

A temporal record of the past with a spectrum of time constants in the monkey entorhinal cortex

Ian M. Bright^{1,*}, Miriam L. R. Meister^{2,3,4,*}, Nathanael A. Cruzado¹, Zoran Tiganj¹, Marc W. Howard^{1,*}, and Elizabeth A. Buffalo^{2,3,4,*}

1. Department of Psychological and Brain Sciences, Boston University

2. Department of Physiology and Biophysics, University of Washington

3. Washington National Primate Research Center

4. University of Washington School of Medicine

* The order of these authors was determined by a coin flip.

Abstract

Episodic memory is believed to be intimately related to our experience of the passage of time. Indeed, neurons in the hippocampus and other brain regions critical to episodic memory code for the passage of time at a range of time scales. The origin of this temporal signal, however, remains unclear. Here, we examined temporal responses in the entorhinal cortex of macaque monkeys as they viewed complex images. Many neurons in the entorhinal cortex were responsive to image onset, showing large deviations from baseline firing shortly after image onset but relaxing back to baseline at different rates. This range of relaxation rates allowed for the time since image onset to be decoded on the scale of seconds. Further, the ensemble carried information about image content suggesting that neurons in the entorhinal cortex carry information not only about when an event took place but also the identity of that event. Taken together, these findings suggest that the primate entorhinal cortex uses a spectrum of time constants to construct a temporal record of the past in support of episodic memory.

Episodic memory, the vivid recollection of an event situated in a specific time and place (Tulving, 1983), depends critically on medial temporal lobe (MTL) structures, including the hippocampus and entorhinal cortex (EC) (Milner, 1959; Eichenbaum, Yonelinas, & Ranganath, 2007; Dedek, Frascino, Wixted, & Squire, 2016; Squire, Stark, & Clark, 2004). Building on pioneering work demonstrating a spatial code in the hippocampus and entorhinal cortex (O'Keefe & Dostrovsky, 1971; Fyhn, Molden, Witter, Moser, & Moser, 2004),

This work was supported by the National Institutes of Health, Grants 2R01MH080007 and R01MH093807 and National Institute of Mental Health Grant P51 OD010425 to E.A.B. and Grants NIBIB R01EB022864 and NIMH R01MH112169 to M.W.H., and the Office of Naval Research, MURI N00014-16-1-2832.

recent research has shown that hippocampal representations also carry information about the time at which past events took place (Pastalkova, Itskov, Amarasingham, & Buzsaki, 2008; MacDonald, Lepage, Eden, & Eichenbaum, 2011, see Eichenbaum, 2017 for a recent review), suggesting that the MTL maintains a representation of spatiotemporal context in support of episodic memory. Although a great deal is known about the temporal coding properties of neurons in the hippocampus, the temporal code in the entorhinal cortex, which provides the majority of the cortical input to the hippocampus is less understood (but see Naya & Suzuki, 2011; Kraus et al., 2015; Tsao et al., 2018).

Hippocampal *time cells* provide a record of recent events including explicit information about when an event occurred. Analogous to hippocampal place cells that fire when an animal is in a circumscribed region of physical space (e.g., Wilson & McNaughton, 1993; O’Keefe & Dostrovsky, 1971), hippocampal time cells fire during a circumscribed period of time within unfilled delays (MacDonald et al., 2011; Kraus, Robinson, White, Eichenbaum, & Hasselmo, 2013). Across studies, there is a remarkable consistency in the properties of hippocampal time cells. Hippocampal time cells peak at a range of times during the delay interval and typically code time with decreasing accuracy as the delay unfolds, as manifest by fewer neurons with peak firing late in the delay and wider time fields later in the delay (e.g., Kraus et al., 2015; Salz et al., 2016; Mau et al., 2018). Hippocampal time cells have been observed in a wide range of tasks, including tasks with and without explicit memory demands during the delay (e.g., Salz et al., 2016), experiments where the animal is fixed in space (MacDonald, Carrow, Place, & Eichenbaum, 2013; Terada, Sakurai, Nakahara, & Fujisawa, 2017), and different stimuli trigger different time cell sequences (Terada et al., 2017; MacDonald et al., 2013). Taken together, time cells provide an explicit record of how far in the past an event, e.g., for instance the beginning of a delay period or a to-be-remembered stimulus, took place. By examining which time cells are active at a particular time, we can easily determine not only what event took place, but how far in the past that event is.

