



# Ramping cells in the rodent medial prefrontal cortex encode time to past and future events via real Laplace transform

Rui Cao<sup>a</sup> , Ian M. Bright<sup>a</sup> , and Marc W. Howard<sup>a,1</sup>

Edited by May-Britt Moser, Norwegian University of Science and Technology, Trondheim, Norway; received February 28, 2024; accepted August 5, 2024

In interval reproduction tasks, animals must remember the event starting the interval and anticipate the time of the planned response to terminate the interval. The interval reproduction task thus allows for studying both memory for the past and anticipation of the future. We analyzed previously published recordings from the rodent medial prefrontal cortex [J. Henke *et al.*, *eLife* 10, e71612 (2021)] during an interval reproduction task and identified two cell groups by modeling their temporal receptive fields using hierarchical Bayesian models. The firing in the “past cells” group peaked at the start of the interval and relaxed exponentially back to baseline. The firing in the “future cells” group increased exponentially and peaked right before the planned action at the end of the interval. Contrary to the previous assumption that timing information in the brain has one or two time scales for a given interval, we found strong evidence for a continuous distribution of the exponential rate constants for both past and future cell populations. The real Laplace transformation of time predicts exponential firing with a continuous distribution of rate constants across the population. Therefore, the firing pattern of the past cells can be identified with the Laplace transform of time since the past event while the firing pattern of the future cells can be identified with the Laplace transform of time until the planned future event.

interval timing | neural dynamics | Laplace transform | motor planning

Imagine singing a favorite song. In the pause before the next line of lyrics, the memory of the previous line gradually recedes, while the planned time to sing the next line approaches. As time progresses, the precise moment to begin singing the next line becomes increasingly clear, culminating at the moment of vocalization. The ability to prepare and execute an action at a specific moment (i.e. motor planning) often requires us to track the flow of time internally, which is explicitly tested in interval reproduction tasks (1–4). In a typical interval reproduction experiment, subjects experience a target duration and must later reproduce the duration through planned actions (e.g. Fig. 1A). Just like the pause between two lines in a song, each moment throughout the reproduction phase is defined relative to a beginning point at  $t = 0$  that recedes into the past as  $t$  increases and an endpoint at  $t = T$ . At each moment after  $t = 0$  the endpoint is a distance  $T - t$  from the present, approaching from the future (Fig. 1B).

A number of hypotheses have been proposed to describe how we keep track of the time of future planned actions internally during tasks such as interval reproduction. Scalar expectancy theory (SET) proposed that time elapsed is tracked through an accumulator that counts the pulses generated by an internal clock after the beginning point (5). While models like SET primarily focus on behavioral data, later models incorporated neural data observed during timing tasks. There is a large animal recording literature that reports ramping activity during timing tasks; individual neurons monotonically increase or decrease their activity as time progresses (1, 4, 6–10). Therefore later theories implemented a single clock that accumulates the firing activities of a group of cells and reaches a threshold at the end of an interval (11, 12). More recently, it was proposed that the changing dynamics of neural population activity, in the form of sequentially activated cells (13, 14) or neural trajectories in low-dimensional space (6, 15), serves as an internal timing mechanism. Researchers applied recurrent neural networks (RNNs) to perform timing tasks (16, 17) and implemented neural-inspired constraints to those RNNs (18, 19) as a way to gain further insights into the population dynamics required for internal time tracking. Researchers specifically focused on restricting the dimension of RNNs by building low-rank RNNs and argued it led to improved generalization to untrained stimuli (20).

These diverse models all share the underlying assumption that the neural dynamic changes with a few characteristic time scales across neurons for a given interval. SET model assumes a single Poisson process as the internal clock (5). Accumulator models

## Significance

We experience the world as a continuous timeline with the past trailing behind and the future stretching before. Previous work suggests that the recent past is remembered using a continuous spectrum of time constants across neurons. Such continuous time constants construct a timeline not unlike our experience of the past. The interval reproduction task here allows for the examination of neural representations of both the time since past events and the time until future events. We found both the past and the future were represented with a continuous set of time constants. Critically, this observation allows us to map the neural population dynamics onto the real Laplace transform, providing theoretical insight into the structures the brain uses for planned actions.

Author affiliations: <sup>a</sup>Department of Psychological and Brain Sciences, Boston University, Boston, MA 02215

Author contributions: R.C. and I.M.B. analyzed data; and R.C., I.M.B., and M.W.H. wrote the paper.

Competing interest statement: M.W.H. is a cofounder of Cognitive Scientific AI, Inc., which could conceivably benefit indirectly from this publication.

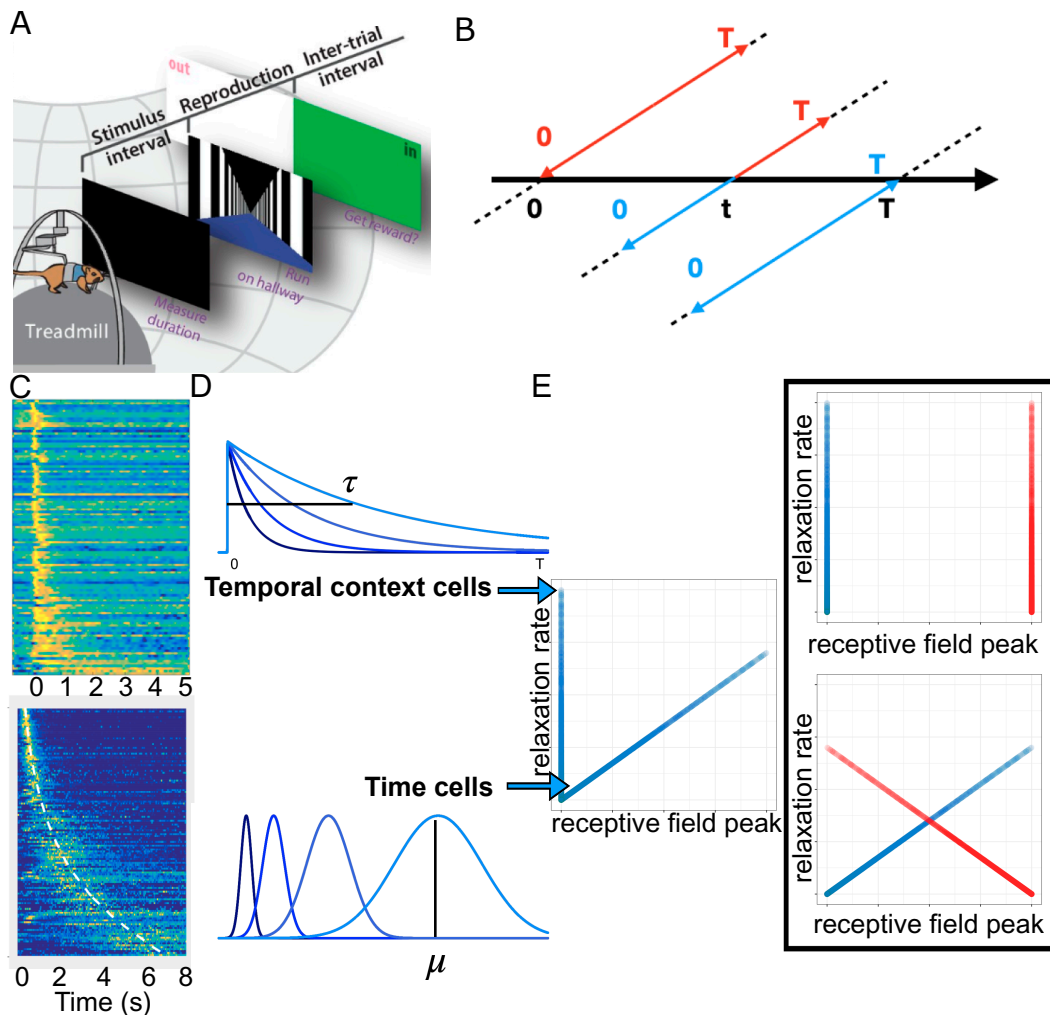
This article is a PNAS Direct Submission.

Copyright © 2024 the Author(s). Published by PNAS. This article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

<sup>1</sup>To whom correspondence may be addressed. Email: marc777@bu.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2404169121/-DCSupplemental>.

Published September 10, 2024.



