nature

SUPPLEMENTARY INFORMATION

https://doi.org/10.1038/s41586-018-0459-6

In the format provided by the authors and unedited.

Integrating time from experience in the lateral entorhinal cortex

Albert Tsao[1,3]*, Jørgen Sugar[1], Li Lu[1,4], Cheng Wang[2], James J. Knierim[2], May-Britt Moser[1] & Edvard I. Moser[1]*

[1]Kavli Institute for Systems Neuroscience and Centre for Neural Computation, NTNU, Trondheim, Norway. [2]Zanvyl Krieger Mind/Brain Institute, Johns Hopkins University, Baltimore, MD, USA. [3]Present address: Department of Biology, Stanford University, Stanford, CA, USA. [4]Present address: Department of Neuroscience, Baylor College of Medicine, Houston, TX, USA. *e-mail: albert.tsao@ntnu.no; edvard.moser@ntnu.no

Supplemental materials and methods

1. Single cell activity in LEC related to time

Mixed selectivity incorporating time

For the BW12 experiment, the GLM analysis showed that a fraction of cells in all three recorded brain areas exhibited selectivity for a mixture of time and non-time variables (8.4%, 7.3% and 4.7% for LEC, CA3 and MEC respectively, percentages averaged across all animals; **Fig. 1c**). We employed a stepwise model-fitting procedure, where the first variable to be added to the model is the one explaining the most response variance. In LEC, 15/40 of time/non-time mixed selective cells were selective first for time, while 2/11 of CA3 time/non-time mixed selective cells were selective first for time and 0/5 of MEC time/non-time mixed selective cells were selective first for time. For CA3 cells exhibiting mixed selectivity, 9/11 satisfied criteria for place cells, while for MEC cells exhibiting mixed selectivity, 2/5 were tuned for a spatial feature (1 speed cell and 1 grid cell). Thus, cells exhibiting mixed selectivity in LEC, CA3, and MEC had differing primary selectivities, with almost half of LEC cells being primarily selective for time. In total, the selectivity of LEC single-cell responses reflected the significant influence of time on LEC activity, in contrast to CA3 and MEC.

LEC single-cell time constants

The distribution of time constants for ramping cells spanned a broad range of time, from a few minutes to hundreds of minutes, and although a greater number of cells had negative time constants, the distribution was not significantly biased (52/80 LEC cells had negative time constants; one-sample *t*-test, $P = 0.46$; **Extended Data Fig. 2c**). There was not any apparent anatomical organization to the distribution of time constants between deep and superficial layers of LEC, as the absolute value of time constants in deep layers was not different than for superficial layers (mean for superficial layers: 48.9 min, mean for deep layers: 41.3; unpaired *t*-test, $P = 0.80$) and the distribution of time constants between deep and superficial layers was not different for trial-selective (short time constant) versus session-selective (long time constant) cells (two-proportion z-test, $P = 0.46$). Up-ramping and down-ramping cells were observed simultaneously on 7 out of the 17 instances where more than one ramping cell was recorded in a single session. The presence of cells with positive time constants, over a wide range of

timescales, suggests that the ramping activity we observe is not simply an effect of habituation[54], and may be generated by intrinsic mechanisms[55,56,57].

Ramping activity is not due to recording instability

One possible concern, given that the vast majority of time-selective cells encoded time across the entire session rather than across single trials, is that this apparent temporal coding merely reflects unstable recordings. This is unlikely, as only cells which passed spike waveform stationarity tests checking for stability across the entire recording session were included in analyses (**Extended Data Fig. 2e, f, see Methods**). Additionally, although cells firing during a few trials in the middle of a session may have reflected legitimate time-influenced activity, only cells which fired spikes in both the first and last trial or intertrial periods were included for analysis. Finally, for some cells showing significant changes in firing rate across trial periods (intertrial), activity during intertrial (trial) periods was relatively consistent across the session (**Extended Data Fig. 2g**), demonstrating that significant changes in activity across an experimental session were not simply an artifact of poor recording stability.