Many of the properties of hippocampal time cells have been observed in other brain regions including prefrontal cortex (Bolkan et al., 2017; Tiganj, Kim, Jung, & Howard, 2017; Tiganj, Cromer, Roy, Miller, & Howard, 2018; Jin, Fujii, & Graybiel, 2009) and striatum (Jin et al., 2009; Mello, Soares, & Paton, 2015; Akhlaghpour et al., 2016) suggesting that the hippocampus is part of a widespread network that carries information about what happened when in the past. A recent report from the rat lateral EC adds important data to this growing body of literature regarding the representation of time in the brain. Tsao et al. (2018) observed a population of neurons that changed slowly and reliably enough to decode time within the experiment over a range of time scales. Unlike time cells, which respond a characteristic time since the event that triggers their firing, lateral EC neurons responded immediately upon entry into a new environment, and then relaxed slowly. The relaxation times across individual neurons were very different, ranging from tens of seconds up to thousands of seconds. To distinguish this population from time cells we will refer to neurons that are activated by an event and then relax their firing gradually as *temporal context cells*. Because these temporal context cells code for time, but with very different properties than time cells, they provide a potentially important clue about the nature of temporal coding in the brain, and thus memory function as well.

Here, we identified temporal context cells in monkey EC (Meister & Buffalo, 2018). We examined EC neuron responses in a five-second period after presentation of an image.

In the time after presentation of the image, a representation of what happened when should carry both time-varying information about when the image was presented, as well as information that also discriminates the identity of the image. To anticipate results, the data demonstrate that neurons in monkey EC are activated shortly after a visual stimulus and then decay with a variety of rates, enabling reconstruction of when the image was presented. This form of temporal coding is similar to temporal context cells observed in rat lateral EC (Tsao et al., 2018), but unlike time cells that have been observed in the hippocampus and other regions. Because different images were each shown twice over the course of the experiment, we were able to assess whether the pattern of activation over neurons depends on the content of the image presented. The results of this series of analysis suggest that these temporal context cells also separately carry information about what happened in addition to when it happened.

Results

A total of 349 neurons were recorded from the entorhinal cortex (EC) in two macaque monkeys during performance of a visual free-viewing task. Each trial began with a required fixation on a small cross, followed by presentation of a large, complex image that remained on the screen for five seconds of free viewing (Figure 1a). Unlike canonical hippocampal time cells, which are activated at a variety of points within a time interval (e.g., Figure 1c), most entorhinal neurons changed their firing shortly after the presentation of the image. Figure 2a shows three representative neurons that responded to the image presentation (more examples in Supplementary Figure S2). While most of the responsive neurons increased their firing rate after the image was presented, some neurons decreased their firing rate in response to image presentation. Although behavior was not controlled during the five second free-viewing period, the response of these neurons was consistent across trials (this can be seen by examination of the trial rasters).

Although the image-responsive neurons in EC responded at about the same time post-stimulus, they relaxed back to their baseline firing at different rates. Whereas some neurons relaxed back to baseline quickly (Figure 2a, top), some relaxed much more slowly. For instance the neuron shown in the bottom of Figure 2a did not return to baseline even after five seconds.

Responses begin at a similar time point and decay at different rates across neurons

Each neuron’s temporal response was quantified relative to the onset of the image using a model-based approach in which the shape of the firing field was described with parameters that corresponded to the latency and decay of the response for each neuron (Figure S1) and were assessed for goodness-of-fit using maximum likelihood. This approach builds on previous work using maximum likelihood methods to estimate time cell activity as a Gaussian firing field (Tiganj, Kim, et al., 2017; Salz et al., 2016; Tiganj et al., 2018). The method for estimating parameters is described in detail in the methods. We can conclude that a neuron’s response was time-locked to the image presentation to the extent a model with a temporal response field fit the neuron’s data better than a model with only constant background firing.

In order to minimize the noise and get the most accurate distribution of response parameters across neurons, we used a conservative criterion to identify neurons that responded

