



Consistent population activity on the scale of minutes in the mouse hippocampus

Document Version:

Accepted author manuscript (peer-reviewed)

Citation for published version:

Liu, Y, Levy, S, Mau, W, Geva, N, Rubin, A, Ziv, Y, Hasselmo, M & Howard, M 2022, 'Consistent population activity on the scale of minutes in the mouse hippocampus', *Hippocampus*, vol. 32, no. 5, pp. 359-372. https://doi.org/10.1002/hipo.23409

Total number of authors: 8

Digital Object Identifier (DOI): 10.1002/hipo.23409

Published In: Hippocampus

License: CC BY-NC-ND

General rights

@ 2020 This manuscript version is made available under the above license via The Weizmann Institute of Science Open Access Collection is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognize and abide by the legal requirements associated with these rights.

How does open access to this work benefit you?

Let us know @ library@weizmann.ac.il

Take down policy

The Weizmann Institute of Science has made every reasonable effort to ensure that Weizmann Institute of Science content complies with copyright restrictions. If you believe that the public display of this file breaches copyright please contact library@weizmann.ac.il providing details, and we will remove access to the work immediately and investigate your claim.

Title: Consistent population activity on the scale of minutes in the mouse hippocampus

Yue Liu¹²³⁴*, Samuel Levy²³⁴, William Mau²³⁴**, Nitzan Geva⁵, Alon Rubin⁵, Yaniv Ziv⁵, Michael Hasselmo²³⁴, Marc Howard¹²³⁴

Department of Physics, Boston University, Boston 02215, USA
 Center for Systems Neuroscience, Boston University, Boston 02215, USA
 Center for Memory and Brain, Boston University, Boston 02215, USA
 Department of Psychological and Brain Sciences, Boston University, Boston 02215, USA
 Department of Neurobiology, Weizmann Institute of Science, Rehovot 76100, Israel

Corresponding author: Yue Liu. Center for Neural Science, New York University. 4 Washington Pl # 809, New York, NY 10003. <u>yl4317@nyu.edu</u>

* Current address: Center for Neural Science, New York University.

** Current address: Neuroscience Department, Icahn School of Medicine.

This is the author manuscript accepted for publication and has undergone full peer review but has not been through the copyediting, typesetting, pagination and proofreading process, which may lead to differences between this version and the Version of Record. Please cite this article as doi: 10.1002/hipo.23409

This article is protected by copyright. All rights reserved.

Abstract

Neurons in the hippocampus fire in consistent sequence over the timescale of seconds during the delay period of some memory experiments. For longer timescales, firing of hippocampal neurons also changes slowly over minutes within experimental sessions. It was thought that these slow dynamics are caused by stochastic drift or a continuous change in the representation of the episode, rather than consistent sequences unfolding over minutes. This paper studies the consistency of contextual drift in three chronic calcium imaging recordings from the hippocampus CA1 region in mice. Computational measures of consistency show reliable sequences within experimental trials at the scale of seconds as one would expect from time cells or place cells during the trial, as well as across experimental trials on the scale of minutes within a recording session. Consistent sequences in the hippocampus are observed over a wide range of time scales, from seconds to minutes. Hippocampal activity could reflect a scale-invariant spatiotemporal context as suggested by theories of memory from cognitive psychology.

10

Introduction

When we remember a particular experience from a trip, other memories from the same 11 trip would also come into mind. Indeed, the retrieval of an episodic memory is believed 12 to involve recovery of the spatiotemporal context associated with that particular episode 13 (Tulving, 1983). The hippocampus has long been implicated in episodic memory (Scoville 14 & Milner, 1957) and it contains single neurons that are active when the animal is at a 15 particular location within an environment (O'Keefe & Dostrovsky, 1971; Moser, Kropff, & 16 Moser, 2008) or at a particular time point during the gap between two stimuli (Pastalkova, 17 Itskov, Amarasingham, & Buzsaki, 2008; MacDonald, Lepage, Eden, & Eichenbaum, 2010; 18 Kraus, Robinson, White, Eichenbaum, & Hasselmo, 2013) (Figure 1a, top). Taken together, 19 this neural population activity can be thought of as a state of spatiotemporal context 20 upon which memories are organized (O'Keefe & Nadel, 1978; Howard, Fotedar, Datey, 21 & Hasselmo, 2005; Polyn & Kahana, 2008; Staresina & Davachi, 2009; Hasselmo, 2012; 22 Eichenbaum, 2017; DuBrow, Rouhani, Niv, & Norman, 2017; Buzsáki & Tingley, 2018; 23 Ekstrom & Ranganath, 2018; Yonelinas, Ranganath, Ekstrom, & Wiltgen, 2019). 24

Episodic memory retrieval is organized according to spatiotemporal proximity at many 25 different scales. When a participant has an episodic memory for an event from a particular 26 temporal position within a list (Kahana, 1996) or spatial position within an environment 27 (Miller, Lazarus, Polyn, & Kahana, 2013), this brings to mind events from nearby posi-28 tions, in time or in space. If episodic memory is indeed associated with the recovery of a 29 spatiotemporal context (Tulving, 1983), then the effect of proximity on behavior could be 30 caused by gradual changes in spatiotemporal context reflected in hippocampal ensembles 31 (Manns, Howard, & Eichenbaum, 2007; Ezzyat & Davachi, 2014; Rubin, Geva, Sheintuch, & 32 Ziv, 2015; Cai et al., 2016). In this view, memories for events close in space or time are linked 33 because of overlap in the spatiotemporal contexts associated with those events. Sequences of 34 time cells or place cells could serve as a spatiotemporal context; because they change slowly 35 over time they could mediate associations between nearby events (Wallenstein, Eichenbaum, 36 & Hasselmo, 1998; Howard et al., 2005; Hsieh, Gruber, Jenkins, & Ranganath, 2014). Thus 37 far, time cell sequences have only been observed over a few seconds within the delay pe-38 riod of an experimental trial embedded in a much longer recording session (Pastalkova et 39 al., 2008; MacDonald, Lepage, Eden, & Eichenbaum, 2011; Kraus et al., 2013). However, 40 behavioral effects linking memories separated in time are observed over many timescales 41 in list learning experiments (Howard, Youker, & Venkatadass, 2008; Unsworth, 2008) and 42 can span days and weeks in memory for real-world events (Healey, Long, & Kahana, 2019; 43 44 Uitvlugt & Healey, 2019). If memories across lists separated by many minutes can be linked, this suggests that hippocampal sequences should also unfold across trials over the scale of 45 minutes. Perhaps the sequence of cells that unfolds in the moments following the beginning 46 of a delay period of a few seconds has an analog in a sequence that unfolds over the entire 47 recording session following the beginning of the session. 48

A series of studies have found that the activity of hippocampal neurons does change 49 slowly over long timescales. For example, it was reported that population neuronal activity 50 in CA1 exhibits gradual changes over multiple trials that span minutes (Manns et al., 2007; 51 Ziv et al., 2013; Mau et al., 2018) (Figure 1a, bottom). It has also been reported that place 52 cell and time cell activity slowly "drift" across hours and days (Ziv et al., 2013; Mankin, 53 Diehl, Sparks, Leutgeb, & Leutgeb, 2015; Mankin et al., 2012; Rubin et al., 2015; Mau et 54 al., 2018; Cai et al., 2016). The observation of these slow changes with multiple recording 55 techniques make it unlikely that they are a recording artifact. However, it is possible that 56 this slow drift is simply caused by stochastic processes in the neural system or perhaps 57 a gradual but continuous change in the representation of events. Stochastic mechanisms 58 would cause changes in firing across trials but there is no reason to expect that they would 59

cause the same sequence over repeated experiences (Figure 1b, bottom left). However, if
slow changes are generated by the same mechanism as time cell sequences, we would expect
the dynamics to be consistent across repeated experiences. In much the same way as time
cell sequences can be understood as coding for the time since the delay period began, slow
changes in firing across trials could contain information about the time since the recording
session began.

A more recent study shows evidence for such coding of progression within a session (Sun, Yang, Martin, & Tonegawa, 2020). In that study, mice were trained to run four consecutive laps to obtain a reward. Some neurons in the hippocampus CA1 show elevated activity during a particular lap, and this firing pattern is consistent across repetitions of the same task (Sun et al., 2020). However, it still remains unclear whether similar neural activity patterns can be observed in tasks without a demand to maintain the task progression. It is also interesting to examine whether other forms of temporal modulations are present in encoding the progression of the task. As will be shown in the rest of the manuscript, the answers are positive to both questions.

Figure 1 describes a strategy for data analyses to distinguish consistent sequences from stochastic drifts. Consider a population of cells being recorded over two separate experiences. During each experience the activity of the population changes gradually. This effect can be demonstrated by measuring the correlation of the population activity patterns at different points in time. As one chooses time points further apart from one another within the experience the population becomes more decorrelated. Now, suppose that the population fires consistently from one experience to the next (Figure 1b, top left). In this case one would observe an analogous decorrelation when examining firing across different experiences (Figure 1b, top right). Although the two experiences could be separated by a time interval much longer than the duration of the experience itself, the population activity from a particular time point in each experience will be similar even if those time points are taken from different experiences. In contrast, if the within-experience correlation were not due to consistent sequences but, say, stochastic variability (Figure 1b, bottom left), the population would still change gradually over time *within* one experience. However, if the population activity simply decorrelates with time, that would also mean that one would not observe correlation between analogous time points within *different* experiences (Figure 1b, bottom right).