Time cell-like activity

In the hippocampus, time cells fire in relation to events[15,16] (e.g. the start of a delay period). Do such cells exist also in LEC? In the context of the BW12 experiment, synchronizing events may be distributed across the entire session without appearing in any obvious way to an observer. If this is the case, because we have not recorded the same cells across days, we would not be able to distinguish time cell activity from activity due to random fluctuations of external stimuli. However, it is also possible that time cells in the BW12 experiment are synchronized to the structure of the task, e.g. the beginning of trial or intertrial periods. If this is the case, we would expect to see repeated activity at the same time relative to the start of a trial or intertrial period. We tested the latter possibility in our data.

We began by ordering cells from individual animals by the latency of their peak firing rate and looking for sequential activity. Although a ridge of peak activity was present, consistent with sequential time cell-like activity, a similar ridge of peak activity was present for both actual and shuffled data across all areas, suggesting that time cell-like activity may not have been present in

LEC during the BW12 experiment (**Extended Data Fig. 3a, b**). The low ratio of peak-to-background activity, in contrast to that of time cells which are mostly silent outside of their time field[15,16], also suggested a lack of time cell-like activity. We quantified these observations by calculating a peak-to-background ratio. The peak was defined as the mean rate within a 10 s window centered on the time point of the cell's peak firing rate, and the background was defined as the mean rate across the rest of the session. Significance was determined by comparing to a distribution generated by calculating peak-to-background ratios for temporally shuffled data (1,000 shuffles per cell, shuffling was performed by time-shifting the firing rate of each cell by a random amount with the end of the trial wrapped around to the beginning) and using the 95th percentile as a significance threshold. The number of cells in LEC with significant temporal information during trial periods was below that expected by chance (number of cells with significant temporal information during trial periods for LEC: 3.5%, 16/451; CA3: 2.0%, 3/148; MEC: 7.5%, 10/133; mean peak time during trial periods for LEC: 67.7 s; CA3: 3 s; MEC: 25.1 s; number of cells with significant temporal information during intertrial periods for LEC: 3.8%, 17/451; CA3: 2.7%, 4/148; MEC: 7.5%, 10/133; mean peak time during intertrial periods for LEC: 58.9 s; CA3: 73.3 s; MEC: 108.9 s, **Extended Data Fig. 3c-e**). A number of cells did have significant peak activity at the beginning of trial periods, which may reflect actual time cell-like coding, given that time cells have been observed to have more accurate timing closer to the start of their sequences[16]. However, we are unable to distinguish between this and the possibility that these cells are merely responding to the change in context from holding pot to recording box.

In a separate quantification, we calculated a temporal information score for each cell, analogous to the commonly-used spatial information score[53]. For a given cell, temporal information was calculated as $\sum P_i \left(\frac{R_i}{R}\right) log_2(\frac{R_i}{R})$, where i is the time bin number, P_i is the probability of occupancy for time bin i (5-s time bins were used), R_i is the mean firing rate for bin i, and R is the overall mean firing rate, analogous to the commonly used spatial information score. Significance was determined by comparing to a distribution generated by calculating temporal information for temporally shuffled data (1,000 shuffles per cell, data was first binned into 1-s time bins, wrap-around shuffling was performed individually on activity from each trial or intertrial prior to taking trial or intertrial averages, and temporal information were then calculated by binning the trial/intertrial-averaged shuffled data into 5-s time bins). The 95th percentile of the

shuffled distribution was taken as a significance threshold. The number of cells in LEC with significant temporal information during trial periods was barely above that expected by chance (number of cells with significant temporal information during trial periods for LEC: 5.5%, 25/451; CA3: 4.7%, 7/148; MEC: 9.8%, 13/133; mean temporal information during trial periods for LEC: 1.11 bits/s; CA3: 3.36 bits/s; MEC: 0.53 bits/s; number of cells with significant temporal information during intertrial periods for LEC: 7.3%, 33/451; CA3: 6.8%, 10/148; MEC: 11.3%, 15/133; mean temporal information during intertrial periods for LEC: 0.67 bits/s; CA3: 1.61 bits/s; MEC: 0.47 bits/s, **Extended Data Fig. 3f-h**). Overall, the small number of temporally-selective cells detected by either method suggests a limited presence of time cell-like activity, with the exception of a few cells with activity locked to the start of trials. However, time cell-related activity may not necessarily be synchronized to the start of trials, and may instead by distributed across the entire experimental session, expressing time relative to any number of unaccounted-for events (e.g. random noises).