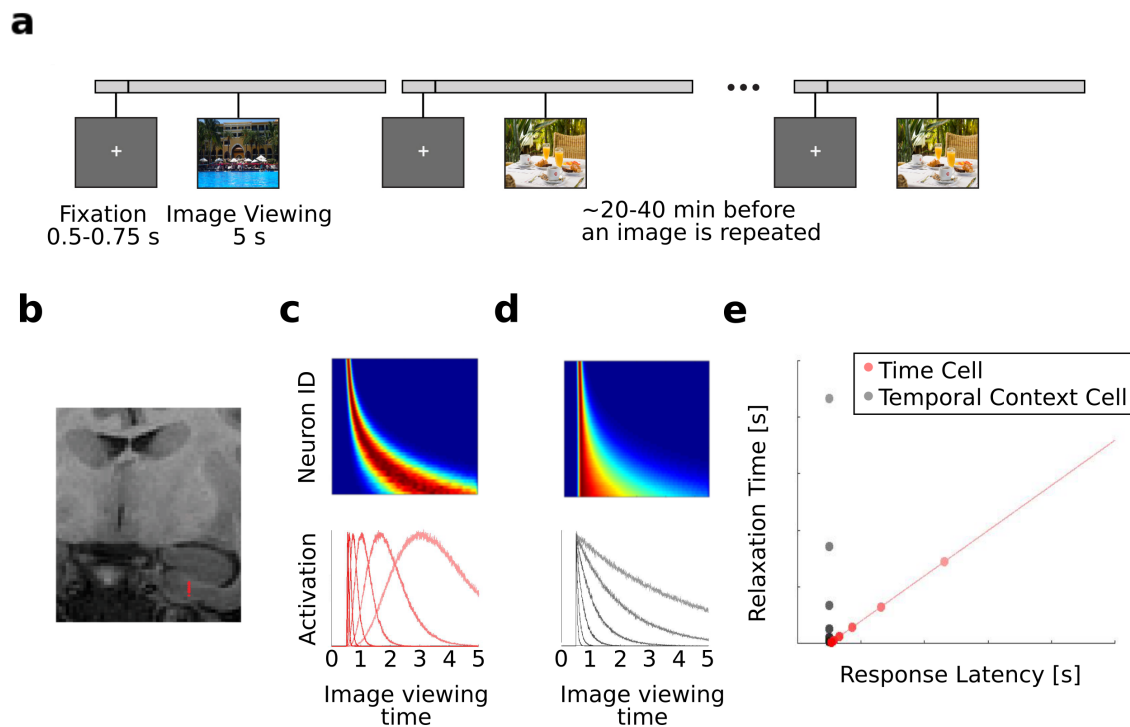


Figure 1. : a-b: Summary of experimental procedures. **a**, Trial schematic for three trials. On each trial, the monkey freely viewed an image. After the monkey viewed an image for 5 seconds, the image disappeared. Between image presentations (not shown), the monkey received a fruit slurry reward for performing gaze calibration trials (see Methods for details). Images were presented twice during an experimental session. Between 20-40 minutes passed before an image was repeated. **b**, Estimated position of recording channels in the entorhinal cortex in one recording session is shown in red on a coronal MRI. **c-e: Two hypotheses for neural representations of time.** **c-d**, Heat plot (top) and tuning curves (bottom) for two hypotheses for how a time interval of image viewing could be coded in neural populations. In the heat plots, cooler colors correspond to low activity while warmer colors correspond to higher activity. **c**, Hypothetical activity is shown for sequentially activated time cells like those observed in the hippocampus. In this population, different neurons exhibit peak responses at different times indicating different firing fields. Because the time of peak response across neurons covaries with the spread of the firing field, neurons with later firing fields display wider firing fields. **d**, Monotonically decaying temporal context cells. Neurons in this population exhibit peak firing at about the same time. However, different neurons decay at different rates. **e**, Properties of neurons representing time passage by the hypothesis shown in **c** (red) or the hypothesis shown in **d** (gray). A population of time cells (red) should exhibit peak responses that occur at different times across a time interval, and these neurons should show a robust correlation between when peak response occurs and the time it takes the response to relax back to baseline. Conversely, a population of exponentially-decaying temporal context cells (gray) should exhibit peak responses that occur in a more restricted time range shortly after the start of a time interval, and these neurons should show no correlation between when peak response occurs and the time it takes the response to relax back to baseline.