**Fig. 1.** (A) Experimental task from ref. 6. On each trial, a gerbil experiences a target stimulus duration, and then reproduces the duration by running on a treadmill. Here, we limit our analysis to the neural activity during the reproduction phase (adapted from ref. 6). (B) Representation of the target interval in the reproduction phase. The horizontal line represents objective time, and the diagonal lines represent the target interval, which consists of time lapsed ( $t$ ) since the beginning of the interval (blue) and time remaining ( $T - t$ , red). The beginning of the interval is in the past while the end of the interval is in the future during the reproduction phase. (C and D) Time elapsed can be encoded by temporal context cells and time cells. (C) The *Top* panel shows previously reported temporal context cells, adapted from ref. 32; the *Bottom* panel shows time cells, adapted from ref. 31. In both panels, each row illustrates the normalized firing rate of a cell where yellow indicates high firing rates and blue indicates low firing rates. (D) Illustration of stimulated temporal context cells (*Top*) and time cells (*Bottom*). Time cells fire sequentially with widening receptive fields after an inciting event, resulting in a wide range of peak locations across cells. By contrast, a group of temporal context cells would reach their peak firing rate shortly after an event and relax back to baseline firing at different rates for different cells. Therefore, when receptive field peaks are plotted as a function of the relaxation rates in the *Left* panel, temporal context cells form a narrow vertical stripe. (E) Hypothetical properties of the neural populations during the reproduction phase with either temporal-context-cell-like representations (*Top*) or time-cell-like representations (*Bottom*). The temporal-context-cell-like representation would result in two vertical lines, one at the beginning and one at the end of the interval. The time-cell-like presentation would result in a wide range of receptive field peaks. The illustrations are presented symmetrically for easy visualization. The neural data are analyzed separately for the beginning and the end of an interval.

typically assume that the neurons are noisy exemplars of the overall rate, and the ramping rate reflects aggregated firing rates across neurons (12). For low-rank RNN models, the dynamics\* have as many time constants as the rank of the RNN has eigenvalues (1, 20). In all of those models, the typical time scale shared by the population only changes for different interval lengths.

Recent work on memory for the past has shown robust evidence that neural activity shows an effectively continuous set of time constants across neurons (Fig. 1 C and D). “Time cells” fire sequentially during a delay following an event (21, 22). Time cells have been widely reported in the hippocampus (21–24), but have also been reported in the medial entorhinal cortex (25), the

prefrontal cortex (23, 26, 27), and striatum (28, 29). The peak times of time cells smoothly tile the delay interval; as the sequence advances time cells fire for progressively longer durations (30, 31). Time cells thus exhibit a continuous distribution of characteristic time constants. “Temporal context cells” (Fig. 1 C, *Top*) also show a continuous distribution of time constants. Temporal context cells all respond shortly after some event, and relax their firing back to baseline at a variety of different rates (32–34). Rather than conveying information about how long ago an event happened via a continuous set of peak times (Fig. 1 D, *Bottom*), temporal context cells show a continuity of relaxation time constants (Fig. 1 D, *Top*).

The effectively continuous time constants in these two types of neural dynamics have similar properties. Within a time cell population, the density of peak times decreases linearly as

\*On the manifold defined by the low-rank recurrence matrix.

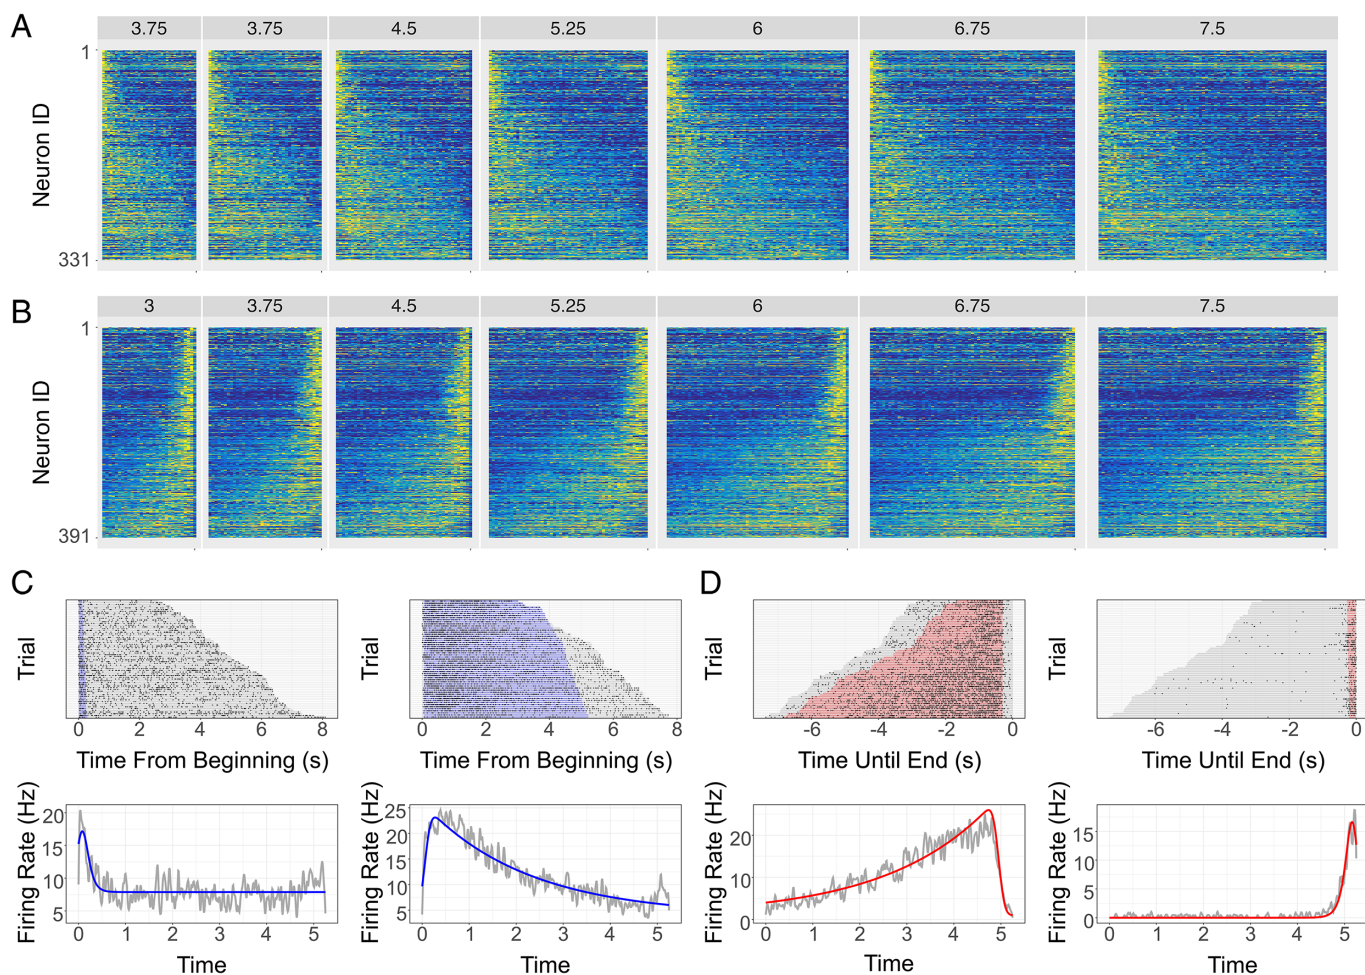
time progresses, while the size of the corresponding temporal receptive windows linearly increases. Consequently, the time of the triggering event is represented on a logarithmic time scale (31). The time constants of temporal context cells also overrepresent short time scales and have a smooth distribution with a long tail. Populations of temporal context cells are well described via exponential relaxation,  $e^{-st}$  with a broad range of real  $s$  (32, 33). Thus one can identify the firing of a population of temporal context cells with the real Laplace transform of the time since the triggering event was experienced (35).

To characterize the temporal information represented in a reproduction task using previously established methods for past time (31, 32), we analyzed recently published recordings (6) from the gerbil medial prefrontal cortex (mPFC) as they performed an interval reproduction task (Fig. 1A). During the measurement phase, the animals were presented with the target interval length  $T$ , sampled randomly from a set ranging from 3 to 7.5 s. During the reproduction phase, they would time their running to match the target interval length. We restricted our analyses to the reproduction phase to investigate the temporal representation of the past action (beginning of the interval) and the future action