[Insert figure 1 around here]

The strategy of this paper is to evaluate whether slow changes in hippocampal activity across trials include consistent sequences extending across multiple trials. In this case, the experience in the example above would be one experimental session that consists of multiple trials. To establish the consistency of the population activity on the scale of a session, it would be necessary to compare the activity between different sessions. By the same reasoning in the example above, if the population activity from a particular time point in each session is similar even if those time points are taken from different sessions, it would indicate that the population firing is consistent between different sessions. Therefore, 99 it requires analysis across multiple sessions across days to analyze the consistency of the 100 neuronal dynamics over minute-level time scales. To compare to the more well-understood 101 sequences—time cells and place cells—we apply the same analyses to the population activity 102 within trials as well. By the same reasoning, it requires analysis across multiple sessions 103 over tens of minutes to examine the consistency of the neuronal dynamics across seconds-104 level time scales. It is impossible to assess whether a sequence is consistent or not if 105 one cannot record from the same population. Therefore, we study populations recorded 106 using the chronic calcium imaging technique that allows identification of the same neurons 107 across recording sessions (Ziv et al., 2013). We found that across three behavioral tasks, 108 populations of neurons in the CA1 region of mouse hippocampus exhibit consistent dynamics 109 both within second-long trials and across trials, spanning many minutes within a recording 110

111 session.

112

Results

To distinguish consistent sequences from stochastic dynamics (Figure 1), we exam-113 ined the consistency of the neuronal dynamics across two timescales while mice performed 114 reward-based navigational tasks (Mau et al., 2018; Levy, Kinsky, Mau, Sullivan, & Has-115 selmo, 2019; Rubin et al., 2015). In all experiments, each session consists of multiple trials, 116 during which mice were trained to navigate through an environment to obtain rewards (Fig-117 ure 2). One-photon endoscope calcium imaging was used to record the activity of many 118 neurons in the CA1 region of the hippocampus across multiple sessions that span days. 119 Consecutive sessions are separated by at least one calendar day (see Methods section for 120 details). Images from different sessions were cross-registered so that the activity of the same 121 neurons could be tracked across sessions (Methods). We develop a series of computational 122 measures for the consistency of activity over seconds-long delay intervals across trials. Not 123 124 surprisingly these measures of consistency detect sequences over seconds that are consistent across trials, driven by time cells and place cells. Next, we apply the same computational 125 measures to detect slow sequences of activity over multiple trials that are consistent across 126 experimental sessions. To the extent these measures show the same kind of consistent dy-127 128 namics described by time cells and place cells, we will establish that hippocampal ensembles exhibit consistent sequences of activity across trials at the time scale of a session. 129

[Insert figure 2 around here]

¹³⁰ Single hippocampal neurons have consistent activity across seconds and minutes

We first extracted the region of interest (ROI), that is, a region of the imaging field 131 believed to correspond to a particular neuron, from the movie obtained from calcium imag-132 ing (see Methods for details). Next, we plotted the normalized calcium transient density 133 of individual ROIs against position or time within a trial. For each ROI we only included 134 trials where it had at least one calcium transient event during the trial period examined. 135 We observed that many ROIs have consistent activity within a trial (Figure 3a-c). For 136 example, some ROIs always have higher activity around a particular time bin (Figure 3a, 137 right) or location bin (Figure 3b-c, right) during each included trial. Other ROIs have 138 higher activity at the start (Figure 3a, left) or end (Figure 3b-c, left) of each active trial. 139 At longer time scales, similar consistent neural activity was observed when activity was 140 plotted against trial number (Figure 3d-f, Methods). For example, some ROIs consistently 141 increase (Figure 3d-e, left) or decrease (Figure 3f, left and 3e, right) their calcium activity 142 across trials within a session. Other ROIs are consistently more active during particular 143 144 trials within the session (Figure 3f, right).

[Insert figure 3 around here]

To quantify the extent to which single neurons fire consistently across the population, we computed a firing consistency score for each neuron, which is a number between zero and one that represents the consistency of that neuron's calcium dynamics between pairs of trials or sessions compared to chance (see Methods for details). We found that in all experiments and for both timescales, the distributions of the firing consistency score are significantly skewed towards one compared to those obtained from the surrogate data where the mean neural activity during different trial bins was randomly shuffled independently for each ROI ¹⁵² (Figure 4). A Kolmogorov-Smirnov test between the distribution from true and shuffled

data showed reliable differences in all cases. This test statistic is the maximum distance 153 between the cumulative distribution functions of the two groups of data, and therefore 154 can be regarded as an effect size that is independent of the number of data points. As a 155 comparison, the same statistic between two Gaussian distributions separated by 1 standard 156 deviation is about 0.38. In order to parallel the more frequently used Cohen's *d* measure 157 for the effect size, we also computed the number of standard deviations *d* that the two 158 159 Gaussian distributions would need to be separated by for the K-S test statistic to be the same as the ones seen in our data. The results for the above statistical tests are: panel a: 160 161 p < .001, D=0.62, n = 1860, d=1.76; panel b: p < .001, D=0.61, n = 4078, d=1.72; panel c: $_{162} p < .001, D=0.52, n = 1202, d=1.42$; panel d: p < .001, D=0.16, n = 1338, d=0.41, panel e: p < .001, D=0.13, n = 3573, d=0.33; panel f: p < .001, d=0.07, n = 1675, d=0.18. This 163 indicates there are significantly more neurons that have consistent dynamics both across 164 trials and across sessions than would be expected by chance. The fraction of neurons with a 165 166 firing consistency score greater or equal to 0.9 among all the neurons included in the analysis 167 is: panel a: 1324/1860 = 71.2%; panel b: 2848/4078 = 69.8%; panel c: 728/1202 = 60.6%; 168 panel d: 324/1338 = 25.6%; panel e: 816/3573 = 22.8%; panel f: 261/1675 = 15.6%.

[Insert figure 4 around here]

We further confirmed the genuity of this robust long timescale firing by adapting the 169 temporal information metric which was previously used for identifying hippocampal "time 170 cells" on the timescale of seconds (Mau et al., 2018). Briefly, a temporal information score 171 between 0 and 1 was computed based on the Shannon entropy, where 0 indicates that the 172 single cell activity does not reliably carry information about trial number within a session 173 compared to neural activity shuffled across trials (see Methods for details). The distribution 174 of the temporal information score is significantly (p < 0.001, Kolmogorov-Smirnov test) 175 different between real and shuffled data in all three datasets (Supplementary figure S20). 176

There could be different types of single cell dynamics that contribute to the observed 177 178 high firing consistency across repeated experiences. For example, some cells could have gradually increasing or decreasing activity, or they could exhibit non-monotonic dynamics 179 over time such as those of time cells. To disentangle these two types of single cell dynamics, 180 we used a similar method as above to construct a firing linearity score for each neuron. The 181 right panels of Figure 4 show the joint scatter plot of the firing consistency score and the 182 firing linearity score for each of the experiments. As can be seen, there is a wide distribution 183 of firing linearity across the population. This indicates that there is a diversity of consistent 184 temporal dynamics both across a trial and a session. 185

We also found that the firing consistency scores on the timescale of minutes and seconds 186 are not significantly correlated with one another across neurons for the linear track task 187 (Supplementary figure S18c, Kendall's τ : τ (1099) = 0.01, p = 0.6). They are weakly but 188 significantly correlated for the treadmill running and the spatial alternation tasks (Supple-189 mentary figure S18a, b. Kendall's τ : τ (1245) = 0.13, $p < 10^{-8}$ for a, τ (3019) = 0.079, $p < 10^{-8}$ 190 10^{-8} for b). This result implies that the second-scale neurons and the minute-scale neurons 191 are from two independent but overlapping populations in the linear track task, and they 192 show slightly more than chance overlap in the treadmill running and the spatial alterna-193 tion tasks. In either case, there exist neurons whose dynamics are jointly modulated by two 194 functions of time, one of which changes on the timescale of seconds and one on the timescale 195 of minutes. In other words, there is a multiplexed code for these two different time scales. 196

Hippocampal population dynamics are consistent both over seconds and min utes

The single cell analysis above left out trials or sessions where a given neuron is not 199 active (does not have any calcium transient during the selected time period). To examine 200 whether the consistency is present when the full ensemble of neurons are considered, we 201 next investigated the consistency of the population-level dynamics across trials and sessions. 202 To this end, we computed the cosine similarity between pairs of population vectors from 203 different trials (Figure 5a-c) or sessions (Figure 5d-f) and assembled them into a matrix 204 (see Methods for details). We found that in all experiments and across both timescales, the 205 matrices exhibit a pattern where the values are highest along the diagonal, which indicates 206 that the population dynamics are consistent across repeated trials and sessions (c.f., Fig-207 ure 1b). Statistical significance was evaluated using a permutation test. A "diagonalness 208 score" was computed to quantify the degree to which a matrix shows a diagonal pattern 209 (see Methods for details). As a result, the diagonalness score for all matrices are greater 210 than the matrices obtained from 10,000 random shuffles of the data. We also computed the 211 effect size as the z-scored statistic with respect to the shuffled distribution of diagonalness 212

²¹⁵ high-valued region spreads out later in the trial, indicating that the population dynamics

scores. The values are: Fig 5a: 29.0, Fig 5b: 18.1, Fig 5c: 50.8, Fig 5d: 14.8. Fig 5e: 41.9, Fig 5f: 6.5. Notably, for the across-trial similarity matrix in the treadmill running task, the

slow down as time progresses (Figure 5a). This is consistent with the observation in the
original study that the number density of sequentially-activated time cells goes down in
time (Mau et al., 2018).