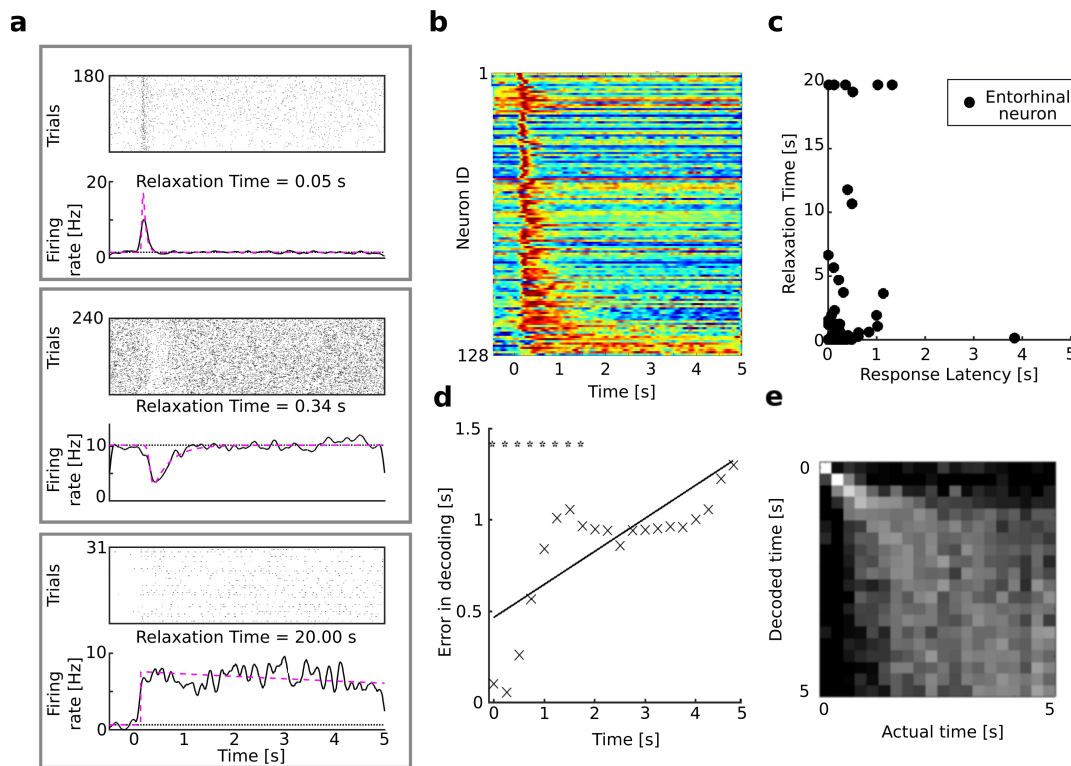


Figure 2. : Temporal context cells in entorhinal cortex respond to the presentation of an image, then relax with a variety of rates, carrying information about time since the image was presented. **a**, Three representative entorhinal neurons that respond to image onset and decay at different rates. (Figure S2 shows more examples). Each rectangle indicates the activity of one neuron relative to image onset in a raster plot (top) and PSTH (bottom). In the raster, a tick mark indicates when the neuron fired an action potential. In the PSTH, the solid black line indicates the firing rate, and the pink line indicates the model estimate of firing rate. Baseline firing rate is indicated by the black dotted line. **b**, The majority of neurons responded shortly after image presentation, relaxing back to baseline with a spectrum of decay rates; some neurons relaxed back to baseline much sooner than other neurons that relaxed more slowly. The heat plot shows normalized firing rates of all entorhinal neurons relative to image onset, sorted by Relaxation Time (duration between response peak and when response returns 63% of the way to baseline). Color scheme is the same as the heat plots in Figure 1c, d. **c**, Joint distribution across neurons shows that the time point the neurons started responding (Response Latency) were clustered within a few hundred ms, whereas Relaxation Times spanned the full 5 s of image presentation. Across neurons, Response Latency was not correlated with Relaxation Time. **d** and **e**, Time since presentation of the image could be decoded from the population of entorhinal neurons. A linear discriminant analysis (LDA) decoder was trained to decode time and then tested on excluded trials. **d**, Decoding error increases with the passage of time, indicated by the increasing values of absolute error between decoded time and actual time as a function of actual time since image onset. The heavy black line is a fitted regression line. The asterisks mark time bins for which a conservative estimate showed above-chance decoding accuracy at that time bin taken in isolation (see Methods and Results for details). **e**, Decoding precision decreases with passage of time, as shown by the wider spread of probability across decoded time estimates as more time passes. The x-axis indicates the actual time bin which the LDA attempted to decode, the y-axis indicates the decoded time, and the shading indicates the log of the posterior probability, with darker shading signaling lower probabilities. Decoding is better than chance and decreases in precision with the passage of time.

to image presentation. This method identified 128/349 neurons as responsive. Of those 128 responsive neurons, 100 neurons showed an increase in their firing rate in response to image onset, whereas 28 showed a decrease in their firing rate.