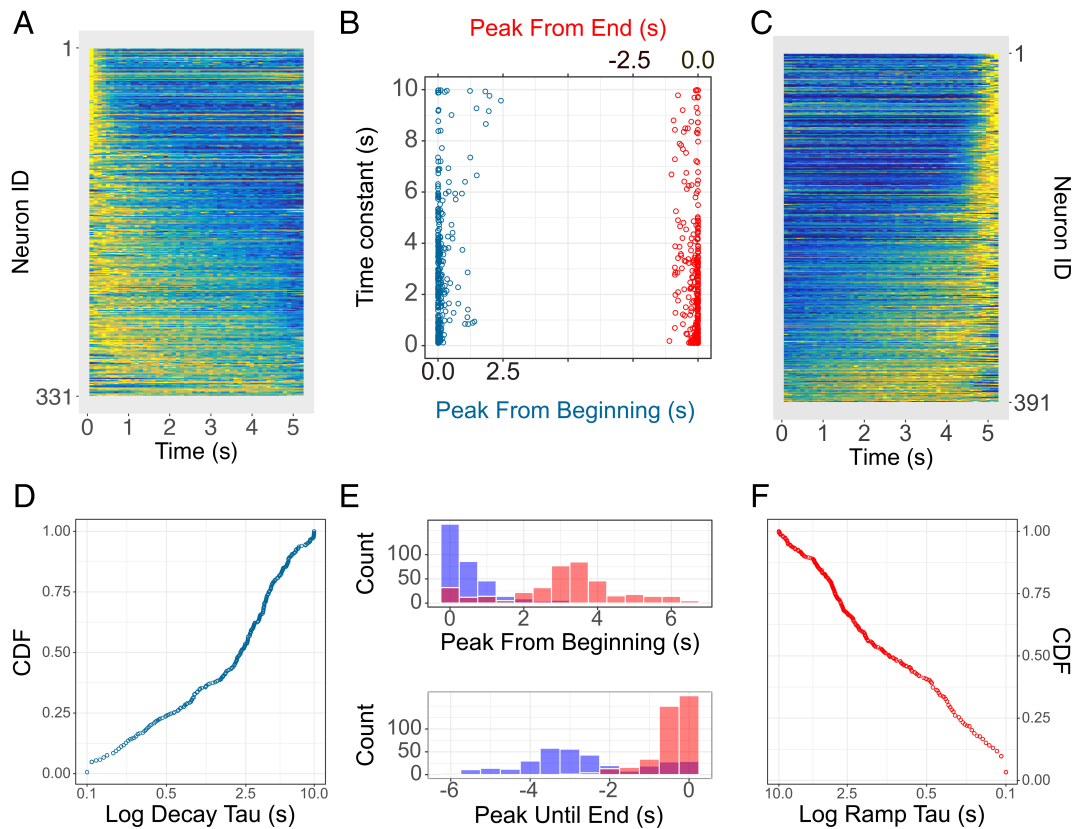
(projected end of the interval). We observed separate groups of units that represent either time since the beginning of the delay or time until the end of the delay. These two groups contained qualitatively similar properties in their respective representations: Both are characterized by a narrow range of peak values and a wide range of time constants, resembling the previously observed temporal context cells (32). Finally, we show that time in these groups is coded continuously, ruling out the possibility that time is represented via a small number of time scales.

## Results

A total of 1,766 mPFC cells were recorded (6). We first identified cells with reliable temporal receptive fields using previously established methods (32). The method revealed two groups of temporal-responsive cells: cells with temporal receptive fields preferentially peaking shortly after the start of the interval (331 cells see Fig. 2A) and cells with temporal receptive fields preferentially peaking shortly before the end of the interval (391 cells; see Fig. 2B). We refer to these populations as past cells and future cells respectively. These two groups can also be appreciated from



**Fig. 2.** (A) Heatmap of past cells during reproduction for each stimulus interval. Each row represents the normalized firing of a cell, linearly rescaled and averaged over all trials under the same stimulus length (marked for each column). Yellow indicates high firing and dark blue represents zero firing rate. Cells are sorted by the corresponding estimated time constant  $\tau$  during the 5.25 s interval. (B) Same as in (A), but for the future cells. (C) Two example cells and model fits from the past cells. On the *Top* panels of each cell, we plot the raster anchored at the beginning of each trial (marked zero) and sorted based on the reproduced interval length (indicated with gray shade). On the *Bottom* panels, we plot the average firing rate of the cell after its activity for every trial linearly rescaled to a 5.25 s interval with a gray line. The color shade in the raster plot marks the estimated receptive field for each trial and the colored line represents the estimated receptive field at 5.25 s from the model. (D) Same as in (C), but for the future cells and the raster plot is anchored at the end of each trial (marked zero).



**Fig. 3.** (A–C) mPFC cells tile the reproduced interval with a broad range of time constants from either the start or the end of the interval. (A) Normalized firing rate of past cells. Cells are sorted based on time constant  $\tau$  and plotted in the same way as figures in Fig. 2, except averaged over all the trials where every trial was linearly rescaled to 5.25 s. (B) Receptive field peaks  $\mu$  plotted as a function of estimated time constants  $\tau$  for both past cells (blue) and future cells (red). For past cells, peak location  $\mu$  was calculated from the start of the interval (blue), and for future cells,  $\mu$  was calculated from the end of the interval (red). (C) Normalized firing rate of future cells. Plotted in the same way as (A). (D) The cumulative density function (CDF) of time constants plotted on a log scale for past cells. Each point represents a cell in the group. (E) Distribution of peak firing locations, estimated without any modeling assumptions, relative to the start (*Top*) and the end (*Bottom*) of intervals. Blue represents the past cells and red represents the future cells. Note that past cells peak times cluster toward the beginning of the interval whereas future cells peak times cluster toward the end of the interval. (F) Same as (D) but for future cells.

Fig. 3 in the previous report of the dataset (6). We proceed to analyze the 722 temporal responsive cells with three hierarchical Bayesian models. We applied the best fitting model from the three to the remaining 1,044 cells and found those cells exhibit similar patterns to the temporal responsive cells (*SI Appendix, section 1*).

Using hierarchical Bayesian models, we tested three hypotheses for the temporal receptive field assumption. For all three models we estimated temporal receptive fields for the past cells as a function of time since the start of the interval; for the future cells, we estimated the temporal receptive fields as a function of time until the end of the studied interval,  $T - t$ . The Gaussian model describes each temporal receptive field via a  $\mu$  and  $\sigma$  for each cell. The exponentially modified Gaussian (ex-Gaussian) model additionally allows a skew to the temporal receptive field of each cell by adding a time constant  $\tau$ . The exponential model is the same as the ex-Gaussian model but with the  $\sigma$  fixed across cells to a small value (100 ms). Additional details about the models can be found in *Materials and Methods*.

The model fits were evaluated with Watanabe–Akaike information criterion (36). The exponential model provided a much better quantitative account of the data compared to Gaussian ( $\Delta\text{WAIC} = 3.36 \times 10^{10}$ ) and full ex-Gaussian ( $\Delta\text{WAIC} = 2.55 \times 10^{10}$ ). The difference is highly significant as it is more than 15 times the SE. Within the population, the majority of cells (647/722) are better fit by an exponential time field than by a Gaussian field according to the log likelihood results. The

log likelihood results and example cells fits from the exponential model and the Gaussian model can be found in *SI Appendix, Fig. S2*. Unless otherwise noted, all results for model parameters are from the exponential model. The fact that the exponential receptive fields fit the data better than Gaussian or ex-Gaussian argues against the hypothesis of sequential time-cell-like firing, which is reasonably well described by Gaussian receptive fields with  $\mu$  and  $\sigma$  covarying across cells (31).

The exponential model estimates temporal receptive fields with two key parameters: 1) the peak parameter  $\mu$  and 2) the time constant parameter  $\tau$ , which indicates the time it takes for the cell to regress 63% back to its baseline firing. Time-cell-like sequential activation would result in a wide range of estimated peak locations  $\mu$  as the sequence continuously tiles the delay (Fig. 1 *E, Bottom*). Temporal context cell populations would have a narrow range of peak locations  $\mu$  but a wide range of time constants  $\tau$  that continuously tile the delay (Fig. 1 *E, Top*).

To account for the changing temporal receptive fields for different interval lengths, we set the time constant  $\tau$  for  $T = 5.25$  s as the baseline and fit the firing pattern at each trial directly by assuming that the time constant for each trial,  $\tau_i$  may change according to the reproduced interval length  $T_i$ . Specifically, we assumed  $\tau_i = \tau(1 + \rho \frac{T_i - 5.25}{5.25})$ , where  $\rho$  determines how much each cell rescales its temporal receptive field according to interval lengths. When  $\rho = 0$ , the cell codes for absolute time and does not change its temporal receptive field for different interval

lengths while  $\rho = 1$  means the cell linearly rescales its time constant for the interval lengths. In addition, the peak firing rate was estimated with a parameter  $a_1$ , which estimates the “threshold” of the future cells just before the end of each interval. We report the results of  $\rho$  and  $a_1$  parameters in *SI Appendix, section 4*.

Example fits from the Exponential model are shown in Fig. 2 C and D. Additional example fits can be found in *SI Appendix, Fig. S3*. While the model successfully captured the firing pattern for the majority of cells, it missed some irregularities in a small subset of cells, such as dual peaks. *SI Appendix, Fig. S3C* includes examples of cells that were poorly fit by the exponential model.

### The Temporal Receptive Fields of mPFC Cells Are Characterized by a Narrow Range of Peak Locations and a Broad Range of Time Constants.