If the population dynamics are consistent across sessions, it should be possible to decode 219 the trial that the animal was in from the population activity. Moreover, this decoder should 220 be generalizable across sessions, meaning that a decoder trained on a subset of sessions 221 should be able to predict the trial number on the rest of the sessions. To verify this, we 222 223 trained a Linear Discriminant Analysis (LDA) classifier to predict the trial bin number within a session using the neural activity of all the other sessions. The posterior probability 224 given by the LDA classifier for each trial bin was plotted against the actual trial bin (Figure 225 6, middle). All three heatmaps show significant diagonal patterns. Statistical significance 226 was evaluated using a permutation test similar to the one used for the cosine similarity 227 measure (Figure 5, see Methods for details). The number of shuffles with higher diagonalness 228 scores than the true data is: Fig 6a: 0/1000, Fig 6b: 0/1000, Fig 6c: 5/1000. An effect 229 size was computed based on the z-scored statistic with respect to the shuffled distribution 230 231 of diagonalness scores, same as for the cosine similarity measure above. The values are: 232 Fig 6a: 4.6, Fig 6b: 11.6, Fig 6c: 2.6. Moreover, the posterior probability given by the LDA for the correct trial bin is above chance for all animals in all tasks (Figure 6, right). These 233

Author Manuscript

240

results show that the LDA classifier trained on all but one sessions can accurately predict the trial bin number of the left-out session. In other words, the minute-scale population dynamics over multiple trials are consistent between different sessions. Taken together, both the cosine similarity and the LDA decoding analysis confirm and reinforce the result obtained from the single cell analysis (Figure 4) that the population exhibits consistent dynamics over the timescale of both seconds and minutes.

Discussion

In this paper, we show that the firing dynamics of hippocampal neurons are consistent 241 over both seconds and minutes. The novel observation is that slow dynamics over minutes 242 include slow sequences and are not simply random drifts (Figure 5d-f, Figure 1b). This 243 population effect results from a significant proportion of neurons with consistent dynam-244 ics over repeated experiences (Figure 4). These neurons have both monotonic and more 245 complex activity modulations across each experience (Figures 3 and 4 and Supplemen-246 tary figures S1 and S2). Therefore, the hippocampal neurons exhibit consistent dynamics 247 over two nested timescales—changing both systematically within a trial and systematically 248 within a session—in each of these experiments. 249

As seen above, the effect sizes for the population analysis are much higher than those for the single cell analysis. This is because the effect size for the single cell analysis is more affected by the variability across cells in terms of their firing consistency, whereas the one for the population analysis measures the significance of the firing consistency on the population level, therefore less affected by single cell variability.

Multiple metrics were used to quantify the firing consistency on the single cell and 255 population levels. The single cell rank metric (Figure 4) directly measures the correlation 256 between the temporal firing patterns of two experiences. The temporal information metric 257 258 (Figure S20) measures the temporal modulation of the average firing across experiences, and is commonly used to identify second-scale time cells (e.g. Mau et al., 2018). The two 259 population analyses also serve slightly different purposes: the cosine similarity metric (Fig-260 ure 5) directly measures the consistency of the population temporal firing patterns across 261 two experiences. On the other hand, the LDA classification accuracy metric (Figure 6) 262 measures the ability for the same downstream population to linearly read out elapsed time 263 on the scale of seconds across trials, and to read out task epoch on the scale of minutes 264 across sessions. 265

This result suggests that the spatiotemporal context as represented by population of 266 neurons in the hippocampus has meaningful dynamics over multiple timescales, from sec-267 onds to minutes. The sensitivity to multiple timescales may enable the hippocampus to 268 269 adaptively encode natural stimuli, which vary at many different scales (Voss & Clarke, 1975; Hasson, Yang, Vallines, Heeger, & Rubin, 2008) and account for the self-similar 270 structure in hippocampal correlations (Meshulam, Gauthier, Brody, Tank, & Bialek, 2019). 271 The responsiveness of hippocampal dynamics provides a constraint for behavioral models 272 of human memory. Models that rely on boundaries and event segmentation (Farrell, 2012; 273 Franklin, Norman, Ranganath, Zacks, & Gershman, 2020) must be able to construct and 274 utilize segments over multiple nested scales. Similarly, neural models for sequence gener-275 ation (Buzsáki & Tingley, 2018; Rajan, Harvey, & Tank, 2016; Howard et al., 2014) must 276 have the capacity to generate sequences at many different scales. 277

A recent study shows that there are neurons in the hippocampus CA1 of mice that 278 encode the number of laps that the animal has traversed in a task where they have to run 279 four consecutive laps to obtain a reward (Sun et al., 2020). These "lap cells" also constitute 280 a reproducible hippocampal sequence, but our results differ from this study in three aspects. 281 First, while some of the neurons reported in our work indeed show elevated activity for a 282 particular trial bin, analogous to the "lap cells" in Sun et al., the majority of neurons that 283 we observed exhibit other types of temporal modulations (Figure 3d-f and Supplementary 284 figures S1- S2). Second, the information about trial number is beneficial for the mice in 285 the task of Sun et al., yet is not useful for the mice in the tasks we analyzed. Yet, slow 286 dynamics are still observed. Third, the timescale of the slow dynamics in our study is an 287 order of magnitude longer than that observed in Sun et al.. 288 Possible causes of slow dynamics 289 Possible non-physiological causes. There are recording artifacts specific to calcium 290 291

imaging that could conceivably cause slow dynamics. For example, photobleaching could cause the calcium fluorescence signal to decrease gradually during each imaging session or gradual heating of the brain could potentially produce stereotypical changes in the apparent calcium fluorescence signal for each ROI over the course of an imaging session. It is difficult to reconcile these artifactual accounts of slow dynamics with non-monotonic patterns of firing over the session or the similarity between effects observed across trials to the effects within trial. The findings within-trial are quite consistent with results using extracellular electrodes. We conclude that it is likely that the slow dynamics are physiological in origin.

Variables correlated with time during a session. It is possible that the slow 299 dynamics observed in the hippocampus are not due to memory *per se* but reflect consistent 300 slow dynamics in the environment or internal state of the animal over the course of the 301 recording session. There are several possibilities for such variables. For example, the satiety 302 of the animal presumably decreases over the course of each recording session. Indeed, it has 303 recently been reported that thirst level has a dramatic impact on the population activity in 304 multiple brain regions in mice over the course of minutes (Allen et al., 2019). Moreover, a 305 recent study shows that slow drifts over minutes in area V4 and prefrontal cortex of monkey 306 are correlated with systematic changes in animal's behavior during a perceptual decision 307 making task (Cowley et al., 2020). In addition, microdialysis of acetylcholine shows higher 308 levels of acetylcholine when an animal is first removed from the home cage and placed in a 309 task (Acquas, Wilson, & Fibiger, 1996). Acetylcholine levels decrease over time on the scale 310 of minutes and have been shown to depolarize hippocampal neurons (Cole & Nicoll, 1984) 311 and increase firing rate (Fu et al., 2014), consistent with the cells that gradually increase or 312 decrease their activity (Figure 3, 4). In all of these cases, the sequential activation of cells 313 over the recording session would require that the hippocampus codes for a monotonically 314 changing variable with a sequence of receptive fields. Indeed, this kind of pattern has been 315 observed for hippocampal receptive fields as a function of smooth changes in frequency of 316 a behaviorally-relevant tone (Aronov, Nevers, & Tank, 2017). 317

In some sense the empirical story for very slow sequences is analogous to the empirical story for place cells or time cells shortly after their initial report. Although a consensus has emerged that place cells and time cells express spatial and temporal relationships between events in the service of memory (e.g., Eichenbaum, 2017), this view only emerged after

r Manusc ALITDO

extensive empirical studies. For instance, a neuron that fires when the animal is in a 322 specific position of an environment could be responding to the visual stimuli available at 323 that location, the particular configuration of auditory stimuli available, or olfactory cues 324 present on the track. Early studies ruled out a series of possible confounds of spatial 325 position (e.g., Quirk, Muller, & Kubie, 1990; Save, Cressant, Thinus-Blanc, & Poucet, 1998). 326 Similarly, it is possible that initial reports of time cells could have been solely a reaction to 327 a behavioral confound during the delay such as a stereotyped behavior. However, time cells 328 have been observed in head-fixed animals and different stimuli trigger distinct sequences 329 (e.g., Pastalkova et al., 2008; Taxidis et al., 2020; Cruzado, Tiganj, Brincat, Miller, & 330 Howard, 2020), ruling out most possible confounds. 331