Figure 2b summarizes the temporal response properties of these 128 neurons. Each row of the figure shows the response of one neuron over the course of a trial. The data demonstrate that almost all of the neurons reached their maximum deviation from baseline within a few hundred milliseconds of the image presentation. This can be appreciated from the vertical red strip along the left edge of the heat map. These results are in striking contrast to analogous plots of hippocampal time cells, a neural population in which different neurons fire in sequence, tiling a time interval (e.g., Figure 1c) and resulting in a curved ridge extending from the upper left to the lower right. In contrast, the variability across neurons in this entorhinal population is not in the time point at which the neurons reach their maximum deviation from baseline, but rather in the time course over which each neuron relaxes. This can be appreciated in the progressive widening of the ridge in Figure 2b from top to bottom.

We modeled temporal responses as a convolution of a Gaussian (latency in neuronal response) and an exponential decay (Figure S1). In this model, we were able to quantify apparent response properties using two key parameters: 1) the parameter μ , which describes the mean of the Gaussian, estimates the time at which each neuron begins to respond (Response Latency) and 2) the parameter τ , which describes the time constant of the exponential term, estimates how long each neuron takes to relax back to 63% of its maximum deviation (Relaxation Time) (Figure 1e). We refer to these two parameters as response latency and relaxation time respectively. A third parameter σ controls the standard deviation of the Gaussian term.

Figure 2c shows the Response Latency and Relaxation Time for each neuron that was categorized as visually responsive. Response Latency values were clustered tightly at small values (median = 0.16 s, interquartile range = 0.13 s to 0.26 s), and did not span the entirety of the 5 s free viewing period. For 90% of neurons, the Response Latency Value was less than 0.46 s. In contrast, neuron Relaxation Time values showed a much wider distribution (median = 0.54 s, interquartile range = 0.22 s to 2.02 s, 90th percentile = 8.50 s), and even included values longer than the 5 s duration of the viewing period. Lastly, the third parameter, which controlled the standard deviation of the Gaussian, was small and tightly clustered across neurons (median = 0.02 s, interquartile range of 0.001 s to 0.06 s, 90th percentile = 0.27 s).

Across neurons, Response Latency and Relaxation Time were not significantly correlated with one another, Kendall's $\tau = 0.05$, $p = 0.40$. To assess whether this null effect was reliable, we computed the Bayes factor, a measure of the likelihood of the null hypothesis. This analysis yielded a Bayes factor of $BF_{01} = 6.09$, providing support that neuron Response Latency and Relaxation Time values are uncorrelated. Across neurons Response Latency and σ were not correlated with one another, Kendall's $\tau = -0.075$, $p = 0.22$, $BF_{01} = 3.97$. Unlike hippocampal time cells, there was no evidence that temporal context cells that peaked later in the time interval showed broader firing fields. In contrast to hippocampal time cells, which show a systematic relationship between the peak time of firing and the width of the temporal firing field (Kraus et al., 2013; Howard et al., 2014; Salz et al., 2016), the overarching conclusion from these analyses is that the firing of entorhinal

neurons deviated from background firing shortly after the presentation of the stimulus and then relaxed exponentially with a variety of time constants, with the time constant of a neuron independent of its peak time.

Decoding time after image presentation from neuron responses

It is well-understood that sequentially activated time cells can be used to decode the time since the beginning of a delay (e.g., Mau et al., 2018, see also Supplementary Figure S3). To assess the temporal information present in the population of entorhinal neurons we trained a linear discriminant analysis (LDA) decoder to estimate time (in 0.25 s bins, for a total of 20 bins) following presentation of an image. To the extent the predicted time bin for out-of-sample data is close to the actual time bin, one can conclude that the population response carried information about time.

Time was decoded better than chance. Figure 2d-e shows the results of the LDA on the neurons from monkey EC. Our first question was whether or not the population contains information about time. For each time bin in Figure 2e, the confidence of the decoder (the posterior distribution) is shown across the range of possible time estimates. Perfect prediction would correspond to a bright diagonal; random uniform decoding would correspond to a gray square. Qualitatively, the non-uniformity of Figure 2e suggests that elapsed time can be decoded from the population of EC neurons. To quantitatively assess this, we found that the posterior distribution from the test data was reliably different from a uniform distribution using a chi-squared goodness of fit test, $\chi^2(380) = 522.17$, $p < 0.001$.

Supporting this result, the mean absolute value of decoding error from the cross-validated LDA was reliably lower than the decoding error from training with a permuted data set. In each of 1000 permutations we randomly reassigned the time bin labels of the training events used to train the classifier. The absolute value of the decoding error for the original data was 1.35 s, which is more accurate than the mean absolute value of the decoding error for all 1000 permutations. As shown in Figure S4, the values for the permuted data were approximately normal with mean 1.66 s and standard deviation 0.04 s, resulting in a z-score of more than seven ($z = 7.75$). These analyses show that time can be decoded from monkey EC.

