GitHub	https://github.com/SharpWave/TENASPIS

CONTACT FOR REAGENT AND RESOURCE SHARING

Inquiries for reagent and resource sharing should be directed to the Lead Contact, William Mau (wmau@bu.edu) and they will be fulfilled, assuming reasonable requests.

EXPERIMENTAL MODEL AND SUBJECT DETAILS

Animal Subjects

All procedures were in compliance with the guidelines of the Boston University Animal Care and Use Committee. Subjects were 4 healthy adult male C57BL6/J mice (Jackson Laboratories), 5 – 10 months of age and weighing 25 – 33 g. Mice were initially socially housed with 1 – 3 cagemates in a vivarium with a 12 hr/12 hr light/dark cycle and lights on at 7am. After surgery, mice were singly housed.

Viral Constructs

For calcium imaging, virus (AAV9-Syn-GCaMP6f.WPRE.SV40) was supplied by U Penn Vector Core at a titer of $\sim 4 \times 10^{13}$ GC/mL, which was diluted down to $\sim 5-6 \times 10^{12}$ GC/mL with 0.05 M phosphate buffered saline prior to surgical infusion into CA1.

METHOD DETAILS

Stereotactic Surgeries

Naive mice underwent two stereotactic surgeries and one base plate implant for calcium imaging [23, 35, 37]. All surgeries were performed on mice anesthetized with $\sim 1\%$ isoflurane with mixed oxygen and 0.05 mL/kg buprenorphine. Mice also received injections of 5.0 mL/kg anti-inflammatory Rimadyl (Pfizer) and 400 mL/kg antibiotic Cefazolin (Pfizer). First, mice received infusions of AAV9-Syn-GCaMP6f (U Penn Vector Core). A small craniotomy was performed (AP –2.0 mm, ML +1.5 mm, DV –1.5 mm relative to bregma) and an infusion needle was inserted. The viral vector was injected at 40 nL/min and allowed 15 min to diffuse and minimize backwash prior to removing the needle. Three weeks after viral infusion, mice were implanted with a gradient index (GRIN) lens (1 mm diameter, 4 mm length; Inscopix). A 2 mm-diameter circular craniotomy centered on AP –2.25 mm, ML +1.8 mm was opened. The neocortex underneath this craniotomy was aspirated until vertical white fiber tracts were visible. Bleeding was controlled via irrigation with cold 0.9% saline solution and GelFoam (Pfizer). Once bleeding was arrested, the GRIN lens was carefully lowered into the craniotomy using a stereotactic device until the bottom of the lens was 200 microns dorsal to the infusion site. Gaps between the lens and the skull were filled in using a non-bioreactive silicone polymer, Kwik-Sil (World Precision Instruments). After the Kwik-Sil set, the lens was affixed to the skull using dental cement Metabond (Parkell) and the top of the lens was covered with a Kwik-Cast cap (World Precision Instruments) to protect the lens and occlude light until base plate attachment. Mice were allowed one week of convalescence before they were implanted with a base plate for camera attachment. The Kwik-Cast cap on the lens of the mouse was removed and a plastic base plate (Inscopix) was magnetically attached to the bottom of the camera. The camera objective was then aligned to the GRIN lens and lowered until visible and focused fluorescence was observed on nVista recording software (Inscopix). Adjustments were manually made to maximize focus of GCaMP6f expressing cells. After an optimal image was obtained,

the camera was raised ~50 microns to account for dental cement shrinkage during curing. The base plate was then affixed to the Metabond surrounding the lens using Flow-It ALC Flowable Composite (Pentron), cured with light, and finally covered with an additional layer of Metabond. The plastic cap of the base plate was then screwed on, and the mouse awoken.

Treadmill Running Behavior

A week after recovery, mice were introduced to a 40 cm x 60 cm rectangular track with an embedded motorized mouse treadmill (Columbus Instruments) as one of its long sides. Mice were acclimated to the environment until they reliably sought 20% sucrose water solution (3-4 days), delivered by a gravity feed. Then, they were trained to run in place on the treadmill for increasing intervals of time in between laps starting with 6 s. For the beginning sessions, running speed was titrated up from ~10 cm/s to 12-24 cm/s depending on the running speed of the subject. Once a stable velocity was reached, run duration was increased every two days by 1 s until the mouse was running for 10 s on the treadmill per lap. Once mice would reliably run for ~30 laps per day, data was then collected for 4 days, with each session lasting approximately 30 min and consisting of ~30 laps of 10 s treadmill running and water retrieval.