332 Slow dynamics as memory for the past. It has been clearly established that hippocampal time cells express memory for the time and identity of past events (e.g., Pastalkova 333 et al., 2008; Taxidis et al., 2020; Cruzado et al., 2020). The most interesting possible cause 334 of the slow dynamics observed here is that they reflect the same computational mechanism, 335 but over much slower time scales than within-trial time cells. How might the same compu-336 tational mechanism generate neural dynamics across a wide range of timescales? It has been 337 suggested (Rolls & Mills, 2019; Shankar & Howard, 2012) that sequential activity across 338 multiple timescales could originate from cells that show exponential decaying activity over 339 the same range of timescales, which have been reported in two recent studies. Tsao et al. 340 (2018) observed slow changes in the firing of neurons in lateral entorhinal cortex (LEC). 341 In that study, LEC neurons changed their firing rate abruptly and then relaxed back to 342 baseline with a broad range of decay rates. For instance, upon entry to a particular envi-343 ronment, a neuron might rapidly increase its firing rate and then decay exponentially back 344 to baseline over several minutes. Other neurons ramped over the entire recording session 345 so there was a variety of decay rates across neurons. This slowly-varying signal in LEC at 346 the scale of minutes could be a cause of the slow sequences we observed in hippocampus. 347 In another study, Bright et al. (2020) studied neurons in monkey EC during a visual task. 348 After an image was presented, the neurons changed firing rate then gradually relaxed back 349 to baseline with a variety of decay rates. Because there was a variety of relaxation rates it 350 was possible to decode time since image onset over several seconds (see also Hyde & Strow-351 bridge, 2012). Very long-lasting firing in EC has been observed *in vitro* (Egorov, Hamam, 352 Fransén, Hasselmo, & Alonso, 2002) and is believed to be caused by the nonspecific calcium-353 sensitive (CAN) cationic current. Computational models show that the CAN current can 354 also induce slowly decaying firing with a variety of decay rates in a simple integrate-and-fire 355 neuron model (Tiganj, Hasselmo, & Howard, 2015). Computational models have shown 356 that the temporal information carried by slowly-decaying activity can be used to generate 357 a population of sequentially-activated time cells (Shankar & Howard, 2012; Howard et al., 358 2014; Liu, Tiganj, Hasselmo, & Howard, 2019; Rolls & Mills, 2019; Liu & Howard, 2020). 359 Of course, the definitive test of whether the slow hippocampal sequences reflect a very slow 360 form of memory is whether the identity of events on previous trials can be decoded. This 361 analysis is not feasible given the design of the tasks analyzed here, but similar analysis 362 have been done on the neural activity from other cortical regions (e.g. Schoenbaum & 363 Eichenbaum, 1995b, 1995a; Bernacchia, Seo, Lee, & Wang, 2011; Morcos & Harvey, 2016). 364 In summary, it was shown that the slow neuronal activity on the timescale of minutes

³⁶⁵ In summary, it was shown that the slow neuronal activity on the timescale of minutes ³⁶⁶ are consistent across repeated sessions. This slow dynamics is most likely related to either a gradual change of the animal's internal state, or a gradual evolution of the animal's memory
trace about past events, or both. The exact function of and the mechanism that generates
this slow consistent activity remains to be elucidated in future experiments.

Methods

371 Behavioral tasks and calcium imaging

The treadmill running task. All procedures were in compliance with the guidelines of the Boston University Animal Care and Use Committee. Four 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 2, Experiment 1). During each session, the mice completed between 35-37 trials. A total of 4 sessions were performed for each mouse.

Mice received infusions of AAV9- Syn-GCaMP6f (U Penn Vector Core). Imaging data 378 379 in dorsal CA1 were acquired using a commercially available miniaturized head-mounted epifluorescence microscope (Inscopix). The raw video was pre-processed using an image 380 segmentation algorithm called Tenaspis (D.W. Sullivan et al., 2017, Soc. Neurosci., ab-381 stract, software available at https://github.com/SharpWave/TENASPIS) to extract ROIs 382 and assign calcium transient events to each ROI. This algorithm is designed to better dis-383 tinguish between overlapping ROIs. The calcium transients it detects correspond to the 384 rising phase of the calcium fluorescence. 296-1136 ROIs were identified during each session. 385 There is a total of 4 imaging sessions spanning 4 calendar days. 386

In order to identify the same neurons across recording sessions that are days apart, 387 ROIs were cross-registered across days. Briefly, this was done by first aligning the field of 388 view of each session to the first session using vasculature as stationary landmarks via image 389 registration software from MATLAB's Image Processing Toolbox, assuming rigid geomet-390 ric transformation. Then, cells were successively registered from each session to the next 391 session (Day 1 to Day 2, Day 2 to Day 3, etc.). Cells were registered by searching for the 392 nearest ROI with a threshold that the ROI centroids must be within 3.3 microns apart. 393 To ensure that neurons do not drift excessively across days, the first day's neurons were 394 registered with the last day's neurons, and any registrations between Day 4 and Day 1 that 395 are different from Day 4 and Day 3 were discarded. Cells are stably tracked across sessions 396 using this method, as illustrated in the original paper (Figure 4, 5 and S3 in Mau et al., 397 2018). More details about the behavioral setup and the calcium imaging experiment can 398 be found in the Methods section of Mau et al. (2018). 399

The spatial alternation task. All procedures presented here were approved by the 401 Institutional Animal Care and Use Committee (IACUC) at Boston University. Four mice 402 received infusions of AAV9-Stn134 GCaMP6f (University of Pennsylvania Vector Core, ob-403 tained at a titer of 4×10^{13} GC/mL and 135 diluted it to 5×10^{12} GC/mL with 0.05 M 404 phosphate buffered saline). They were trained on a spatial alternation task, during which 405 they alternated between "study" and "test" trials. On study trials, mice were placed on the 406 center stem of maze, ran to the crossroads, where a removable barrier forced them to run 407 down one of the two return arms and received a reward of chocolate sprinkle. They were 408 then moved into the delay area located at the bottom of the center stem, waited through a 409

370

400

427

20-second delay, and the delay barrier was lifted to start the test trial. On a test trial, mice 410 again ran up the center stem to the crossroads but this time there was no barrier and they 411 had to remember the direction they traveled on the study trial and turn to the return arm 412 opposite to the preceding study trial in order to receive a reward. They then moved to the 413 delay area, and were placed in their home cage to wait through a 15-25 second inter-trial 414 interval while the next study trial was set up (Figure 2, Experiment 2). Mice completed 415 between 25 and 40 study-test trial pairs per session. Each of these trial pairs is considered 416 a "trial" in the analysis in the main text. There is a total of 9-11 imaging sessions spanning 417 up to 17 calendar days 418

419 The experimental procedures for calcium imaging and the data pre-processing pipeline are the same as the treadmill running task. 1149-3165 ROIs were identified for each ses-420 sion. The cell cross-registration procedure is slightly different from the treadmill running 421 task. Sessions were aligned to a "base" session from the middle of the recording schedule 422 using 25-40 "anchor" cells. Cells with centers within 3 microns were identified as the same 423 cell. Cells are stably tracked across sessions using this method, as illustrated in the original 424 paper (Figure 1 and Supplemenatary Figure 1 in Levy et al., 2019). More details about the 425 experimental setup can be found in Levy et al. (2019). 426

The linear track task. All procedures were approved by the Weizmann Institute 428 IACUC. Three mice (2 were injected with AAV2/5-CaMKIIa-GCaMP6f and one was a 429 Thy1-GCaMP6f transgenic; Jackson stock number 025393) were trained to run back and 430 forth on an elevated 96 cm long linear track. They received water sweetened with lemon 431 flavored fruit juice concentrate at each end of the track. An overhead camera (DFK 33G445, 432 The Imaging Source, Germany) was used to record mouse behavior. Each session consisted 433 of five 3-min trials with 3-min intertrial intervals. There are 7-8 imaging sessions conducted 434 every other day for each mouse. Sessions are in the morning and the afternoon in alternation. 435 An integrated miniature fluorescence microscope (nVistaHD, Inscopix) was used to ob-436 tain the imaging data from the CA1 region of the hippocampus. Imaging data was pre-437 processed using commercial software (Mosaic, Inscopix) and custom MATLAB routines as 438 previously described in Ziv et al., 2013. Spatial filters corresponding to individual ROIs 439 were first identified using a cell-sorting algorithm that utilizes principal component analysis 440 and independent component analysis (PCA and ICA, Mukamel, Nimmerjahn, & Schnitzer, 441 2009) and then subjected to further manual cell sorting (see the "Materials and methods" 442 section of Rubin et al. (2015) for more details). Calcium transient events were identified 443 when the amplitude of the calcium traces $\frac{dF}{F}$ crossed a threshold of 5 median absolute 444 deviations (MAD). Further measures were taken to avoid the detection of multiple peaks 445 as well as the spillover of the calcium fluorescence to neighboring cells. More details about 446 the method can be found in the "Materials and methods" section of Rubin et al. (2015) 447

Registration of cells across sessions was performed by first aligning the field of view in each session to the first session and then computing the spatial correlation between ROI centroids in the reference coordinate system. Pairs with spatial correlation > 0.7 or distance < 5 μ m were registered as the same neuron. Cells can be stably tracked over days, as illulstrated in the original article (Figure 1 and Figure 1-figure supplment 3 in Rubin et al., 2015). For the full detail on the experimental setup please refer to the "Materials and methods" section of Rubin et al., 2015.