The precision of the time estimate decreased as the interval unfolded. Although the population response in entorhinal cortex could be used to reconstruct time, inspection of Figure 2e suggests that the precision of this reconstruction was not constant throughout the interval. Figure 2d shows the average absolute value of the decoding error at each time bin. These data suggest that this error increased as a function of time. A linear regression of decoding error as a function of time showed a reliable slope, 0.18 ± 0.04 , as well as intercept 0.5 ± 0.1 , both $p < 0.001$, $R^2 = 0.56$, $df = 18$. Thus the temporal information in the population of entorhinal cortex neurons decreases as time elapses.

Time can be decoded well past the peak firing of temporal context cells. To assess how far into the interval time could be reconstructed, we repeated the LDA analysis excluding progressively more time bins starting from zero. If the LDA can reconstruct time using only bins corresponding to times $\geq t$, then we can conservatively conclude that time can be reconstructed at least time t into the interval. To assess this quantitatively, the actual data

were compared with permuted data for each repetition of the LDA using absolute error to assess performance (see Methods for details). This analysis showed that time can be reliably decoded even if the first 1.75 s were excluded. Thus the population of temporal context cells contains information about time at least 1.75 s after image onset. This conservative estimate is an order of magnitude longer than the median value of the peak time (0.160 s), suggesting that the gradual decay of temporal context cells could be used to reconstruct information about time.

Activity of individual EC neurons was similar for presentations of the same image

In this experiment, each image was presented twice. Although it is not practical to assess image coding using a classifier, it is possible to exploit the repetition of images to determine if EC neurons contained information about image identity. To assess this we assembled, for each neuron, an array giving the firing rate during the first presentation of each image (5 s) and asked whether this array was correlated with the firing rate of second presentations of the same images. If the firing rate of a neuron depends on the identity of the image, we would expect to see a positive correlation using this measure. Note that because this analysis compares the correlation between first and second presentations across stimuli it is not sensitive to repetition effects (e.g., Xiang & Brown, 1998; Meyer & Rust, 2018; Jutras & Buffalo, 2010). For this analysis we restricted our attention to neurons ($n = 270$) that were recorded long enough to be observed for both first and second presentations of a block of stimuli (repetitions were separated by 20-40 minutes).

The mean correlation coefficient (Kendall's τ) across neurons was significantly greater than zero, $\tau = 0.06 \pm 0.02$, $t(269) = 7.69$, $p < 0.001$, Cohen's $d = 0.47$, indicating that the spiking activity of many neurons depended on image identity (Fig. S5c). This comparison was confirmed by a Wilcoxon signed rank test on the values of Kendall's τ , $V = 27577$, $p < 0.001$. This finding was also observed for the subset of visually-responsive neurons ($N = 104$) that we describe as temporal context cells. Taken in isolation, the temporal context neurons showed a mean significantly greater than zero, as measured by t-test, 0.09 ± 0.02 , $t(103) = 7.70$, $p < 0.001$, Cohen's $d = 0.76$ and Wilcoxon signed rank test, $V = 4696$, $p < 0.001$. Neurons that were not temporal context cells ($n = 166$) also had a mean correlation coefficient different from zero 0.04 ± 0.02 , $t(165) = 4.18$, $p < 0.001$, Cohen's $d = 0.32$, $V = 9525$, $p < 0.001$. Taken together, this suggests that EC neuron spike activity contains information about stimulus identity.

The population of EC neurons distinguished image identity

The preceding analyses show that the firing of many EC neurons distinguished image identity above chance. If the response of the entire population contained information about stimulus identity, we would expect, all things equal, that pairs of population vectors corresponding to presentations of the same image would be more similar to one another than pairs of population vectors corresponding to presentations of different images. To control for any potential repetition effect, these analyses compared the similarity between the repetition of an image with its original presentation to the similarity between the second presentation of an image with the presentation of the images adjacent to original presentation. To be more concrete, if we label a sequence of images initially presented in sequence as A, B, and C, we would separately compare the population response to the repetition of

B to the response to the initial presentation of A, B, and C. To the extent that the response to the second presentation of B is more similar to the initial response to B (lag 0) than to the response to A (lag - 1) or C (lag + 1), we can conclude that the population vector distinguishes image identity.