In addition to looking for time cell-like activity specifically, we also ran an expanded GLM which included single-trial ramps and symmetrical ramps, which have been suggested as inputs to time cells[31] (**Extended Data Fig. 3i**). Both predictors were synchronized to the start of trials, and the symmetrical ramps spanned a range of slopes, dependent on which trial was being synched to. These two additional predictors would reveal cells with increased activity within a specific range of time, whereas the trial and session ramping variables of the original GLM would identify cells with activity that changed continuously across the session. This type of activity would be more similar to time cells, which fire at specific points in time, albeit at a much shorter timescale. Many cells were classified as exhibiting single-trial ramps, raising the possibility that ramping cells in LEC provide input to time cells (**Extended Data Fig. 3i-l**).

2. Population activity in LEC related to time

Regressing for time using dimension-reduced data

Given that a progression of population states across time could be seen in the dimension-reduced visualization of LEC data, was it possible to regress for time using the first few principle components from this dimension-reduced data? Regression was carried out using the linear model $y = \beta_1 * PC1 + \beta_2 * PC2$, with *y* being time, *PC*1 and *PC*2 the first two principal

components, and β_1, β_2 the fitted coefficients. Regressing the first two principal components from PCA results for individual animals against time lead to significant fits for all areas, but substantially higher explained variance for LEC (70.9%, 10.5% and 12.9% mean explained variance for LEC, CA3, and MEC respectively; **Extended Data Fig. 4b, c**). The improved fit for all LEC animals is due to one of the top two PCs exhibiting the session-length exponential decay seen in PC1 of the example LEC fit, and suggests a significant impact of time on LEC activity during unconstrained behavior. For the specific example shown, the poorer fit towards the end of the session for LEC reflects the flattening of the first PC, while the oscillating nature of the fit for CA3 reflects rate remapping due to wall color being picked up by both PCs, and the predicted time for both CA3 and MEC being centered on the mean overall time reflects the minimization of error by the regression.

Decoding controls for across-day pooling and behavior

Given that decoding accuracy was determined using cell populations which were pooled across recording sessions, the concern arises that pooling cells across days would reduce the covariance amongst neurons and drive an artificial increase in decoding accuracy for temporal epochs. We first computed decoding accuracy for temporal epoch using groups of simultaneously-recorded cells, and found that all single-day decoding accuracies were significantly above chance (16.9% median decoding accuracy compared to 4.2% chance level, n = 46 single days). We then compared single-day decoding accuracy against a distribution of decoding accuracies formed using subpopulations with the same number of cells, but randomly drawn from the overall population pooled across days. This was done by calculating a z-score for the single-day decoding accuracy using the mean and standard deviation of the randomly-drawn distribution. Red line indicates 95th percentile boundary for null hypothesis that single-day decoding accuracy was not significantly lower than across-day decoding accuracy. Only 2 out of 46 single-day decoding accuracies were significantly lower than randomly-drawn decoding accuracies (mean z-score -0.39, **Extended Data Fig. 4d**). Comparison of randomly drawn data to data taken from consecutive pairs of days (middle, n = 72 days), or data taken from pairs of days which were separated by half of the overall number of recording days (right, n = 44 days), gave similar results (mean z-scores -0.24 and -0.19 for consecutive day-pairs and across-experiment day-pairs respectively, **Extended Data Fig. 4e**).

If decoding of temporal epoch were in part due to changes in animals' behavior across time, it would be possible to decode temporal epoch with accuracy greater than chance using just behavior tracking data. We trained separate decoders using behavior tracking data and compared against decoders trained using neural activity. Both tracking data and neural activity were binned into 500 ms bins. Decoding accuracy using just behavior tracking data did not exceed chance, indicating that the high decoding accuracy for temporal epoch observed using LEC population activity was not due to behavioral changes across experimental sessions (**Extended Data Fig. 4f**).

Decoding errors for CA3 and MEC reflect non-temporal coding

In addition to lower decoding accuracy for temporal epoch, both CA3 and MEC also exhibited significantly more erroneous predictions for trials with the same wall color than trials with different colors (14.9%, 36.1% and 23.9% mean percentage of same-colored errors for LEC, CA3, and MEC respectively; unpaired t-test, LEC vs. MEC: $P < 0.05$, LEC vs. CA3: $P < 0.01$; **Fig. 3c**), consistent with a role for CA3 and MEC in representing aspects of the environment that are constant over time.