Fig. 3A summarizes the temporal response properties of the 331 past cells. To facilitate inspection of the overall firing pattern of each cell across all intervals, we linearly rescaled cell activities for each trial to 5.25 s and plotted the averaged activity for one cell in each row. The figure demonstrates that almost all the cells reached peak firing shortly after the start of an interval but different cells relax back to baseline firing at different rates within the interval. Fig. 3C summarizes the temporal response properties of the 391 future cells. Almost all the future cells reached peak firing shortly before the end of the interval. Different cells ramped up their activity at different rates within the interval, mirroring the pattern for past cells. Distributions of peak locations  $\mu$  and time constants  $\tau$  estimated through formal modeling of individual cells from the two groups confirmed these qualitative observations (Fig. 3B). We reported and performed subsequent statistical analysis on the most frequent observations from the posterior distribution.

**The Peaks of Past Cells Cluster at the Start of the Interval.** As can be appreciated from Fig. 3B, the estimated peak locations  $\mu$  for past cells were distributed over a narrow range of small values (median = 0.01 s, interquartile range = 0 to 0.10 s, 90th percentile = 0.40 s). Indeed most of the past cells reached peak firing rate shortly after the beginning of the interval (303/331 cells peaked within 0.5 s). This pattern was true within each animal: For the 76 cells from animal 1, the median of estimated peak locations was 0.07 s with 90th percentile = 1.01 s; for the 156 cells from animal 2, the median of estimated peak locations was 0.01 s with 90th percentile = 0.39 s; for the 99 cells from animal 3, the median of estimated peak locations was 0.01 s with 90th percentile = 0.19 s. In addition to model parameters, we directly calculated the peak locations from the spike train data, defined as the highest firing rate within 100 ms intervals (Fig. 3E, blue). To aggregate across all trials, we aligned the trials either with the beginning (*Top*), or the end (*Bottom*) of the reproduced intervals. Consistent with the estimated peak parameters, the peak firing bins for the past cells clustered at the beginning of the interval (median = 0.3 s with 90th percentile = 1.4 s). This provides strong evidence against the hypothesis that time elapsed since the start of the interval is coded by a group of sequentially firing cells that tile the interval as one would expect for a population of time cells.

**The Time Constants of Past Cells Are Distributed Over a Broad Range.** Contrary to the narrow range for  $\mu$ , the estimated time constants  $\tau$  for the past cells resulted in a wide range of values (median = 2.25 s, interquartile range = 0.59 to 4.00 s, 90th percentile = 6.31 s; see Fig. 3B). This pattern was observed for each animal: For the 76 cells from animal 1, the median of estimated time constant was 1.92 s with interquartile range from 0.63 to 3.60 s; for the 156 cells from animal 2, the median of estimated time constant was 2.07 s with interquartile range from 0.53 to 3.57 s; for the 99 cells from animal 3, the median of

estimated peak locations was 3.54 s with interquartile range from 0.93 to 5.28 s. In fact, the estimated  $\tau$  for some cells reached the far end of the parameter search boundary (0.01 to 10 s), suggesting the time constants for those cells could include values beyond 10 s. We eliminated cells with  $\tau > 9$  s (16/331 cells were excluded) when evaluating the correlation between  $\mu$  and  $\tau$  with Bayesian Kendall's test. For cells where an accurate  $\tau$  value can be measured, we found evidence in favor of the null hypothesis (Kendall's  $\tau = 0.069$ ,  $BF_{01} = 2.5$ ), indicating that there was no systematic relationship between  $\mu$  and  $\tau$  for past cells.

**The Peaks of Future Cells Cluster Closely Near the End of the Interval.** Similarly to the peak location of the past cells, the estimated peak locations for future cells (*Right* side of Fig. 3B) were distributed over a narrow range of values shortly before the end of the interval (median =  $-0.02$  s, interquartile range =  $-0.01$  to  $-0.13$  s, 90th percentile =  $-0.47$  s). 354/391 future cells reached peak firing within  $-0.5$  s of interval length  $T$ . This pattern was true within each animal: For the 67 cells from animal 1, the median of estimated peak locations is  $-0.01$  s with 90th percentile =  $-0.12$  s; for the 194 cells from animal 2, the median of estimated peak locations is  $-0.02$  s with 90th percentile =  $-0.66$  s; for the 130 cells from animal 3, the median of estimated peak locations is  $-0.05$  s with 90th percentile =  $-0.39$  s. Fig. 3E plots the peak firing bins for the future cells in red. The peak firing bins for the future cells clustered near the end of the interval (median =  $-0.3$  s with 90th percentile =  $-1.1$  s). This provides strong evidence against the hypothesis that time remaining before the end of the interval is coded by a group of sequentially firing cells that tile the interval.

**The Time Constants of Future Cells Were Distributed Over a Broad Range.** Just like past cells, we observed a wide range of time constants  $\tau$  for future cells: the median of this group was 1.03 s with interquartile range between 0.26 s and 3.00 s and the 90th percentile was 5.39 s (Fig. 3B). This pattern was true within each animal: For the 67 cells from animal 1, the median of the estimated time constant was 1.52 s with interquartile range from 0.51 to 2.71 s; for the 194 cells from animal 2, the median of estimated time constant was 0.60 s with interquartile range from 0.29 to 2.87 s; for the 130 cells from animal 3, the median of estimated peak locations was 1.70 s with interquartile range from 0.18 s to 4.19 s. Similarly to past cells, there were a few future cells with time constant  $\tau$  near the far end of the boundary for the future cells (10/391 cells were excluded), and were eliminated from the subsequent Bayesian Kendall's correlation test. We found no evidence supporting a correlation between the  $\mu$  and  $\tau$  using the Bayesian Kendall test (Kendall's  $\tau = 0.06$ ,  $BF_{01} = 3.44$ ) for cells where an accurate  $\tau$  value can be measured.

### Both Time Elapsed and Time Remaining Show Continuous Time Constants.

It is of tremendous theoretical importance to know whether the time constants of past cells and future cells are continuous or if they have a characteristic time scale. Visual inspection of the time constants for both past cells and future cells suggests that time constants are distributed smoothly over the parameter space (Fig. 3B). To formally test this observation, we fit the observed time constants with two types of commonly used distributions: normal distribution and power-law distribution. The normal distribution is characterized by its expected value  $\mu$  plus noise  $\sigma$  distributed symmetrically around the expected value. While the distribution is defined from  $-\infty$  to  $\infty$ , the probability of observing a value that is two  $\sigma$  from the mean is low ( $<5\%$ ). Power-law distribution assumes the probability of observing  $x$  follows  $p(x) \propto x^{-\alpha}$ , suggesting that there is a high concentration of small values while larger values become less frequent

monotonically. The distribution is characterized by its long tail where the probability of an atypical observation, no matter how far away from the concentrated small values, never asymptotes to zero. The size of the tail is determined by  $\alpha$ , where a smaller  $\alpha$  indicates a larger portion of the observations falls under the tail. Because the long tail will continuously add more area under the density function for those less frequent but still possible (in the statistical sense) observations, power-law is only defined as a distribution between a minimum value and a maximum value for  $\alpha < 2$ . In our case, a power-law distributed time constant means larger time constants gradually and continuously become less frequent without a clear cut-off point or distinct clusters within the bounds. Therefore one could not separate power-law distribution with a characteristic value plus noise like the normal distribution.

The time constants from the past cells and future cells were fitted separately. The fit was evaluated based on WAIC (36). We found that the power-law distribution provided a significantly better fit than the single normal distribution for both the past cells (WAIC for power law = 1,257, WAIC for normal distribution = 1,565) and the future cells (WAIC for power law = 1,133, WAIC for normal distribution = 1,819). The difference between power law and normal distribution is significant as it is about 10 times the SE for both groups. We further tested whether time constants should be quantified with a mixture of *two* normal distributions and found the fit improved marginally from the fit of a single normal distribution, but was still significantly lower than the power law fit for both the past cells (WAIC for mixed normal distributions = 1,345) and the future cells (WAIC for mixed normal distributions = 1,435). This is striking as the mixed normal distribution has five free parameters while the power law distribution only has one free parameter, the shape parameter  $\alpha$ . These results provided strong quantitative evidence against the hypothesis that there exists one or few typical time scales in the population. When the power-law exponent  $\alpha = 1$ , the density function becomes  $p(x) \propto x^{-1}$ , where  $x$  are distributed according to log scales. We found the mean of the posterior of the exponent for past cells was 0.76 with a 95% credible interval [0.68, 0.84]. The mean of the posterior of exponent for future cells was 1.05 with 95% credible interval [0.97, 1.12]. Both of the posteriors are close to 1, suggesting that the time constants were distributed close to logarithmic compression. When plotted against a log scale, logarithmically distributed observations would form a straight line (31). We plot the cumulative density function (CDF) on a log scale for the past cells (Fig. 3D) and the future cells (Fig. 3E). The curve for future cells at least is close to being straight, consistent with a power-law exponent of one.