Figure S5d shows the results of this population analysis. The similarity for lag 0 pairs (comparing an image with itself) was greater than the similarity for lags ± 1 (comparing a repeated image to the neighbors of its original presentation). Statistical comparisons to lags ± 1 each showed a reliable difference. A paired t-test comparing population similarity at the block level showed that similarity at lag 0 was reliably larger than both lag +1, 0.012 ± 0.005 , $t(63) = 5.11$, $p < 0.001$, Cohen's $d = 0.64$, and lag -1, 0.014 ± 0.006 , $t(63) = 5.02$, $p < 0.001$, Cohen's $d = 0.63$. To evaluate the same hypothesis using a non-parametric method, we performed a permutation analysis by randomly swapping within-session pairs of lag 0 and lag ± 1 and calculating the mean difference between the pairs 100,000 times. The observed value exceeded the value of 100,000/100,000 permuted values for both lags +1 and -1. We conclude that the population response was more similar for presentations of the same image than for presentations of different images. This analysis controlled for repetition and recency.

Discussion

Episodic memory requires information about both the content of an event as well as its temporal context (Tulving, 1983; Eichenbaum et al., 2007; Eichenbaum, 2017). In this study, many EC neurons changed their firing in response to the onset of the image. These *temporal context cells* changed their firing at about the same time, within a couple hundred milliseconds after the image was presented. However, different temporal context cells showed variable rates of relaxation back to baseline (Figure 2). Information about time since the image was presented could be decoded due to gradually decaying firing rates over a few seconds (Figure 2d-e). Notably, the relaxation rate was not constant across neurons, but rather showed a spectrum of time constants. The population vectors following repeated presentations of the same image were more similar to one another than to presentations of different images. This, coupled with several control analyses, show that the firing of entorhinal neurons also distinguished stimulus identity. Taken together the results demonstrate that in the time after image presentation, the population of EC neurons contained information about what happened when; temporal information could be distinguished due to gradually-decaying firing rates.

Sequentially activated time cells, such as have been observed in the hippocampus (Pastalkova et al., 2008; MacDonald et al., 2011; Salz et al., 2016), medial entorhinal cortex (Kraus et al., 2015) and many other brain regions (Jin et al., 2009; Mello et al., 2015; Tiganj et al., 2018; Tiganj, Shankar, & Howard, 2017) also contain temporal information about the past. Entorhinal temporal context cells provide a temporal record of recent events, but with different firing properties than canonical hippocampal time cells. As a population, time cells convey the passed time since the occurrence of an event by firing at different temporal delays after the triggering event. In contrast, EC temporal context cells here all responded at about the same time but relaxed at different rates. The range of relaxation times enables the population to convey temporal information at different time scales. For instance, a temporal context cell that decays back to baseline firing within 1 second cannot

be used to distinguish a 5 second delay from a 10 second delay. In contrast, a cell that only decays back to baseline after 7 seconds can be used to distinguish this longer delay. In this way, a range of decay rates enables the population of temporal context cells to decode time over a wide range of time scales.

The pattern of results observed here aligns well with a recent report from rodent lateral EC by Tsao et al. (2018). In the Tsao et al. study, lateral EC neurons responded to a salient event, i.e., the animal entering a new environment, and then changed their firing monotonically. Notably, Tsao et al. observed that different neurons responded at different rates from the scale of tens of seconds to many minutes. However, despite the many methodological differences between the Tsao et al. study and this one—rats moving through a series of open enclosures *vs* seated monkeys observing a series of images—the response properties shared striking similarities, suggesting a common computational function for EC across species.

Our results are also consistent with studies that have shown long-lasting responses in EC neurons in different preparations in rodents. For example, sustained responses have been observed in both slice (Egorov, Hamam, Fransén, Hasselmo, & Alonso, 2002) and anesthetized (Leitner et al., 2016) EC preparations. Similarly, slow changes in firing rate have been observed in EC during classical conditioning (Pilkiw et al., 2017). Coupled with studies of navigating rodents demonstrating that EC neurons can code body position relative to parts of the environment (e.g., Deshmukh & Knierim, 2011; Wang et al., 2018), these findings suggest that EC neurons code for stimuli in both temporal and spatial domains.