Temporal information in CA3 and MEC is limited

Although decoding accuracy for temporal epoch was significantly higher for LEC than CA3 or MEC, decoding accuracies for CA3 and MEC were nonetheless greater than chance. Was this simply due to a more widely distributed temporal code in each of these structures, such that with a sufficient number of observed cells, both MEC and the CA areas would reach similar decoding accuracies as LEC? We examined this possibility in the BW12 experiment by quantifying decoding accuracy for temporal epochs as a function of population size, and found that CA3 and MEC began to plateau before reaching accuracies close to what LEC achieved (**Fig. 3d**). Extending population sizes by pooling cells across animals gave similar results (**Extended Data Fig. 4i**). Thus, although the decoding accuracies for time in CA3 and MEC were above chance, temporal information was greater in LEC, requiring significantly fewer cells to reach high decoding accuracy.

Additional evidence for temporal coding in LEC using BW4 data

Existing BW4 data offered us a chance to compare temporal coding properties across CA2 and CA1, which have been shown to exhibit more significant place cell decorrelation than CA3[22,24]. Decoding accuracies for both trial identity and wall color using matched BW12 data restricted to the first four trials were comparable, suggesting that there were no conspicuous differences between the two data sets besides the reduced number of trials (trial identity: 88.0% versus 83.1% mean decoding accuracy for BW4 vs. BW12; wall color: 91.2% versus 86.8% for BW4 vs. BW12; **Extended Data Fig. 4g**).

Subsequently, we compared decoding accuracy across all areas at several epoch lengths. Decoding accuracy for trial-length epochs using LEC data was significantly better than all hippocampal areas (mean accuracy for LEC: 75.1%; MEC: 39.8%: CA3: 44.4%; CA2: 51.6%; CA1: 51.1%; chance level 25%; **Extended Data Fig. 4h**). Decoding accuracy for 20-s epochs using LEC data was significantly better than all hippocampal areas (mean accuracy for LEC: 16.4%; MEC: 12.2%: CA3: 8.5%; CA2: 8.3%; CA1: 9.2%; chance level 0.4%; **Extended Data Fig. 6a**). Decoding accuracy was higher for LEC and MEC than CA3 for 10-s epochs (mean accuracy for LEC: 12.4%; MEC: 11.4%: CA3: 6.5%; CA2: 7.4%; CA1: 8.1%; chance level 0.2%; **Extended Data Fig. 6a**). Decoding accuracy was similar across all areas for 1-s epochs (mean accuracy for LEC: 0.54%; MEC: 0.99%: CA3: 0.76%; CA2: 0.65%; CA1: 0.82%; chance level 0.02%; **Extended Data Fig. 6a**). The parity between areas at shorter epoch lengths is potentially due to similar dynamics across all areas at these short timescales but may also reflect the fact that time cells in MEC and hippocampus are capable of coding for time at these short timescales.

Reduced temporal precision in LEC at shortest timescales

The differences in decoding accuracy across epoch-lengths suggest a basic limitation in temporal precision within LEC (**Extended Data Fig. 6c**). We examined whether temporal information was in fact reduced at shorter timescales compared to longer, trial-length timescales. First, we re-binned confusion matrices for shorter-epoch decoders to match the size of confusion matrices for trial/intertrial-length epoch decoders (e.g. for 10-s decoders, prediction counts within 25×25 bin areas and 14×14 bin areas were summed to match trial- and intertrial-epochs of the

trial/intertrial-length epoch decoder, so that the 39×12 classes of the 10-s decoder were converted into 24 classes, matching the trial/intertrial-length epoch decoder). We then calculated the decoding accuracy for this re-binned confusion matrix from its diagonal, and found that decoding accuracy was in fact decreased when epoch length was shortened, suggesting that decoders using shorter temporal epochs were comparatively worse than decoders using longer temporal epochs (mean accuracy for trial: 88.6%; 20-s: 76.6%: 10-s: 56.5%; 1-s: 24.9%; chance level 4%; **Extended Data Fig. 6c**). Even though decoding accuracy for shorter temporal epochs was worse, decoding errors may nonetheless have been close to the correct epoch. In order to examine whether errors were close to the correct epoch, we shuffled the confusion matrix entries for each row, but kept the diagonal original. We then applied the same re-binning procedure and compared decoding accuracies to those found using unshuffled data. If errors were close to the correct epoch, decoding accuracies should decrease for shuffled data compared to the decoding accuracies found using unshuffled data, which is in fact what we observed (**Extended Data Fig. 6c**).