Freely-Moving Calcium Imaging and Mouse Tracking

Calcium imaging

A commercially available miniaturized epifluorescence microscope (Inscopix) was used to collect imaging videos of CA1 activity at a frame rate of 20 Hz. Digital gain (1.0-1.75) and LED intensity (~10%) was adjusted for each mouse to maximize dynamic range. Frames were spatially down-sampled from 1440 x 1080 pixels to 720 x 540 pixels (1.1 microns/pixel) to accelerate post-acquisition processing steps. Microscope attachment was done on awake, restrained mice. Optical focus and recording settings were kept consistent for each mouse each day. TIF videos collected via nVista were saved and transferred to a permanent workstation for pre-processing steps. First, videos were motion corrected and cropped using Mosaic (Inscopix). Cropping excluded areas with no GCaMP6f activity (usually resulting in 500 x 500 pixel videos).

Videos were then passed through a custom image segmentation algorithm called Tenaspis (software available at <https://github.com/SharpWave/TENASPIS>) that has been optimized to reduce optical crosstalk between regions of interest (ROIs) (see D.W. Sullivan et al., 2017, Soc. Neurosci., abstract). ROI-based segmentation algorithms detect events based on threshold crossings within an ROI, which is susceptible to Ca^{2+} transients bleeding in from a nearby overlapping cell. To contrast, Tenaspis detects events frame by frame, separates overlapping fluorescent regions first, and then afterward assigns these events to neuronal ROIs. Briefly, Tenaspis utilizes heuristics about neuron shape and size, and then iteratively (by frame) detects fluorescent regions that fit the description of a neuron. After identifying these regions, Tenaspis collects timestamps for the rising phase of each Ca^{2+} transient ($\Delta F/F > 2$ standard deviations above the mean) and allocates them to a neuronal ROI, then collapses images into ROI masks. Thus, all Ca^{2+} transients described in this article refer to times of increasing fluorescence, ignoring the slow decay of the Ca^{2+} indicator, the former more accurately reflecting action potentials in neurons.

Longitudinal cell tracking

In order to track neurons captured during recording sessions one or more days apart, neuron ROIs were registered across days. First, the minimum time projection for each session was computed to utilize vasculature as stationary landmarks during image alignment. Using these landmarks, each session's field of view was aligned to the first session's minimum projection via image registration software from MATLAB's Image Processing Toolbox, assuming rigid geometric transformation and rotation. Then, we successively took each session (reference sessions) and registered the neurons from that session to the next day's neurons (registered sessions; i.e., we would register Day 1 to Day 2, Day 2 to Day 3, etc.). Cell registration was done by searching for the nearest ROI, with a threshold that the displacement between ROI centroids must be under 3.3 microns. In rare cases where multiple neurons on the registered session were the same distance away from a neuron in the reference session, a spatial correlation was done for each candidate mask and the neuron with the higher correlation coefficient was ultimately matched to the reference neuron. To ensure that neurons did not drift excessively over the course of the experiment, for each mouse, the first day's neurons were registered to the last day's neurons to check for large deviations. Any neuron registrations from this condition (Day 4 versus Day 1) that differed from the first condition (Day 4 versus Day 3) were discarded. Additionally, in analyses involving neurons across multiple days, if a neuron on one day did not have a corresponding registered neuron on the subsequent day, it was discarded from the analysis.

Mouse tracking

The mouse's position was recorded using an overhead camera (30 Hz) and video tracking software CinePlex (Plexon). The tracking video was synchronized to the imaging using a TTL signal from Cineplex to trigger data acquisition on the microscope. Frames were linearly interpolated to match the sampling rate of the microscope. Position tracking was carefully reviewed and errors were manually corrected using a custom MATLAB script.

Histology and Epifluorescent Microscopy

After data collection, mice were perfused transcardially with 10% phosphate buffered formalin. Brains were extracted and then submerged in formalin for an additional two days, followed by 30% sucrose solution in phosphate buffered saline for another two days. Brains were then flash-frozen and sliced into 40 μm sections on a cryostat (Leica CM 3050S). Sections were mounted and coverslipped with Vectashield Hardset mounting medium with DAPI (Vector Laboratories) to visualize cell nuclei. Slides were imaged on a widefield epifluorescence microscope (Nikon Eclipse Ni-E) with a 10x and 20x objective to verify viral expression and lens tract

localization to the CA1 region. Viral expression was confirmed by examining native fluorescence of the GCaMP6f fluorophore under the anterior-posterior and medial-lateral extent of the lens tract.