455 Data analysis methods

456

Coarse-graining of calcium activity.

Coarse-graining for the across session dynamics. To extract the slow neuronal 457 dynamics across multiple trials while filtering out the fast within-trial dynamics, the neural 458 activity was first temporally coarse-grained before further analysis. When comparing a pair 459 of sessions, the session with more trials was first truncated to have the same number of trials 460 as the other session. Then, all the remaining trials within a session were evenly divided into 461 5 trial bins by using the array_split function in Numpy. Then the neural activity for each 462 463 ROI is the calcium transient density averaged over that trial bin. Therefore, the temporally coarse-grained activity of an ROI *n* during a session *i* was represented by a time series with 464 length 5: $v_{n,i} = [v_{n,i,1}, v_{n,i,2}, v_{n,i,3}, v_{n,i,4}, v_{n,i,5}]$. Furthermore, since we are interested in 465 the temporal modulation of the neural activity rather than the absolute magnitude of the 466 activity, the coarsed-grained activity of each cell was z-scored across trial bins. We chose 467 5 as the number of time bins within a session since each session in the linear track task 468 consists of 5 running trials (Figure 2, Experiment 3), and we wish to keep the way trials 469 are divided consistent across experiments. Similar results were obtained for the treadmill 470 running task and the spatial alternation task by using 10 trial bins. Furthermore, we only 471 averaged over the calcium activity during time periods when the animal's behavior is under 472 experimental control. In the treadmill running task, the time periods used are when the 473 animal is running on the treadmill for 10 seconds. In the spatial alternation experiment, 474 the time periods used are when the animal is running along the start arm. In the linear 475 track experiment, the time periods used are when the animal's position is within the middle 476 60% of the linear track. 477

Coarse-graining for the within-trial dynamics. To extract the within-trial neu-478 ronal dynamics, coarse-graining was done in a similar way by computing the calcium tran-479 sient density over 10 time bins or location bins within each trial. For the treadmill running 480 task, calcium event rate was averaged over each second during the 10-second running pe-481 riod. For the spatial alternation task, the start arm was evenly divided into 10 location 482 bins and total number of calcium transient events within each bin divided by the amount 483 of time the animal spent in that bin was computed. Unless otherwise specified, all analysis 484 was performed separately for the two task epochs (study and test) and two trial types (turn 485 left and turn right) and the results were averaged. For the linear track task, the within-trial 486 neural activity was computed by first computing for each individual run the number of 487 calcium transient events within each location bin divided by the amount of time the animal 488 spent in that bin, and then averaging this quantity over all runs within a 5-minute trial. 489 This was done for the two running directions separately, and the results were averaged. The 490 10 location bins span the middle 30% of the track. We chose the middle 30% of the linear 491 track because this is similar to the length of the start arm in the spatial alternation task. 492 Lastly, for all experiments, the activity of each neuron was z-scored across all spatial or 493 time bins for each trial. 494

Population analyses. To quantify the consistency of the population dynamics across sessions (Figure 5d-f), we computed the cosine similarity between pairs of population activity vectors during different sessions after they were coarse-grained and z-scored as described above. We then built a matrix where each element represents the cosine similarity between ⁴⁹⁹ a pair of population vectors at two trial bins during two different sessions, averaged over all⁵⁰⁰ pairs of sessions and all animals.

To quantify the consistency of the population dynamics across trials (figure 5a-c), population vectors were computed by averaging neural activity over time or location bins within each trial, as described above. Then a similar correlation matrix was constructed where each element is the cosine similarity between a pair of population vectors from different trials.

To test that the matrix shows a significant diagonal pattern, neural activity across 506 all bins within each session (Figure 5d-f) or trial (Figure 5a-c) was shuffled 10000 times 507 independently for each neuron and matrices from this shuffled data were constructed. We 508 characterized the degree to which each matrix shows a diagonal pattern by an index d, 509 which equals the difference between the average value of the near-diagonal matrix elements 510 to that of the off-diagonal matrix elements. The near-diagonal matrix elements are those 511 whose row and column indices are differed by less than half the dimension of the matrix. 512 Mathematically, $d = \langle M_{ij} \rangle_{|i-j| \ge \frac{N}{2}} - \langle M_{ij} \rangle_{|i-j| \ge \frac{N}{2}}$, where N is the dimension of the 513 matrix

and indicates the mean. Then we counted how many matrices constructed from the shuffled data have a value *d* greater than the matrix obtained from the true data. As a result, none of the 10000 matrices from the shuffled data has a higher *d* than the matrices in Figure 5.

The same method was used to quantify the significance of the diagonal patterns of the matrices obtained from the decoding analyses (Figure 6), except that the shuffling was done 1000 times.

Firing consistency score. To assess the consistency of the single neuron dynamics 521 across repeated trials or sessions, we computed a firing consistency score for each neuron. 522 For each cell n and each pair of sessions or trials (for example i and j), we computed 523 the Pearson correlation coefficient between the coarse-grained activity vectors $v_{n,i}$ and $v_{n,j}$ 524 obtained from the method described above. Then we shuffled the entries in each activity 525 vector and computed the Pearson's correlation coefficient again. This was repeated for 100 526 times. The Pearson correlation coefficients from all pairs of sessions (or trials) were then 527 averaged to obtain a mean Pearson correlation coefficient across pairs of sessions, for the 528 true data and each shuffle. For the second-level scores (Figure 4a-c), the Pearson correlation 529 coefficients were only computed between pairs of trials with the same trial type. The firing 530 consistency score was computed as the percentile where the true mean Pearson correlation 531 is at among all the shuffles (if there are multiple shuffles that yield the same Pearson's 532 correlation as the true data, the median percentile was used). Sessions or trials where the 533 neuron does not have any calcium transient event during the selected time period were 534 excluded from the analysis. 535

Firing linearity score. To disentangle the gradually ramping/decaying activity from 536 more complex temporal modulations, we computed a firing linearity score for each neuron. 537 For a given neuron n and session (or trial) i_i , we fitted a linear model as a function of 538 the bin number for the coarse-grained activity $v_{n,i}$ of that neuron. The F-statistic of this 539 linear model was computed along with those obtained from 100 shuffled activity vectors 540 (shuffling was performed in the same way as in computing the firing consistency score). 541 The F-statistics for both the true and shuffled data were then averaged across all sessions 542 (or trials) to obtain a mean F-statistic for the true data and for each shuffle. The firing 543

linearity score was computed as the percentile of the true mean F-statistic among all the
shuffled mean F-statistic was computed (if the F statistic of multiple shuffles equal the true
F statistic, the median percentile was used). For each neuron, sessions (or trials) where no
calcium transients were observed were excluded from the analysis.

Temporal information score. To assess the single cell basis for the coding of trial number, the temporal information metric, which was used as a criteria for identifying hippocampal "time cells" (Mau et al., 2018), was adapted to identify cells that robustly carry information about trial number during a session. First, the average activity during each trial bin was computed (see "Coarse-graining for the across session dynamics". The *z*-score step was not performed), and further averaged across sessions. This results in a time series of length *N* where *N* is the number of trial bins (*N* = 5 in our analysis). $v_n = \{v_{n,1}, v_{n,2}, v_{n,3}, v_{n,4}, v_{n,5}\}$. Then the temporal information for neuron *n* is the negative Shannon entropy of the normalized time series

$$TI(n) = \sum_{k=1}^{N} p_{n,k} \log p_{n,k},$$

where $p_{n,k} = \frac{v_{n,k}}{\sum_{k=1}^{N} v_{n,k}}$

If a cell is not modulated by trial number, or it is modulated in a way that is not consistent across sessions, its session-averaged activity $v_{n,i} = [v_{n,i,1}, v_{n,i,2}, v_{n,i,3}, v_{n,i,4}, v_{n,i,5}]$ would be weakly modulated by trial number. Since the Shannon entropy is the largest for uniform probability distributions, the TI (which is the negative Shannon entropy) would be small. The TI was then compared with the TIs of 1000 shuffles where neural activity was randomly shuffled across trial bin number before averaging over sessions, and the percentile of the true TI among all the shuffled was defined as the "temporal information score".

556

560

561

548

⁵⁵⁷ Data availability

The data that support the findings of this study are available from the corresponding ⁵⁵⁹ author upon reasonable request.