Distributed population coding of environmental context in LEC

In addition to temporal information, it was clear from 2D visualizations of LEC population activity during the BW12 experiment that environmental context was also encoded within LEC population states (**Fig. 2**). The amount of information present about the current environmental context (B or W) could be measured using a linear binary classifier. Decoding accuracy for wall color was very high for LEC population activity (91.9% mean accuracy). Similar decoding accuracies were observed for CA3, suggesting that environmental context is represented robustly at the population level in both regions (78.6% versus 83.1% mean accuracies for LEC and CA3, matched population sizes; **Extended Data Fig. 4j**).

High decoding accuracy for wall color was not simply an artifact of the separation of population states across temporal epochs, as decoding accuracy using data which was shuffled temporally while maintaining the original wall color condition was similarly high, suggesting that wall color defines two regions in population state space, within which different temporal epochs are encoded (91.8 versus 91.9% mean accuracy for original and temporally shuffled data, **Extended Data Fig. 4k**). Information about wall color in LEC was also distributed across the population,

as wall color could be decoded with high accuracy (77.9% versus 91.4%, mean accuracy for wall color using population with no wall color-selective cells, compared to a size-matched randomly-drawn population; **Extended Data Fig. 4l**).

Information about the immediately preceding context was also present in LEC, as decoders trained on intertrial period activity could infer the preceding wall color with high accuracy when the data were labeled by the wall color of the previous trial (74.1% mean accuracy), in agreement with previous work[30,47]. Thus, while temporal information was encoded through a continuously-changing population state, other streams of information such as environmental context and recent experience had a significant impact on this changing population state, constraining it to a limited range for any given period of time.

3. LEC population dynamics

High decoding accuracy for temporal epochs could be caused by purely random fluctuations in LEC activity, raising the possibility that temporal information emerges from a stationary noise-related process (e.g. white noise) that may not have any actual connection to temporal coding. The presence of a fair number of ramping cells argues against this possibility. Nonetheless, we further examined the structure of neural trajectories by measuring the distance between population states across trial and intertrial periods to population states at specific points in time.

First, we measured the distance between population states at all time bins within a trial period or intertrial period to the middle time bin of the same trial or intertrial period (**Extended Data Fig. 5a**), the last time bin of the preceding intertrial or trial period (**Extended Data Fig. 5b**), or the first time bin of the same trial or intertrial period (**Extended Data Fig. 5c**). If LEC activity was generating temporal information purely through stationary noise-related processes, one would expect no difference in population dynamics at the start versus the end of a trial. If LEC activity was instead representing time as it related to experience, one might expect differential dynamics at prominent temporal landmarks such as the start and the end of a trial or intertrial periods, given the significant change in external environment at the start of a trial. Across all three measures, changes in distance were indeed larger at the start of trials than the end, suggesting that the temporal information did not arise purely from trivial noise sources.

For LEC and MEC, the distance between the population state for the first time bin and all subsequent time bins in the same trial or intertrial period could be fit by an autoregressive model (**Extended Data Fig. 5c**). A fifth-order autoregressive model was used, with the order determined by looking at the autocorrelations of the distance measures for all three areas. The LEC model incorporated significant terms for lags of 1, 4, and 5, suggesting that activity at time points of 1, 4, and 5 time bins prior to the current time bin influence current time. The MEC model incorporated a significant term for lag 1. Applying the autoregressive model to CA3 data did not lead to any significant terms. This influence of preceding activity may be due to decay of externally-driven activity that takes longer than the time bin length, with more lag terms implying that longer periods of preceding time are influencing current activity.