It is worth noting that the present study is not well suited to estimate the precise value of  $\alpha$ . As demonstrated in *SI Appendix, section 4*, a large portion of the population rescale the temporal receptive field with the length of reproduced interval. While our model controlled for the rescaling with additional parameters, it is possible that there is residual trade-off between the estimated  $\tau$  and the rescale parameter  $\rho$ , especially for small  $\tau$ . An experiment with a fixed reproduction interval would allow for a more accurate estimation of  $\alpha$ .

## Discussion

We analyzed previously published recordings from mPFC during an interval reproduction task (6) using hierarchical Bayesian models to estimate the temporal receptive fields of individual cells and the populations. Consistent with previous research on interval timing (4, 11), we observed cells preferentially reach

peak firing either when the animal started running, denoted as past cells, or when the animal was about to stop, denoted as future cells. Consistent with previous work (7), we saw strong evidence against a time-cell-like sequential temporal representation in mPFC cells in these data. We observed exponentially decaying temporal receptive fields for past cells and exponentially ramping temporal receptive fields for future cells. Critically, both populations showed a continuous distribution of time constants over a wide range. The firing pattern of the past cells resembled temporal context cells from the entorhinal cortex (32, 33); future cells behave like a mirror image of the past cells. The pattern of firing of temporal context cells resembles the real Laplace transform of the time since the interval began (34). Whereas the population of past cells codes for the Laplace transform of the time since a past event, the population of future cells codes for the Laplace transform of the time until a future event. Together the two groups form a continuous timeline encoded in the Laplace domain, as predicted by theory (37).

There are several limitations to the present study. First, it should be made clear the results presented in this paper do not straightforwardly generalize to the measurement phase data. We applied an ex-Gaussian model to analyze these data and, despite our best efforts, identified no clear pattern nor systematic relationship between the parameters of neurons in the measurement phase and the reproduction phase. It is possible that the measurement data encode stimulus intervals at the population level, as suggested in the original study (6), or that the information about the duration of the measurement phase is stored in another brain area, such as the medial temporal lobe. Second, the present study is not well-suited to distinguish subtleties between hypotheses, such as whether the neural population represents the interval as a whole or represents the temporal information in relation to a past event and a planned future action separately. Additionally, it is difficult to determine whether the animal was planning for timed future actions or replaying the memory of the stored interval during the measurement phase. These variables are highly correlated in this study. An experiment where the reproduced interval could be dissociated from the trained interval in memory, perhaps by instructing animals to produce not the interval  $T$ , but some factor of the interval, would allow these two variables to be dissociated (15).

## Laplace Transform As a Basis Set for Neural Computation.

The results observed here add to a growing body of work showing exponential receptive fields with a continuity of time constants in a broad range of brain regions serving different functions. Previous work has shown memory for past events with exponential receptive fields with a variety of time constants in the entorhinal cortex (32) and precuneus (34), regions believed to be important in episodic memory encoding and retrieval. The future cells observed here resemble neurons in the frontal motor cortex during the waiting period before execution of a planned lick (38). Dopamine neurons in the midbrain are believed to be important for learning associations between past events and predicted future rewards (39). Recent evidence shows that the firing of dopamine neurons during presentation of a cue codes for the time of future outcomes (40, 41). As the time  $\tau_o$  between a cue and an outcome is increased, the firing of dopamine neurons in response to the cue decreases exponentially as  $e^{-s\tau_o}$ . Critically, researchers observed an effectively continuous distribution of rate constants  $s$  across different dopamine neurons. This is just as would be expected if the firing of dopamine neurons expressed the time of future outcomes via real Laplace transform (37, 41–43).

The present study is thus part of a larger literature showing evidence for real Laplace transform of the time of past events and the time of predicted future events. Why would the brain commit to this coding scheme for the time of events in such diverse brain regions and tasks? One answer is that Laplace transform provides a temporal basis set to represent functions of time. All possible functions over time can in principle be represented as linear combinations of exponential functions with different time constants. Moreover, properties of the Laplace transform that make it useful in the engineering curriculum—efficient formulae for data-independent operators, turning differential equations into algebraic equations—may also be useful for neural computation (44).

**Constructing Laplace Transform of the Past and Future with Recurrent Neural Networks.** The results presented here present two challenges for RNNs that have been extensively studied as a model of brain dynamics supporting cognition. First, the effectively continuous time constants of neurons in the network present a substantial challenge for low-rank RNNs. Second, although the Laplace transform of functions of time can readily be generated using diagonal RNNs, the roughly exponential growth of the future cells present a substantial challenge for these high-rank RNNs. Finally, we review recent work proposing a continuous attractor neural network that is able to describe these dynamics.

**Challenges for low-rank RNNs.** Low-rank RNNs have been proposed as models for timed behavior. The recurrent matrix of low-rank RNNs is composed of a few randomly chosen eigenvectors. The assumption of low-rank recurrence allows the mathematical properties of the network to be thoroughly understood (20). However, the eigenvalue spectrum of an RNN controls the number of distinct time constants it can exhibit. We observed an effectively continuous distribution of time constants for both past cells and future cells. The empirical observation of an effectively continuous distribution of time constants is difficult to reconcile with low-rank recurrence.

The linear dimensionality of a population with a continuous distribution of time constants is in principle unbounded. Consider the spanning dimension for a population of past cells measured from the moment after the population peaks out to time  $t$ . The spanning dimension to time  $t$  should go like  $\int_0^t p(\tau) d\tau$  where  $p(\tau)$  is the probability of observing a time constant  $\tau$ . If  $p(\tau)$  were uniform—which is clearly falsified by the results in the current paper—the spanning dimension would go up linearly with  $t$ . If the time constants were distributed on a logarithmic scale  $p(\tau) \sim \tau^{-1}$ , then the spanning dimension would increase like  $\log t$ , a decelerating function that grows without bound. Studying neural data from working memory experiments in monkeys, Cueva and colleagues observed that the dimensionality spanned by neural population indeed increased sublinearly with recording time (18). The present paper suggests that, rather than saturating at the rank of recurrent matrix, the dimensionality spanned by the neural population should grow without bound as the recording duration increases (42).

**Diagonal RNNs for Laplace transform are falsified by exponential growth.** It is straightforward to write down an RNN that gives out a population of exponentially decaying cells with a continuous spectrum of time constants as a solution (45). For instance, a population of past cells could be simply constructed from a recurrent connection with  $\frac{dx}{dt} = -\mathbf{R}x + f(t)$ , where  $x(t)$  is the recurrent state and  $f(t)$  is the external input at time  $t$ . If  $\mathbf{R}$  is a diagonal matrix with the appropriate values of  $s$  along

the diagonal, then one obtains a continuous eigenvalue spectrum and  $x(s, t) = \int_{-\infty}^t e^{-st} f(t') dt'$ , the Laplace transform of the past leading up to the present. There are a number of ways to implement slow time constants within neurons (46, 47) and via network effects (48), so this diagonal RNN is a reasonable computational model to construct Laplace transform of the past.

However, diagonal RNNs are ill-suited to describe exponential growth as required by the future cells described in this paper. Even tiny amounts of noise in the initial state would be amplified by exponential growth. After a short period of time, the pattern of activation across different cells would no longer be reliable or coherent. Exponential growth with a spectrum of time constants, i.e., the future cells observed in this paper, requires something more subtle than a linear RNN.


**Continuous attractor networks for Laplace transforms of delta functions.** The results of the present paper suggest that memory for the past and predictions of the future are both described as Laplace transform of a delta function at a particular point in time. That is, at time  $t$  into the delay, the function representing the time at which the reproduction phase starts is a delta function at time  $t$  in the past. If  $f(\tau) = \delta(\tau - t)$ , then the Laplace transform of  $f$  is  $\mathcal{L}\{f\}(s) = e^{-st}$ , resulting in exponential decay as the start of the reproduction phase recedes further into the past. Similarly, a firing rate of  $e^{-s(T-t)}$  across neurons is what one would expect from the (real) Laplace transform of a delta function at time  $T - t$  in the future.

If a neural network does not need to represent the Laplace transform of every possible function, but simply to represent the Laplace transform of a delta function, then the problem is much more constrained. If the goal is to generate a logarithmic distribution of time constants, as observed empirically in time cells in the hippocampus (31) and consistent with the value of  $\alpha$  observed for the future cells in this paper (Fig. 3F), then the network exhibits translation equivariance, making it a strong candidate for continuous attractor neural networks.

Fig. 4 illustrates the basic insight. Suppose that the values of  $s$  in the population are distributed logarithmically, such that  $s_n$ , the value of  $s$  for the  $n$ th cell, obeys  $n \propto \log \frac{s_n}{s_0}$ . Imagine that we can observe the activity of all the neurons simultaneously and plot them as a function of  $n$ . Now,  $e^{-s_n t}$  as a function of  $n$  appears as an edge for any particular value of  $t > 0$  (within bounds specified by the minimum and maximum values of  $s$ ). Changing the value of  $t$  simply shifts the location of the edge. Because of the logarithmic distribution of time constants, the rate at which the edge moves depends on its current location.