References

- Acquas, E., Wilson, C., & Fibiger, H. C. (1996). Conditioned and unconditioned stimuli
 increase frontal cortical and hippocampal acetylcholine release: effects of novelty,
 habituation, and fear. *Journal of Neuroscience*, *16* (9), 3089–3096.
- ⁵⁶⁵ Allen, W. E., Chen, M. Z., Pichamoorthy, N., Tien, R. H., Pachitariu, M., Luo, L., &
 ⁵⁶⁶ Deisseroth, K. (2019). Thirst regulates motivated behavior through modulation of
 ⁵⁶⁷ brainwide neural population dynamics. *Science*, 364 (6437).
- Aronov, D., Nevers, R., & Tank, D. W. (2017, Mar). Mapping of a non-spatial dimension
 by the hippocampal-entorhinal circuit. *Nature*, 543 (7647), 719-722. doi: 10.1038/na ture21692
- ⁵⁷¹ Bernacchia, A., Seo, H., Lee, D., & Wang, X. J. (2011). A reservoir of time constants for
 ⁵⁷² memory traces in cortical neurons. *Nature Neuroscience*, 14(3), 366-72.
- ⁵⁷³ Bright, I. M., Meister, M. L., Cruzado, N. A., Tiganj, Z., Buffalo, E. A., & Howard, M. W. (2020). A temporal record of the past with a spectrum of time constants in the monkey entorhinal cortex. *Proceedings of the National Academy of Sciences*.

⁵⁷⁴ Buzsáki, G., & Tingley, D. (2018). Space and time: The hippocampus as a sequence ⁵⁷⁵ generator. *Trends in cognitive sciences*, 22 (10), 853–869.

Cai, D. J., Aharoni, D., Shuman, T., Shobe, J., Biane, J., Song, W., . . . Silva, A. (2016).
A shared neural ensemble links distinct contextual memories encoded close in time. *Nature*, 534 (7605), 115–118.

⁵⁷⁹ Cole, A. E., & Nicoll, R. (1984). Characterization of a slow cholinergic post-synaptic
 ⁵⁸⁰ potential recorded in vitro from rat hippocampal pyramidal cells. *The Journal of* ⁵⁸¹ *physiology*, 352 (1), 173–188.

Cowley, B. R., Snyder, A. C., Acar, K., Williamson, R. C., Byron, M. Y., & Smith, M. A.
(2020). Slow drift of neural activity as a signature of impulsivity in macaque visual and prefrontal cortex. *Neuron*, 108 (3), 551–567.

Cruzado, N. A., Tiganj, Z., Brincat, S. L., Miller, E. K., & Howard, M. W. (2020). Conjunc tive representation of what and when in monkey hippocampus and lateral prefrontal
 cortex during an associative memory task. *Hippocampus, 30* (12), 1332-1346. Retrieved
 from https://onlinelibrary.wiley.com/doi/abs/10.1002/hipo.23282
 https://doi.org/10.1002/hipo.23282

- ⁵⁹⁰ DuBrow, S., Rouhani, N., Niv, Y., & Norman, K. A. (2017). Does mental context drift or ⁵⁹¹ shift? *Current opinion in behavioral sciences*, 17, 141–146.
- Egorov, A. V., Hamam, B. N., Fransén, E., Hasselmo, M. E., & Alonso, A. A. (2002).
 Graded persistent activity in entorhinal cortex neurons. *Nature*, 420 (6912), 173-8.
- Eichenbaum, H. (2017). On the integration of space, time, and memory. *Neuron*, 95 (5), 1007–
 1018.
- Ekstrom, A. D., & Ranganath, C. (2018). Space, time, and episodic memory: The hip pocampus is all over the cognitive map. *Hippocampus*, 28 (9), 680–687.

Ezzyat, Y., & Davachi, L. (2014). Similarity breeds proximity: Pattern similarity within
 and across contexts is related to later mnemonic judgments of temporal proximity.
 Neuron, *81*(5), 1179–1189.

Farrell, S. (2012). Temporal clustering and sequencing in short-term memory and episodic
 memory. *Psychological Review*, 119 (2), 223-71. doi: 10.1037/a0027371

 Franklin, N. T., Norman, K. A., Ranganath, C., Zacks, J. M., & Gershman, S. J. (2020).
 Structured event memory: A neuro-symbolic model of event cognition. *Psychological Review*, 127 (3), 327.

Fu, Y., Tucciarone, J. M., Espinosa, J. S., Sheng, N., Darcy, D. P., Nicoll, R. A., ... Stryker,
 M. P. (2014). A cortical circuit for gain control by behavioral state. *Cell*, 156 (6), 1139–
 1152.

- Hasselmo, M. E. (2012). *How we remember: Brain mechanisms of episodic memory*.
 Cambridge, MA: MIT Press.
- Hasson, U., Yang, E., Vallines, I., Heeger, D. J., & Rubin, N. (2008). A hierarchy of temporal
 receptive windows in human cortex. *Journal of Neuroscience*, 28 (10), 2539-50. doi:
 10.1523/JNEUROSCI.5487-07.2008
- Healey, M. K., Long, N. M., & Kahana, M. J. (2019). Contiguity in episodic memory.
 Psychonomic bulletin & review, 26 (3), 699–720.
- Howard, M. W., Fotedar, M. S., Datey, A. V., & 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. *Psychological Review*,

-Author Manuscrip 619 *112* (1), 75-116.

Howard, M. W., MacDonald, C. J., Tiganj, Z., Shankar, K. H., Du, Q., Hasselmo, M. E.,
& Eichenbaum, H. (2014). A unified mathematical framework for coding time, space,
and sequences in the hippocampal region. *Journal of Neuroscience*, 34 (13), 4692-707.
doi: 10.1523/JNEUROSCI.5808-12.2014

Howard, M. W., Youker, T. E., & Venkatadass, V. (2008). The persistence of mem ory: Contiguity effects across several minutes. *Psychonomic Bulletin & Review*,
 15 (PMC2493616), 58-63.

Hsieh, L.-T., Gruber, M. J., Jenkins, L. J., & Ranganath, C. (2014). Hippocampal activity
patterns carry information about objects in temporal context. *Neuron*, *81* (5), 1165–
1178.

Hyde, R. A., & Strowbridge, B. W. (2012). Mnemonic representations of transient stimuli
and temporal sequences in the rodent hippocampus in vitro. *Nature Neuroscience*,
15 (10), 1430-8. doi: 10.1038/nn.3208

Kahana, M. J. (1996). Associative retrieval processes in free recall. *Memory & Cognition*,
 24, 103-109.

Kraus, B. J., Robinson, R. J., 2nd, White, J. A., Eichenbaum, H., & Hasselmo, M. E. (2013).
Hippocampal "time cells": time versus path integration. *Neuron*, 78 (6), 1090-101. doi: 10.1016/j.neuron.2013.04.015

Levy, S. J., Kinsky, N. R., Mau, W., Sullivan, D. W., & Hasselmo, M. E. (2019). Hip pocampal spatial memory representations in mice are heterogeneously stable. *bioRxiv*,
 843037.

Liu, Y., & Howard, M. W. (2020). Generation of scale-invariant sequential activity in linear recurrent networks. *Neural Computation*, *32* (7), 1379–1407.

Liu, Y., Tiganj, Z., Hasselmo, M. E., & Howard, M. W. (2019). A neural microcircuit model
 for a scalable scale-invariant representation of time. *Hippocampus*, 29 (3), 260–274.

MacDonald, C. J., Lepage, K., Eden, U., & Eichenbaum, H. (2010). Hippocampal neurons encode the temporal organization of non-spatial event sequences. *Society for Neuroscience Abstracts*, *36*, 100.15.

⁶⁴⁸ MacDonald, C. J., Lepage, K. Q., Eden, U. T., & Eichenbaum, H. (2011). Hippocampal
⁶⁴⁹ âĂIJtime cellsâĂİ bridge the gap in memory for discontiguous events. *Neuron*, 71 (4),
⁶⁵⁰ 737 - 749.

Mankin, E. A., Diehl, G. W., Sparks, F. T., Leutgeb, S., & Leutgeb, J. K. (2015). Hip pocampal CA2 activity patterns change over time to a larger extent than between
 spatial contexts. *Neuron*, 85(1), 190-201. doi: 10.1016/j.neuron.2014.12.001

 Mankin, E. A., Sparks, F. T., Slayyeh, B., Sutherland, R. J., Leutgeb, S., & Leutgeb, J. K.
 (2012). Neuronal code for extended time in the hippocampus. *Proceedings of the National Academy of Sciences*, 109, 19462-7. doi: 10.1073/pnas.1214107109

⁶⁵⁷ Manns, J. R., Howard, M. W., & Eichenbaum, H. B. (2007). Gradual changes in hippocam-⁶⁵⁸ pal activity support remembering the order of events. *Neuron*, *56* (3), 530-540.