We also measured pairwise distances between consecutive population states, both directly (**Extended Data Fig. 5d**) and in terms of the angle formed between vectors defined by the consecutive population states (**Extended Data Fig. 5e**). For distance measured directly using the population states themselves, we found significantly longer distances at the start of epochs, which may reflect significantly more variability in activity when the animal is first placed in the recording box, as we observed a number of cells with higher firing rates for the first 5-10 seconds of each trial compared to the rest of the trial (data not shown). The short period of initial variability was followed by settling into a stable state, as indicated by the roughly uniform pairwise distances across the rest of the trial or intertrial periods.

For measuring pairwise angles across consecutive points in time, population state vectors were defined as the difference between consecutive population states, and the angle between them was calculated (**Extended Data Fig. 5e**). If trajectories evolved according to a stationary process (e.g. white noise), one would expect the distribution of angles across an epoch to be uniform. For LEC, we observed clustering around angles which were greater than 90 degrees although not quite reaching 180 degrees, indicating that LEC activity did not evolve as a stationary process. Both CA3 and MEC exhibited a wider range of angles, including instances which were less than 90 degrees (corresponding to a reversal in direction), suggesting that they were less structured.

Finally, we measured the distance between the overall mean population state for the first trial or intertrial period and the overall mean population states of all subsequent trial or intertrial periods for LEC (**Extended Data Fig. 5f**). A non-stationary progression was observed across trials. As a whole, these results indicate that population activity in LEC evolved in a non-stationary manner, as expected for a system expressing the evolution of time.

4. Temporal coding within a fixed environmental context

A robust temporal code for episodic experience should function across a range of external conditions. To what extent do the results observed during the BW12 experiment generalize to other contexts? We looked at LEC activity in a separate experiment where the environment was kept constant, and an object was present on a subset of trials (**Extended Data Fig. 7a**). First, we examined the activity of single cells using the same GLM method that we used for the BW12 experiment, incorporating the presence/absence of the object, trial time, and session time as variables. More than half of the cells we recorded were found to be selective for at least one of the variables used in the model (**Extended Data Fig. 7b-e**), with 31.6% of cells exhibiting selectivity for time and 8.0% exhibiting selectivity for the presence of the object (number of cells significantly influenced by at least one variable: 150/263, 57.0%; **Extended Data Fig. 7d**). For cells exhibiting mixed selectivity, 34/46 were primarily selective for time.

Next, we visualized LEC population activity and observed once again a separation of the three trials across the session, suggestive of temporal coding at the population level (**Extended Data Fig. 7f**). Decoders trained to determine trial identity using LEC population activity achieved near-perfect accuracy (98.7% mean accuracy, chance level 33.3%; **Extended Data Fig. 7g**). Object presence could also be decoded with near-perfect accuracy (98.4% mean accuracy, chance level 50%; **Extended Data Fig. 7g**). In addition to decoding trial-length epochs, it was also possible to decode shorter epoch-lengths within trials with accuracies well-above chance (14.7% mean accuracy for 20-s epochs, chance level 1.1%; 11.1% mean accuracy for 10-s epochs, chance level 0.6%; 1.1% mean accuracy for 1-s epochs, chance level 0.05%; **Extended Data Fig. 7h**).

Finally, we examined whether temporal information was distributed across the entire population, as observed in the BW12 experiment. We trained decoders for trial identity using only cells which were not characterized by the GLM to be selective for time. Decoding accuracy was reduced compared to decoding accuracy for size-matched populations which were randomly drawn from the full population, but still well above chance (84.2% versus 93.8% for data with time-selective cells removed vs. randomly drawn control; **Extended Data Fig. 7i**). Similarly, decoding accuracy for object presence using cells not selective for object presence was reduced but still well above chance (81.5% versus 96.9% for data with object-selective cells removed vs. randomly drawn control; **Extended Data Fig. 7i**). Consequently, the encoding of information within LEC was again distributed across the general population. Importantly, the GLM and decoding results together suggest that under a different set of external conditions than the BW12 experiment, time nonetheless significantly influences LEC activity and both the population and single-cell level.