Continuous attractor neural networks (CANNs), such as those believed to map onto the head direction circuit in *Drosophila* (49), excel at allowing a particular shape of activation to translate smoothly along a network. The Laplace transform of a delta function in time requires a CANN with 1) an edge (of a particular shape) as a solution 2) the rate at which the edge moves depends on its location. In this case, the primary distinction between representing the past and the future is the direction and rate at which the edge moves. A recent paper (50) developed a minimal CANN to represent the Laplace transform of delta functions.

**Future Time Cells Can Be Constructed from Future Temporal Context Cells.** While we observed a Laplace domain timeline in the present study, it is entirely possible, expected even, that other brain regions maintain sequentially coded information about the time of future events, i.e., “future time cells.” Indeed, there is some evidence for such a phenomenon. Previous researchers observed highly sequential dynamics in the dorsolateral striatum



**Fig. 4.** (A) The interval begins at  $t = 0$  and ends at  $t = T$ . During this period the memory for the beginning of the interval recedes from the present toward the past and the expectation of the end of the interval approaches from the future closer to the present. (B) The plots on the *Left* show  $e^{-st}$  for several choices of  $t$  as a function of  $\log s$ . The plots on the *Right* show function  $e^{-s(T-t)}$  for several choices of  $t$  as a function of  $-\log s$ . The value of  $t$  is communicated by the shade of the line, corresponding to the moments in (A). Note that the shape of the pattern of activity over neurons does not change, for either the past or the future. Thus a continuous attractor network would be able to account for these results. This is only possible if time constants are evenly distributed on a log scale. Although the difference between the values of  $t$  is the same, the distances between adjacent edges are not the same. This is because  $s$  is on a logarithmic scale. Note that whereas the distance between successive lines becomes progressively smaller as the edge recedes into the past, the distance between successive lines becomes progressively larger as the future approaches.

while mice performed a two-interval timing task (14). Recording from the cerebellum, Wagner and colleagues observed sequentially activated cells tile time until an anticipated outcome (51). In the hippocampus and related regions, neuroscientists have found sequential codes that appear to tile distance in physical space leading to a goal location (52, 53).

Future cells observed here could be used directly to construct future time cells, cells that fire in sequence coding for the time until a planned event. Work on temporal coding of past events provides a strong analogy for this. It has been shown that temporal context cells, exponentially decaying cells with a continuous spectrum of time constants, can be used to generate time cells, sequentially activated cells with a continuous spectrum of peak times (54–56). The CANN described above greatly simplifies the task of inverting the Laplace transform. If the Laplace transform of a delta function is described as an edge across a network (Fig. 4), then the inverse transform is well-described by a bump that moves with the edge (50). More broadly, whatever neural circuit takes temporal context cells in the entorhinal cortex to time cells in the hippocampus (assuming for the sake of argument that these phenomena are causally related), would generate future time cells if given the future cells observed here as input.

## Materials and Methods

Because we took previously published recording data, the specific details about the data acquisition can be found in the original publication (6). In the section below we describe the computational analysis applied to the data in detail.

**Data Analysis.** The goal of these analyses is to identify parameters that best describe the properties of time fields of temporally responsive cells at both the trial and population levels. First, we identify and characterize temporally

responsive cells using previously established methods (27, 32). Then we apply Bayesian models to capture the temporal properties across trials and of the population.

**Selecting temporally responsive cells as input to hierarchical bayesian model.** In order to initially identify temporally responsive cells, spikes were first analyzed in Matlab 2016a using a custom maximum likelihood script. This approach has previously been used to identify both time cells and temporal context cells (27, 32). Fits were restricted to spikes that occurred during the reproduction phase. We fit the spike trains to five models. The first model was a constant firing model, which indicates a lack of temporal responsiveness. The constant firing model,

$$M_{\text{CON}}(t; a_0) = a_0 \quad [1]$$

contained a single parameter  $a_0$ , allowed to range between  $10^{-7}$  and 1, which predicted a constant probability of firing a spike at time  $t$ .

We then considered an exponentially modified Gaussian model, which describes the firing field as a convolution of an exponential distribution and a Gaussian distribution. This model was fit to both time elapsed since the start of the reproduction phase and time remaining until the end of the reproduction phase. The exponentially modified Gaussian model for time since the start of the reproduction phase,

$$M_{\text{EMG}}(t; a_0, a_1, \sigma, \mu, \tau) = a_0 + \frac{a_1}{2} e^{\frac{(2\mu + \frac{\sigma^2}{\tau} - 2t)}{2\tau}} \operatorname{erfc} \left( \frac{\mu + \frac{\sigma^2}{\tau} - t}{\sqrt{2}\sigma} \right) \quad [2]$$

contained the parameters  $t$ , which tracked time since the start of the reproduction interval,  $a_0$  (allowed to range between  $10^{-7}$  and 0.5) describes the baseline firing rate,  $a_1$  (allowed to range between  $10^{-7}$  and 0.5) describes the amplitude of the temporal field,  $\mu$  (allowed to range between 0 and the average midpoint across reproduced intervals for a given cell) describes the temporal receptive field peak location of the cell,  $\sigma$  (allowed to range between 0.01 and 0.5)



describes the variability of the peak,  $\tau$  (allowed to range between 0.1 and 16) describes the time constant of the exponential decay in firing, and  $\text{erfc}$  is the complementary error function.

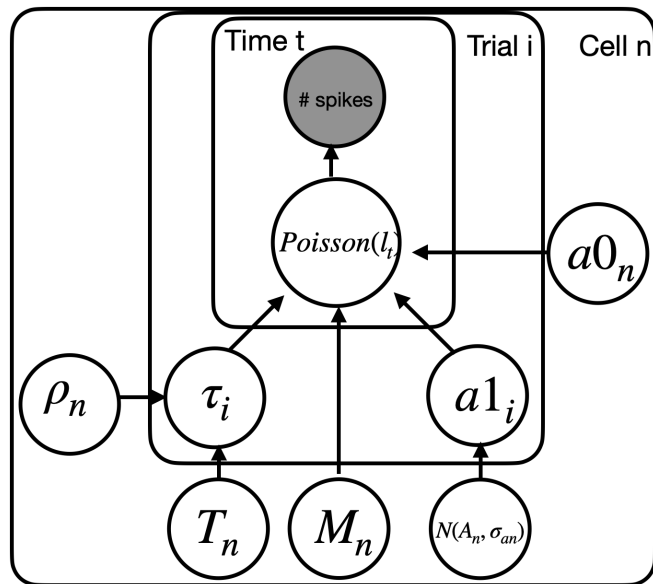
In addition, we also considered an exponentially modified Gaussian model that modulated the exponential time constant by the length of the reproduced interval. This reproduction length modulated model replaced  $\tau$  with  $\tau_i$ , such that  $\tau_i = \tau \frac{T_i}{5.25}$ , where  $\tau_i$  linearly rescale with reproduced interval length  $T_i$ . Parameters fit by the model spanned the same ranges as in the prior exponentially modified Gaussian model.

We also fit the exponentially modified Gaussian model for time until the end of the reproduction phase, replacing  $t$  with  $T - t$ , with models including both with  $\tau$  and  $\tau_i$ . All parameters fit by the maximum likelihood models for the end of the delay spanned the same ranges as the parameters in prior models.

We evaluated the models for each neuron using a likelihood ratio test and counted the number of neurons that were better fit by an ex-Gaussian model at the 0.05 level, Bonferroni-corrected by the total number of neurons. We further required that both even and odd trials for a neuron were significantly fit by the same model, and that those fits had a Pearson's correlation coefficient greater than 0.4. In the event that multiple ex-Gaussian models met these requirements for a given cell, we selected the model with the best log-likelihood. Therefore the temporal responsive cells are further categorized into cells best fit with an ex-Gaussian model with time elapsed ( $t$ , the "past cells"), and cells best fit with an ex-Gaussian model with time remaining ( $T - t$ , "future cells").