⁶⁵⁹ Mau, W., Sullivan, D. W., Kinsky, N. R., Hasselmo, M. E., Howard, M. W., & Eichenbaum,
 H. (2018). The same hippocampal ca1 population simultaneously codes temporal
 information over multiple timescales. *Current Biology* 28(10), 1400, 1508

information over multiple timescales. *Current Biology*, 28 (10), 1499–1508.
Meshulam, L., Gauthier, J. L., Brody, C. D., Tank, D. W., & Bialek, W. (2019). Coarse

graining, fixed points, and scaling in a large population of neurons. *Physical review*

6

670

671

672

673

674

675

676

677

689

690

- 664 *letters*, 123 (17), 178103.
- Miller, J. F., Lazarus, E. M., Polyn, S. M., & Kahana, M. J. (2013). Spatial clustering
 during memory search. *Journal Experimental Psychology: Learning, Memory and Cogntion*, 39 (3), 773-81. doi: 10.1037/a0029684
- Morcos, A. S., & Harvey, C. D. (2016). History-dependent variability in population dynamics during evidence accumulation in cortex. *Nature neuroscience*, *19* (12), 1672–1681.
 - Moser, E. I., Kropff, E., & Moser, M.-B. (2008). Place cells, grid cells, and the brain's spatial representation system. *Annual Review of Neuroscience*, 31, 69-89. doi: 10.1146/an-nurev.neuro.31.061307.090723

Mukamel, E. A., Nimmerjahn, A., & Schnitzer, M. J. (2009). Automated analysis of cellular signals from large-scale calcium imaging data. *Neuron*, 63 (6), 747-60. doi: 10.1016/j.neuron.2009.08.009

- O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map. preliminary evidence from unit activity in the freely-moving rat. *Brain Research*, *34* (1), 171-175.
- O'Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map*. New York: Oxford
 University Press.
- Pastalkova, E., Itskov, V., Amarasingham, A., & Buzsaki, G. (2008). Internally generated
 cell assembly sequences in the rat hippocampus. *Science*, 321 (5894), 1322-7.
- Polyn, S. M., & Kahana, M. J. (2008). Memory search and the neural representation of context. *Trends in Cognitive Science*, *12*(1), 24-30.
- Quirk, G. J., Muller, R. U., & Kubie, J. L. (1990). The firing of hippocampal place cells
 in the dark depends on the rat's recent experience. *Journal of Neuroscience*, 10 (6),
 2008-17.
- ⁶⁸⁷ Rajan, K., Harvey, C. D., & Tank, D. W. (2016). Recurrent Net ⁶⁸⁸ work Models of Sequence Generation and Memory. *Neuron*, 90(1), 128–
 - 142. Retrieved from http://dx.doi.org/10.1016/j.neuron.2016.02.009 doi: 10.1016/j.neuron.2016.02.009
- Rolls, E. T., & Mills, P. (2019). The generation of time in the hippocampal memory system.
 Cell reports, 28 (7), 1649–1658.
- ⁶⁹³ Rubin, A., Geva, N., Sheintuch, L., & Ziv, Y. (2015). Hippocampal ensemble dynamics ⁶⁹⁴ timestamp events in long-term memory. *eLife*, *4*, e12247.
- Save, E., Cressant, A., Thinus-Blanc, C., & Poucet, B. (1998). Spatial firing of hippocampal
 place cells in blind rats. *Journal of Neuroscience*, *18* (5), 1818-1826.
- Schoenbaum, G., & Eichenbaum, H. (1995a). Information coding in the rodent prefrontal cortex. II. Ensemble activity in orbitofrontal cortex. *Journal of Neurophysiology*, 74 (2), 751-62.
- Schoenbaum, G., & Eichenbaum, H. (1995b). Information coding in the rodent prefrontal cortex. I. Single-neuron activity in orbitofrontal cortex compared with that in pyriform cortex. *Journal of Neurophysiology*, 74 (2), 733-50.
- Scoville, W. B., & Milner, B. (1957). Loss of recent memory after bilateral hippocampal
 lesions. *Journal of Neurology, Neurosurgery, and Psychiatry*, 20, 11-21.
- Shankar, K. H., & Howard, M. W. (2012). A scale-invariant internal representation of time.
 Neural Computation, 24 (1), 134-193.
- Staresina, B. P., & Davachi, L. (2009). Mind the gap: binding experiences across space and
 time in the human hippocampus. *Neuron*, *63* (2), 267-76.

Sun, C., Yang, W., Martin, J., & Tonegawa, S. (2020). Hippocampal neurons represent
 events as transferable units of experience. *Nature Neuroscience*, 23 (5), 651–663.

- 711 Taxidis, J., Pnevmatikakis, E. A., Dorian, C. C., Mylavarapu, A. L., Arora, J. S., Samadian,
 K. D., . . . Golshani, P. (2020). Differential emergence and stability of sensory and
 temporal representations in context-specific hippocampal sequences. *Neuron*.
- Tiganj, Z., Hasselmo, M. E., & Howard, M. W. (2015). A simple biophysically plausible
 model for long time constants in single neurons. *Hippocampus*, 25 (1), 27-37.
- ⁷¹⁶ Tsao, A., Sugar, J., Lu, L., Wang, C., Knierim, J. J., Moser, M.-B., & Moser, E. I. (2018).
 ⁷¹⁷ Integrating time from experience in the lateral entorhinal cortex. *Nature*, 561 (7721),
 ⁷¹⁸ 57.
- 719 Tulving, E. (1983). Elements of episodic memory. New York: Oxford.
- Uitvlugt, M. G., & Healey, M. K. (2019). Temporal proximity links unrelated news events
 in memory. *Psychological science*, 30(1), 92–104.
- ⁷²² Unsworth, N. (2008). Exploring the retrieval dynamics of delayed and final free recall:
 ⁷²³ Further evidence for temporal-contextual search. *Journal of Memory and Language*,
 ⁷²⁴ 59, 223-236.
- 725 Voss, R. F., & Clarke, J. (1975). 1/f noise in music and speech. Nature, 258, 317-318.
- Wallenstein, G. V., Eichenbaum, H. B., & Hasselmo, M. E. (1998). The hippocampus as
 an associator of discontiguous events. *Trends in Neuroscience*, 21, 317-323.
- Yonelinas, A. P., Ranganath, C., Ekstrom, A. D., & Wiltgen, B. J. (2019). A contex tual binding theory of episodic memory: systems consolidation reconsidered. *Nature Reviews Neuroscience*, 20(6), 364–375.
- Ziv, Y., Burns, L. D., Cocker, E. D., Hamel, E. O., Ghosh, K. K., Kitch, L. J., Schnitzer,
 M. J. (2013). Long-term dynamics of CA1 hippocampal place codes. *Nature Neuro- science*, *16* (3), 264-6. doi: 10.1038/nn.3329

Figure captions

Figure 1. **Distinguishing slow consistent sequences from random drifts. a**. Top: The firing of time cells changes across seconds in sequences that are consistent across trials (left), which contributes to the decorrelation of population activity pattern over the timescale of seconds (right). Bottom: The firing of hippocampal neurons also changes slowly over trials (right), but it is not known if this is driven by consistent sequences on the timescale of minutes (left). **b**. Two possibilities for the nature of the slow dynamics over minutes. Top: it may reflect the animal's experience. If so, the neural activity would be similar if the animal goes through the same experience twice (left), analogous to the sequences on the timescale of seconds. In this case, the correlation between a pair of population vectors from different experiences will decay as the difference in their time, each within its experience, increases (right). Bottom: alternatively, the slow dynamics may be solely driven by the stochastic noise in neural systems and therefore drift randomly during different experiences (left). In this case, the correlations described above would not have any pattern (right).

Figure 2 . **Two timescales in the structure of the experiments.** For each experiment studied in this paper, the animals are trained to perform some task for several seconds-long trials in a recording session spanning tens of minutes. The calcium activity of the same neurons are recorded across sessions. During the treadmill running task (Experiment 1), mice are trained to run on the treadmill for 10 seconds before going to the opposite side of the maze to collect a water reward. The mice perform the same task for tens of trials each session for a total of around 20 minutes. For the spatial alternation task (Experiment 2), mice are trained to alternate between left and right turns in a T-maze to collect food rewards. Each trial consists of a study and test phase where mice have to turn to opposite directions at the choice point. Mice perform tens of trials for a total of around 30 minutes during each session. For the linear track experiment, mice are trained to run back and forth on a linear track to collect water rewards at both ends of the track. Each trial is about 3 minutes long and is separated by 3-minute resting periods where mice are placed in a separate box. Each session consists of 5 pairs of running and resting trials for a total of 30 minutes.

Figure 3 . Many hippocampal neurons exhibit consistent dynamics across seconds and across minutes. a-c. Example neurons with consistent dynamics across trials for each of the experiments. In a, the 10-second running period is evenly divided into 10 time bins. In b, the start arm of the maze is evenly divided into 10 location bins. In c, the linear track is evenly divided into 10 location bins, and neural activity is averaged over all runs within a 3-minute trial. See Methods for details. d-f. Example neurons with consistent activity across sessions. Trials within each session are evenly divided into 5 trial bins. Each line represents the z-scored calcium transient rate over one trial/session. Darker lines indicate earlier trials (for a-c)/sessions (for d-f). Inactive trials/sessions are not shown. See Supplementary figures S1 and S2 for more example neurons.

Figure 4. Many hippocampal neurons have consistent responses across seconds and across minutes. For each neuron, a firing consistency score was computed to estimate how consistent the single cell dynamics are for within trial dynamics (**a**-**c**) and for across-trial dynamics (**d**-**f**). The histograms show the distribution of firing consistency score relative to a surrogate distribution. To the extent these distributions differ, one can conclude

that there are consistent sequences. To evaluate the degree to which dynamics were simply monotonic changes in firing rate, we also computed a firing linearity score. The scatter plots show the consistency score and linearity score for each neuron. Red dots indicate the example single neurons shown in Figure 3. L and R refer to the example neuron on the left and right of each panel in Figure 3, respectively. See Methods for more details. Supplementary figures S3-S17 show the same analyses for each trial type, session and animal.