5. Additional evidence that temporal information arises inherently in LEC

Within the BW12 experiment, repetition of experience was potentially present in the pairs of trials for which wall color remained constant. We examined whether the population response to trials with repeated wall color differed from trials with changing wall color by measuring the Euclidean distance between mean LEC population states for pairs of adjacent trials with different or same wall color. Mean population states were calculated by taking the population activity vector (dimensionality equal to number of recorded cells pooled from all sessions) for each 10-s bin, and averaging across all 25 of these vectors for each trial (trials 250 s long). Distances were z-scored to allow comparison between animals. Population states for pairs of consecutive trials with the same wall color in the BW12 experiment were in fact more similar than population states for pairs of trials in which wall color changed, consistent with temporal information arising inherently in LEC through changing experience (**Extended Data Fig. 9a**).

Along this line, we also measured the distance between the population state for the first time bin and all subsequent time bins in each trial (**Extended Data Fig. 10a**). We observed fluctuations in distance which may have reflected significant changes in activity corresponding to when the animal was encountering features within the task (e.g., time bin 5 corresponded to when animal

was at base of central stem, time bin 10 corresponded to when animal was turning onto the goal arm). The consistency of these fluctuations was again consistent with a change in dynamics compared to BW12 data: whereas activity evolved in a roughly uniform manner during free foraging in the BW12 experiment tracking the general flow of time, during performance of a stereotyped, learned behavior in the figure-eight task, activity instead became aligned to time within individual trials.

Finally, we examined the variability of temporal representation in LEC in a second experiment in which animals were trained to run back and forth on a circular track, alternating between clockwise and counter-clockwise laps for reward at a turnaround point (**Extended Data Fig. 9c**). As with the figure-eight experiment, animal behavior was much more structured than during free foraging in an open field, and was likely goal-oriented towards the reward location. Trials were aligned temporally by the time-point at which the animal reached the middle of the track halfway to the turnaround point, and trials consisted of activity spanning from 3 s before to 3 s after reaching the midpoint. Looking at 2D projections of LEC population states for each trial, there was a clear decrease in the separation between trials when compared to matched BW data (**Extended Data Fig. 9d**). Consistent with this, although decoding accuracy for trial identity across the entire session was above chance, it was far lower than for matched BW data (17.6% versus 30.5% mean accuracy, chance level 6.6%; **Extended Data Fig. 9e**). In conjunction with the figure-eight experiment, these results suggest that constraining behavior to be more stereotyped reduced the amount of separation between population states encoding different epochs of time.

6. Interpretation of discrepancies in past studies of LEC responses

Overall, our results provide an explanation for discrepancies in previous recordings from LEC. During free exploration, LEC cells often display inconsistent responses to objects and other local cues[26,27,58,59]. In contrast, reliable selectivity for task-related features including objects has been observed during experiments where animals performed a structured behavioral task[29,44]. Since the representation of physical stimuli encountered during free foraging (e.g. objects or walls) is mixed with a representation of time, inconsistent responses would be expected across repeated exposures to these stimuli. During structured behavior, the influence of time decreases as mixed

selectivity is reduced, leading to more consistent representations of task-related features. The mixing of time with other variables may be present in other brain areas as well, given that time is not important solely for episodic memory.

54. Brown, M. W., Wilson, F. A. W. & Riches, I. P. Neuronal evidence that inferomedial temporal cortex is more important than hippocampus in certain processes underlying recognition memory. *Brain Res.* **409**, 158-162, (1987).
55. Tahvildari, B., Fransén, E., Alonso, A.A., Hasselmo, M.E. Switching between "On" and "Off" states of persistent activity in lateral entorhinal layer III neurons. *Hippocampus* **17**, 257-63, (2007)
56. Tiganj, Z., Hasselmo, M. E. & Howard, M. W. A simple biophysically plausible model for long time constants in single neurons. *Hippocampus* **25**, 27-37, (2015).
57. Sohal, V. S. & Hasselmo, M. E. A model for experience-dependent changes in the responses of inferotemporal neurons. *Network-Comp Neural* **11**, 169-190, (2000).
58. Yoganarasimha, D., Rao, G. & Knierim, J. J. Lateral entorhinal neurons are not spatially selective in cue-rich environments. *Hippocampus* **21**, 1363-1374, (2011).
59. Neunuebel, J. P., Yoganarasimha, D., Rao, G. & Knierim, J. J. Conflicts between local and global spatial frameworks dissociate neural representations of the lateral and medial entorhinal cortex. *J. Neurosci.* **33**, 9246, (2013).