**Hierarchical Bayesian models describe temporal receptive field across trials.** We fit hierarchical Bayesian models inspired by the results from the maximum likelihood methods, with additional parameters estimated at the trial level so cells that code for absolute time and cells that rescale with different reproduced intervals can be captured with the same model. We fit three hierarchical Bayesian models: full ex-Gaussian, exponential, and Gaussian. Because ex-Gaussian contains both exponential and Gaussian models, we first describe the full ex-Gaussian model and detail the other two models after. A graphic diagram of the Bayesian model we reported in the main section can be found in Fig. 5

For cells categorized as past cells by the maximum likelihood methods, we assume a Poisson process binned at 10 ms is applied to evaluate the latent model  $M_{\text{EMG}}$ . The model is largely the same as Eq. 2 except that both time constant  $\tau$  parameter and amplitude of time field  $a_1$  are estimated for each trial  $i$ :



**Fig. 5.** Illustration of the hierarchical Bayesian model. Each node represents a variable in the model; the filled node represents the observed number of spikes in each time bin (10 ms) and open nodes represent latent variables. Arrows represent relationships between variables and plates indicate whether the variable is estimated at the trial level, or cell level.

$$M_{\text{EMG}}(t; a_0, a_1, \sigma, \mu, \tau_i) = a_0 + \frac{a_1}{2} e^{\frac{(2\mu + \frac{\sigma^2}{\tau_i} - 2t)}{2\tau_i}} \text{erfc}\left(\frac{\mu + \frac{\sigma^2}{\tau_i} - t}{\sqrt{2}\sigma}\right) \quad [3]$$

The trial level time constant  $\tau_i$  is modified according to the time constant  $\tau$  (allowed to range between 0.1 to 10) at 5.25 s with an additional free parameter  $\rho$  and the interval length  $T_i$ :

$$\tau_i = \tau \left(1 + \rho \frac{T_i - 5.25}{5.25}\right), \quad [4]$$

where  $\rho$  quantifies how much and in which directions time constant  $\tau_i$  changes relative to the change of the produced interval length from the baseline estimation  $\tau$  when the interval is 5.25 s. A positive  $\rho$  indicates the cell increases its time constant for longer intervals while a negative  $\rho$  indicates the opposite. It is obvious that a  $\rho$  close to zero means the cell code for absolute time and does not change its time constant for different intervals. Another value of interest is when  $\rho = 1$ . In this case the time constant  $\tau_i$  linearly rescales with the interval: For example, if the interval is doubled ( $T_i = 10.5$  s), the corresponding  $\tau_i$  would also be twice the  $\tau$ . When the absolute value of  $\rho$  is smaller than 1, it means  $\tau_i$  changes relatively less compared to the change in the interval length.  $\rho$  are allowed to take any value as far as the resulted  $\tau_i$  from Eq. 4 is positive for the temporal receptive field estimation.

Another trial-level parameter  $a_{1i}$  is allowed to vary freely across trials. The only constraint was applied indirectly through marginal probability with the assumption that  $a_{1i}$  is sampled from a normal distribution for a given cell. The Gaussian width  $\sigma$  is allowed to range between 0 and 3. For cells categorized as future cells, we applied the same model as Eq. 3 except replace time elapsed  $t$  with time remaining  $T_i - t$  for each trial. For the past cells the peak time  $\mu$  is allowed to range between 0 to 3 s from the start of each trial; for the future cells the peak time  $\mu$  is allowed to range between 0 to 3 s to the end of each trial.

**Exponential and Gaussian models.** We constrained certain parameters in the full ex-Gaussian model to test alternative assumptions about receptive fields. For exponential fields, we fixed the Gaussian  $\sigma$  at 0.1; everything else remains the same from the model described above. For Gaussian fields, we fitted a Gaussian curve, which is equivalent to fixing the time constant  $\tau$  at 0. Just like the model described above, the Gaussian  $\mu$  is calculated from the beginning of interval for past cells and from the end of interval for future cells. To add additional flexibility to the Gaussian model, enabling it to fit individual trial data similarly to the full ex-Gaussian model, we assume that  $\mu$  can rescale with interval length in the same way as time constant parameter  $\tau$  in the ex-Gaussian model:  $\mu_i = \mu \left(1 + \rho \frac{T_i - 5.25}{5.25}\right)$ .  $\rho$  here can be interpreted the same way as the rescale parameter in Eq. 4. We also considered a linear ramp model with a threshold. However, because exponential functions  $e^x$  are well approximated by linear functions  $1 + x$  for small  $x$ , it is in practice difficult to distinguish exponential functions from linear functions with the amount of trial variability in spike train data, especially for this task.

The posterior distributions of estimated parameters were generated through the rStan package (57) with 4 independent Markov-Chain-Monte-Carlo (MCMC) chains (5,200 warm-up iterations and 50 post warm-up samples in each of the MCMC chains). This procedure returns posteriors of the likelihood of the data as well as posteriors for the parameters of the model.

**Evaluating the time constants population with Bayesian models.** We compared three models (power-law distribution, single normal distribution, and a mixture of two normal distributions) to describe the time constants population for the past cells the future cells separately, after the hierarchical Bayesian models described above were applied to quantify the temporal receptive field for each cell. The distribution parameters were estimated through rStan package with eight independent MCMC chains (4,000 warm-up iterations and 4,000 post warm-up samples in each of the MCMC chains).

**Power-law distribution.** The probability density function of Power law distribution can be derived within the range  $[\text{min}, \text{max}]$  of time constants  $\tau$ :

$$p(\tau) = \tau^{-\alpha} / C, \quad [5]$$

where

$$C = \int_{\min}^{\max} \tau^{-\alpha} d\tau \quad [6]$$

is the area under the density function over  $[\min, \max]$ .

When  $\alpha \neq 1$ , the result of the integration is

$$C = \frac{\max^{1-\alpha} - \min^{1-\alpha}}{1-\alpha} \quad [7]$$

When  $\alpha = 1$ , the result is

$$C = \ln(\max) - \ln(\min) \quad [8]$$

When fitting the Power-law model, we adopt a straightforward approach by setting the upper and lower bound of the power-law to be the same as the parameter space of  $\tau$ .

**Normal distribution.** In the second model, we quantify the time constants with a single normal distribution:

1. M. Jazayeri, M. N. Shadlen, A neural mechanism for sensing and reproducing a time interval. *Curr. Biol.* **25**, 2599–2609 (2015).
2. M. D. Mauk, D. V. Buonomano, The neural basis of temporal processing. *Annu. Rev. Neurosci.* **27**, 307–340 (2004).
3. H. Merchant, W. Zarco, O. Pérez, L. Prado, R. Bartolo, Measuring time with different neural chronometers during a synchronization-continuation task. *Proc. Natl. Acad. Sci. U.S.A.* **108**, 19784–19789 (2011).
4. N. S. Narayanan, Ramping activity is a cortical mechanism of temporal control of action. *Curr. Opin. Behav. Sci.* **8**, 226–230 (2016).
5. J. Gibbon *et al.*, Scalar timing in memory. *Ann. N. Y. Acad. Sci.* **423**, 52–77 (1984).
6. J. Henke *et al.*, Distributed coding of duration in rodent prefrontal cortex during time reproduction. *eLife* **10**, e71612 (2021).
7. J. Kim, J. W. Ghim, J. H. Lee, M. W. Jung, Neural correlates of interval timing in rodent prefrontal cortex. *J. Neurosci.* **33**, 13834–13847 (2013).
8. J. Kunimatsu, T. W. Suzuki, S. Ohmae, M. Tanaka, Different contributions of preparatory activity in the basal ganglia and cerebellum for self-timing. *eLife* **7**, e35676 (2018).
9. A. Mita, H. Mushiaki, K. Shima, Y. Matsuzaka, J. Tanji, Interval time coding by neurons in the presupplementary and supplementary motor areas. *Nat. Neurosci.* **12**, 502 (2009).
10. M. Xu, S. Zhang, Y. Dan, M. Poo, Representation of interval timing by temporally scalable firing patterns in rat prefrontal cortex. *Proc. Natl. Acad. Sci. U.S.A.* **111**, 480–485 (2014).
11. H. Merchant, B. B. Averbeck, The computational and neural basis of rhythmic timing in medial premotor cortex. *J. Neurosci.* **37**, 4552–4564 (2017).
12. P. Simen, F. Balci, L. deSouza, J. D. Cohen, P. Holmes, A model of interval timing by neural integration. *J. Neurosci.* **31**, 9238–9253 (2011).
13. K. I. Bakhurin *et al.*, Differential encoding of time by prefrontal and striatal network dynamics. *J. Neurosci.* **37**, 854–870 (2017).
14. S. Zhou, S. C. Masmanidis, D. V. Buonomano, Neural sequences as an optimal dynamical regime for the readout of time. *Neuron* **108**, 651–658 (2020).
15. E. D. Remington, D. Narain, E. A. Hosseini, M. Jazayeri, Flexible sensorimotor computations through rapid reconfiguration of cortical dynamics. *Neuron* **98**, 1005–1019 (2018).
16. R. Laje, D. V. Buonomano, Robust timing and motor patterns by taming chaos in recurrent neural networks. *Nat. Neurosci.* **16**, 925–933 (2013).
17. E. D. Remington, S. W. Egger, D. Narain, J. Wang, M. Jazayeri, A dynamical systems perspective on flexible motor timing. *Trends Cogn. Sci.* **22**, 938–952 (2018).
18. C. J. Cueva *et al.*, Low-dimensional dynamics for working memory and time encoding. *Proc. Natl. Acad. Sci. U.S.A.* **117**, 23021–23032 (2020).
19. S. Zhou, S. C. Masmanidis, D. V. Buonomano, Encoding time in neural dynamic regimes with distinct computational tradeoffs. *PLoS Comput. Biol.* **18**, e1009271 (2022).
20. M. Beiran, N. Meirhaeghe, H. Sohn, M. Jazayeri, S. Ostojic, Parametric control of flexible timing through low-dimensional neural manifolds. *Neuron* **111**, 739–753 (2023).
21. C. J. MacDonald, K. Q. Lepage, U. T. Eden, H. Eichenbaum, Hippocampal “time cells” bridge the gap in memory for discontinuous events. *Neuron* **71**, 737–749 (2011).
22. E. Pastalkova, V. Itskov, A. Amarasingham, G. Buzsáki, Internally generated cell assembly sequences in the rat hippocampus. *Science* **321**, 1322–1327 (2008).
23. N. A. Cruzado, Z. Tiganj, S. L. Brincat, E. K. Miller, M. W. Howard, Conjunctive representation of what and when in monkey hippocampus and lateral prefrontal cortex during an associative memory task. *Hippocampus* **30**, 1332–1346 (2020).
24. W. Mau *et al.*, The same hippocampal CA1 population simultaneously codes temporal information over multiple timescales. *Curr. Biol.* **28**, 1499–1508 (2018).
25. B. J. Kraus *et al.*, During running in place, grid cells integrate elapsed time and distance run. *Neuron* **88**, 578–589 (2015).
26. W. Ning, J. H. Bladon, M. E. Hasselmo, Complementary representations of time in the prefrontal cortex and hippocampus. *Hippocampus* **32**, 577–596 (2022).
27. Z. Tiganj, J. A. Cromer, J. E. Roy, E. K. Miller, M. W. Howard, Compressed timeline of recent experience in monkey IPFC. *J. Cogn. Neurosci.* **30**, 935–950 (2018).
28. H. Akhlaghpour *et al.*, Dissociated sequential activity and stimulus encoding in the dorsomedial striatum during spatial working memory. *eLife* **5**, e19507 (2016).
29. G. B. Mello, S. Soares, J. J. Paton, A scalable population code for time in the striatum. *Curr. Biol.* **25**, 1113–1122 (2015).
30. B. J. Kraus, R. J. Robinson II, J. A. White, H. Eichenbaum, M. E. Hasselmo, Hippocampal “time cells”: Time versus path integration. *Neuron* **78**, 1090–1101 (2013).