Figure 5. Population dynamics in the hippocampus are consistent across both seconds and minutes. a-c. Consistent population dynamics within seconds-long trials. Each element of the matrix is the cosine similarity between a pair of population vectors at different binned locations within a trial, average over pairs of trials, trial types (for b and c), sessions and animals. Critically, the population vectors are taken from the same time or location bin, but from *different* trials. See Methods for the details of how binning was performed for the different tasks. All three experiments (a: treadmill task, b: spatial alternation, c: linear track) show higher correlation along the diagonal, indicating that the population goes through a consistent sequence within each trial. This is as we would expect from the known properties of time cells (a) and place cells (b and c). d-f. Consistent **population dynamics** *across* trials. Similar to *a-c*, except each element of the matrix is the cosine similarity between a pair of population vectors from two different sessions, averaged over all pairs of sessions and animals. Rather than computing population vectors from bins of time or space within a trial, population vectors were computed by averaging over entire trials (see Methods for details). The similarity between population vectors from different recording session was then computed for different pairs of trial bins. The elements of all three matrices are highest on the diagonal and gradually decrease off-diagonal, similar to the matrices over a trial (**a-c**). This indicates that the population dynamics are also consistent across minute-long sessions.

Figure 6 . **Cross-session decoding for trial bin number.** An Linear Discriminant Analysis (LDA) classifier was trained to predict the trial bin number within a given session from the mean population neural activity during that trial bin. The classifier was trained on all but one sessions and tested on the left-out session. Middle: all three heatmaps exhibit a diagonal pattern, indicating that the classifier can correctly predict the trial bin number within the left-out session. The diagonal pattern in the heatmaps indicate successful decoding. Statistical significance was evaluated by computing a diagonalness metric (same as the one used for the heatmap of correlations. See Methods) and comparing with shuffled data. Right: the posterior probability for the correct trial bin given by the LDA classifier (each point represents an animal) compared to chance $(\frac{1}{5}= 0.2, \text{ dashed line})$.

Supplementary figure captions

Figure S1 . Additional hippocampal neurons that fire consistently across trials. Each line represents the z-scored transient rate of that neuron during a trial. Darker lines indicate earlier trials within a session. Trials where the neuron is not identified or is inactive are not plotted.

Figure S2 . Additional hippocampal neurons that fire consistently across multiple sessions. Each line represents the z-scored transient rate of that neuron during a session. Darker lines indicate earlier sessions. Sessions where the neuron is not identified or is inactive are not plotted.

Figure S3. The distribution of the across-trial firing consistency score for each individual session in the treadmill running task.

Figure S4. The across-trial firing consistency score plotted against the across-trial firing linearity score for each individual session in the treadmill running task.

Figure S5. The distribution of the across-session firing consistent score for real and shuffled data (top) and he joint distribution of the across-session firing consistency score and firing linearity score (bottom) for each individual mouse in the treadmill running task.

Figure S6. The distribution of the across-trial firing consistent score for real and shuffled data for each individual session and trial type in the spatial alternation task. Data for mouse Bellatrix.

Figure S7. The joint distribution of the across-trial firing consistency score and firing linearity score for each individual session and trial type in the spatial alternation task. Data for mouse Bellatrix.

Figure S8. The distribution of the across-trial firing consistent score for real and shuffled data (left) for each individual session and trial type in the spatial alternation task. Data for mouse Calisto.

Figure S9. The joint distribution of the across-trial firing consistency score and firing linearity score for each individual session and trial type in the spatial alternation task. Data for mouse Calisto.

Figure S10. The distribution of the across-trial firing consistent score for real and shuffled data. Data for mouse Nix.

Figure S11. The joint distribution of the across-trial firing consistency score and firing linearity score for each individual session and trial type in the spatial alternation task.

Data for mouse Nix.

Figure S12. The distribution of the across-trial firing consistent score for real and shuffled data. Data for mouse Polaris.

Figure S13. The joint distribution of the across-trial firing consistency score and firing linearity score for each individual session and trial type in the spatial alternation task. Data for mouse Polaris.

Figure S14. The distribution of the across-session firing consistent score for real and shuffled data (top) and the joint distribution of the across-session firing consistency score and firing linearity score (bottom) for each individual mouse in the spatial alternation task.

Figure S15. The distribution of the across-trial firing consistency score for each individual session in the linear track task. Data shown separately for left and right runs.

Figure S16. The across-trial firing consistency score plotted against the across-trial firing linearity score for each individual session in the linear track task.

Figure S17. The distribution of the across-session firing consistent score for real and shuffled data (top) and the joint distribution of the across-session firing consistency score and the firing linearity score (bottom) for each individual mouse in the linear track task.

Figure S18. The correlation between the firing consistency score on the timescales of seconds and minutes for the treadmill running (a), spatial alternation (b) and linear track (c) experiments. The same for the firing linearity score (d-f). Kendall's τ : **a**: τ (1245) = 0.13, $p < 10^{-8}$. **b**: τ (3019) = 0.079, $p < 10^{-8}$. **c**: τ (1099) = 0.01, p = 0.6. **d**: τ (1245) = 0.014, p = 0.5. **e**: τ (3019) = 0.12, $p < 10^{-15}$. **f**: τ (1099) = 0.011, p = 0.6.

Figure S19. Firing consistency ranks are not significantly correlated with x or y position in the field of view, either on the timescale of seconds (a, b), or on the timescale of minutes (c, d). A statistical test using Pearson's correlation coefficient was conducted. **a** top: r = -0.02, p = 0.12; bottom: r = 0.0004, p = 0.98. **b** top: r = -0.01, p = 0.68; bottom: r = -0.02, p = 0.43. **c** top: r = 0.003, p = 0.84; bottom: r = -0.02, p = 0.23. **d** top: r = 0.01, p = 0.67; bottom: r = -0.03, p = 0.14.

Figure S20. The distributions of the temporal information score are significantly biased towards larger values compared with shuffled data. See Methods section for the definition of the temporal information score.

Figure S21. The cross-trial correlations for each individual session in the treadmill running task.

Figure S22. The cross-session correlations for each individual mouse in the treadmill running task.

Figure S23. The cross-trial correlations for each individual session, task phase and turn direction in the spatial alternation task. Data for mouse Bellatrix.

Figure S24. The cross-trial correlations for each individual session, task phase and turn direction in the spatial alternation task. Data for mouse Calisto.

Figure S25. The cross-trial correlations for each individual session, task phase and turn direction in the spatial alternation task. Data for mouse Nix.

Figure S26. The cross-trial correlations for each individual session, task phase and turn direction in the spatial alternation task. Data for mouse Polaris.

Figure S27. The cross-session correlations for each individual mouse in the spatial alternation task.

Figure S28. The cross-trial correlations for each individual session and running direction in the linear track task.

Figure S29. The cross-session correlations for each individual mouse in the linear track task

Author Manuscript

Fig 1



r Manu

1 trial = 10 seconds running on the treadmill Exp 1: Treadmill running (Mau et al., 2018) SCTD 1 session ~ 20 mins 1 trial = study phase (blue) + test phase (red) Exp 2: Spatial alternation (Levy et al., 2019) 1 session ~ 30 mins 1 trial = 3 minutes of running Exp 3: Linear track Rest (3 mins) (Rubin et al., 2015) 1 session = 30 mins Fig 3 Seconds **Minutes** Animal 1, session 1 neuron 121 nal 1, session 1 neuron 314 Animal 1, neuron 34 Animal 4, neuron 418 ate transient rate 0.75 d. ate a. 1.0 0.50 transient v 0.5 0.25 Z-scored trans 0.00 0.0 Z-scored Z-scored t -0.2 3 4 5 6 7 8 9 10 Time bin 1 2 3 4 5 6 7 8 9 10 Time bin 3 Trial bin Trial bin Animal Calisto, session 1, study phase, right turn trials, start arm, neuron 472 Animal Calisto, session 1 study phase, left turn trials, start arm, neuron 1859 Animal Nix, neuron 2191 nal Calisto, neuron 804 rate 2 1.5 rate 0.6 rate e. b. 1.0 0. nsient Z-scored transient 2-scored transient 1 0.2 0.5 0.0 Z-scored tra 0.0 0 -0.5 -0.2 -1.0 1 2 3 4 5 6 7 8 9 10 Location bin 3 4 5 6 7 8 9 10 Location bin 3 Trial bin i ż Trial bin Animal C6M4, session 1, right runs, neuron 21 Animal c68m3, session 1, right runs, neuron 94 Animal C6M4, neuron 331 Animal c16m4, neuron 172 ate c. f. transient v Z-scored transient

2 3 4 5 6 7 8 9 10 Location bin

i

1 2 3 4 5 6 7 8 9 10 Location bin

3 Trial bin

Z-scored



Fig 5

2



Representations in timescales of seconds

-0.32 -0.24 cosine similarity

Cosine similarity

Cosine similarity

0.06

-0.03

-0.00

-0.18

0.12

0.06

0.00

-0.06

Representations in timescales of minutes



Cosine similarity

Fig 6









0.5

0.0

•

0.5

c.