QUANTIFICATION AND STATISTICAL ANALYSIS

Time Cell Selection

Tuning curves were constructed using temporally binned (250 ms) activity profiles of each cell during treadmill running and taking the mean across treadmill run trials. Time-shuffled tuning curves were also constructed by circularly shuffling activity timestamps for each trial 1,000 times and averaging across these trials. Temporal information (TI) was computed using the following equation:

$$TI = \frac{1}{\lambda} \sum_i \lambda_i \log_2 \left(\frac{\lambda_i}{\lambda} \right) P_{t_i}$$

where:

- λ is the average transient rate of the cell.
- λ_i is the average transient rate of the cell in time bin t_i (50 ms bins from 0 to 10 s).
- P_{t_i} is the probability the mouse is in time bin t_i .

The TI was then computed 1,000 times for iterations of shuffled calcium event timestamps within the treadmill run epochs. A neuron was called a time cell if it met the following criteria:

1. The neuron's TI was higher than 99% of the shuffled TIs.
2. The neuron fired for at least a quarter of the total completed treadmill runs.
3. The neuron had at least two consecutive time bins where its tuning curve exceeded the time shuffled tuning curve's 99% of the time.

Within-Session Trial Bias Score

To quantify trial preference, and thus characterize the within-session activity dynamics of single cells, we calculated each cell's trial bias score. This score was the mean of all the trial numbers that the cell was active within its temporal receptive field, divided by the total number of runs. Thus, a lower trial bias score indicated more activity closer to the beginning of the session whereas higher scores indicated more activity near the end of the session. Cells that were consistently active over the course of the entire session had a trial bias score of 0.5 or near 0.5. Examples in [Figures 3A](#) and [3B](#) had trial bias scores in the 5th (early-session cells) or 95th (late-session cells) percentiles of the distribution. Similar results were found when this analysis was repeated with data that excluded even-numbered trials to verify that it was not being driven by noise.

Population Correlations

To measure the similarity of the time and place cell population across trials and days, normalized Ca^{2+} transient traces for each trial were correlated with each other (Pearson correlation), and the correlation coefficients averaged across the population.

Naive Bayes Classifiers

Naive Bayes classifiers were built around the MATLAB function fitcnb with population activity as predictors and temporal variables as response variables, within each mouse and session. To avoid overfitting and to assess classifier performance, we utilized a cross-validation scheme where we trained the classifier on Ca^{2+} transient activity from a random 50% of available treadmill runs and tested it the other 50%. Below are the procedures we used for temporal decoding on the scale of seconds, minutes, and days:

Seconds ([Figure 2](#))

We trained the classifier on Ca^{2+} transient timings from a random 50% of treadmill runs each session, and tested on the remaining 50%. Chance was calculated by training classifier on same data with cell identities shuffled 50 times each per mouse and session.

Trials ([Figure 3F](#))

We trained the classifier on the number of Ca^{2+} transients on each treadmill run from a random 50% of treadmill runs sampled evenly from each of six trial blocks and tested on the remaining 50%. Trials were categorized into blocks due to technical restraints on the classifier. Chance was calculated by shuffling treadmill run blocks 50 times each per mouse and session. Analyses were repeated on different sizes of trial blocks and similar results were found for two trial blocks up to and including seven trial blocks.

Seconds across days ([Figure 4G](#))

We trained the classifier on Ca^{2+} transient timings on all treadmill runs from one session, and tested on all the treadmill runs from the other session. Chance was calculated by training the classifier on the same data with cell identities shuffled 50 times each per mouse and session.

Days (Figure 5F)

We trained the classifier on the number of Ca^{2+} transients each treadmill run from a random 50% of treadmill runs sampled evenly from each of four recording sessions and tested on the remaining 50%. Chance was calculated by shuffling day identities 50 times each per mouse.

Statistical Tests

All statistics were done with one- or two-way ANOVAs, Mann-Whitney U tests, or permutation tests by shuffling data along one dimension. All critical p values were corrected for multiple comparisons with Bonferroni corrections when applicable.

DATA AND SOFTWARE AVAILABILITY

Raw imaging and behavioral videos are available at https://drive.google.com/open?id=1r4lpk_Z6rttc4btQAlml-S0uSmv0pTPW. The processed data reported in this paper is available from Mendeley Data at <https://dx.doi.org/10.17632/f9fmrj98n3.1>.