$$p(\tau) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{1}{2}\left(\frac{\tau - \mu}{\sigma}\right)^2\right), \quad [9]$$

where  $\mu$  is the estimated stereotypical time constant for the population.

**Mixture of two normal distributions.** The third model we tested is a mixture model of two normal distributions with the probability of  $\tau$  sampled from the first normal distribution to be  $\lambda_1$  and the probability of the second normal distribution to be  $1 - \lambda_1$ :

$$p(\tau) = \lambda_1 \text{normal}(\tau|\mu_1, \sigma_1) + (1 - \lambda_1) \text{normal}(\tau|\mu_2, \sigma_2) \quad [10]$$

**Data, Materials, and Software Availability.** Analysis code data have been deposited in Open Science Framework (58).

**ACKNOWLEDGMENTS.** We gratefully acknowledge support from R01MH132171532 and a Google-AI Faculty Research Award.

31. R. Cao, J. H. Bladon, S. J. Charczyski, M. E. Hasselmo, M. W. Howard, Internally generated time in the rodent hippocampus is logarithmically compressed. *eLife* **11**, e75353 (2022).
32. I. M. Bright *et al.*, A temporal record of the past with a spectrum of time constants in the monkey entorhinal cortex. *Proc. Natl. Acad. Sci. U.S.A.* **117**, 20274–20283 (2020).
33. A. Tsao *et al.*, Integrating time from experience in the lateral entorhinal cortex. *Nature* **561**, 57–62 (2018).
34. S. Zuo *et al.*, Neural signatures for temporal-order memory in the medial posterior parietal cortex. *bioRxiv* [Preprint] (2023). <https://doi.org/10.1101/2023.08.17.553665> (Accessed 24 August 2024).
35. K. H. Shankar, M. W. Howard, Optimally fuzzy temporal memory. *J. Mach. Learn. Res.* **14**, 3753–3780 (2013).
36. S. Watanabe, M. Opper, Asymptotic equivalence of Bayes cross validation and widely applicable information criterion in singular learning theory. *J. Mach. Learn. Res.* **11**, 3571–3594 (2010).
37. M. W. Howard, Z. G. Esfahani, B. Le, P. B. Sederberg, Learning temporal relationships between symbols with Laplace Neural Manifolds. *arXiv* [Preprint] (2024). <http://arxiv.org/abs/2302.10163> (Accessed 26 August 2024).
38. R. O. Affan *et al.*, Ramping dynamics in the frontal cortex unfold over multiple timescales during motor planning. *bioRxiv* [Preprint] (2024). <https://doi.org/10.1101/2024.02.05.578819> (Accessed 24 August 2024).
39. W. Schultz, P. Dayan, P. R. Montague, A neural substrate of prediction and reward. *Science* **275**, 1593–1599 (1997).
40. M. Sousa *et al.*, Dopamine neurons encode a multidimensional probabilistic map of future reward. *bioRxiv* [Preprint] (2023). <https://doi.org/10.1101/2023.11.12.566727> (Accessed 24 August 2024).
41. P. Masset *et al.*, Multi-timescale reinforcement learning in the brain. *bioRxiv* [Preprint] (2023). <https://doi.org/10.1101/2023.11.12.566754> (Accessed 24 August 2024).
42. I. Momennejad, M. W. Howard, Predicting the future with multi-scale successor representations. *bioRxiv* [Preprint] (2018). <https://doi.org/10.1101/449470> (Accessed 24 August 2024).
43. P. Tano, P. Dayan, A. Pouget, A local temporal difference code for distributional reinforcement learning. *Adv. Neural Inf. Proces. Syst.* **33**, 13662–13673 (2020).
44. M. W. Howard, M. E. Hasselmo, Cognitive computation using neural representations of time and space in the laplace domain. *arXiv* [Preprint] (2020). <http://arxiv.org/abs/2003.11668> (Accessed 24 August 2024).
45. Y. Liu, M. W. Howard, Generation of scale-invariant sequential activity in linear recurrent networks. *Neural Comput.* **32**, 1379–1407 (2020).
46. Z. Tiganj, M. E. Hasselmo, M. W. Howard, A simple biophysically plausible model for long time constants in single neurons. *Hippocampus* **25**, 27–37 (2015).
47. Y. Liu, Z. Tiganj, M. E. Hasselmo, M. W. Howard, A neural microcircuit model for a scalable scale-invariant representation of time. *Hippocampus* **29**, 260–274 (2019).
48. M. Stern, N. Istrate, L. Mazzucato, A reservoir of timescales emerges in recurrent circuits with heterogeneous neural assemblies. *eLife* **12**, e86552 (2023).
49. H. Kim, J. Homann, D. W. Tank, M. J. Berry, A long timescale stimulus history effect in the primary visual cortex. *bioRxiv* [Preprint] (2019). <https://doi.org/10.1101/585539> (Accessed 24 August 2024).
50. B. C. Daniels, M. W. Howard, Continuous attractor networks for laplace neural manifolds. *arXiv* [Preprint] (2024). <http://arxiv.org/abs/2406.04545> (Accessed 26 August 2024).
51. M. J. Wagner, T. H. Kim, J. Savall, M. J. Schnitzer, L. Luo, Cerebellar granule cells encode the expectation of reward. *Nature* **544**, 96–100 (2017).
52. J. L. Gauthier, D. W. Tank, A dedicated population for reward coding in the hippocampus. *Neuron* **99**, 179–193.e7 (2018).
53. A. Sarel, A. Finkelstein, L. Las, N. Ulanovsky, Vectorial representation of spatial goals in the hippocampus of bats. *Science* **355**, 176–180 (2017).
54. K. H. Shankar, M. W. Howard, Timing using temporal context. *Brain Res.* **1365**, 3–17 (2010).
55. M. W. Howard *et al.*, A unified mathematical framework for coding time, space, and sequences in the hippocampal region. *J. Neurosci.* **34**, 4692–4707 (2014).
56. E. T. Rolls, P. Mills, The generation of time in the hippocampal memory system. *Cell Rep.* **28**, 1649–1658 (2019).
57. Stan Development Team, RStan: The R Interface to Stan (R package version 2.32.6, The Comprehensive R Archive Network, 2024).
58. R. Cao, Past and future Laplace. Open Science Framework. <https://osf.io/rkxps/>. Deposited 26 August 2024.