Exponentially-decaying neurons with a spectrum of time constants is the Laplace transform of time

This study in primates and the rodent study by Tsao et al. (2018) both observed entorhinal neurons that code time by gradually changing firing rate; different neurons changed at different rates. Time cells in the hippocampus and other brain regions also carry information about time, but with neurons with very different response properties. Why would the brain use two distinct coding schemes to represent the same kind of information? It has been proposed (Shankar & Howard, 2012, 2013; Howard et al., 2014) that the brain estimates a temporal record of the past—a function over passed time—by first computing the real Laplace transform of that function. The real Laplace transform is just a set of exponentially-decaying kernel functions with a spectrum of rate constants convolved with the input signal. In this study, if the input signal were the onset of the image, the real Laplace transform of the time in the past at which the image was presented would manifest as a set of exponentially-decaying cells with a spectrum of time constants (Howard et al., 2014). This is very much like the results observed here in entorhinal neurons (Figure 2). The slowly-relaxing cells observed in rats by Tsao et al. (2018) also resemble the Laplace transform of the past in response to a new spatial environment.

Although it may seem inefficient to estimate the Laplace transform of a function and then invert the transform rather than estimating the function directly, there are several computational advantages to this approach. For instance, in much the same way that the Fourier transform has useful computational properties that make it widely used in signal processing and manipulation, many computations can be more efficiently calculated in the Laplace domain (Howard, Shankar, & Tiganj, 2015). The inverse transform, which

takes a set of exponentially-decaying temporal context cells into a set of sequentially activated time cells, can be implemented using simple center-surround receptive fields in a feedforward circuit (Shankar & Howard, 2013; Liu, Tiganj, Hasselmo, & Howard, 2019). Notably, one can use the same computational framework to compute the Laplace transform of functions other than time, including spatial variables (Howard et al., 2014), accumulated evidence (Howard, Luzardo, & Tiganj, 2019) and temporal distance to expected outcomes (Momennejad & Howard, 2018). For instance, neurons that fire along an environmental border in rodent and monkey EC (Solstad, Boccara, Kropff, Moser, & Moser, 2008; Kilian, Jutras, & Buffalo, 2012; Hardcastle, Ganguli, & Giocomo, 2015), have properties like one would expect for the Laplace transform of distance to an environmental boundary. In the spatial case, the inverse transform of distance to a boundary would appear as one-dimensional place cells (Gothard, Skaggs, Moore, & McNaughton, 1996; Lever, Burton, Jeewajee, O’Keefe, & Burgess, 2009; Burgess & O’Keefe, 1996). Perhaps temporal context cells and time cells—the Laplace transform and inverse transform of functions of time—are a special case of a more general principle of computational cognitive neuroscience in which abstract continuous variables are coded using a continuous spectrum of rate constants in the Laplace domain.

References

- Akhlaghpour, H., Wiskerke, J., Choi, J. Y., Taliaferro, J. P., Au, J., & Witten, I. (2016). Dissociated sequential activity and stimulus encoding in the dorsomedial striatum during spatial working memory. *eLife*, 5, e19507.
- Bolkan, S. S., Stujenske, J. M., Parnaudeau, S., Spellman, T. J., Rauffenbart, C., Abbas, A. I., ... Kellendonk, C. (2017). Thalamic projections sustain prefrontal activity during working memory maintenance. *Nature Neuroscience*, 20(7), 987–996. Retrieved from <http://dx.doi.org/10.1038/nn.4568>
- Burgess, N., & O’Keefe, J. (1996). Neuronal computations underlying the firing of place cells and their role in navigation. *Hippocampus*, 6(6), 749–62.
- Dede, A. J., Frascino, J. C., Wixted, J. T., & Squire, L. R. (2016). Learning and remembering real-world events after medial temporal lobe damage. *Proceedings of the National Academy of Sciences*, 113(47), 13480–13485.
- Deshmukh, S. S., & Knierim, J. J. (2011). Representation of non-spatial and spatial information in the lateral entorhinal cortex. *Frontiers in behavioral neuroscience*, 5, 69.
- Egorov, A. V., Hamam, B. N., Fransén, E., Hasselmo, M. E., & Alonso, A. A. (2002). Graded persistent activity in entorhinal cortex neurons. *Nature*, 420(6912), 173–8.
- Eichenbaum, H. (2017). On the integration of space, time, and memory. *Neuron*, 95(5), 1007–1018. doi: 10.1016/j.neuron.2017.06.036
- Eichenbaum, H., Yonelinas, A., & Ranganath, C. (2007). The medial temporal lobe and recognition memory. *Annual Review of Neuroscience*, 30, 123–152.
- Fyhn, M., Molden, S., Witter, M. P., Moser, E. I., & Moser, M. B. (2004). Spatial representation in the entorhinal cortex. *Science*, 305(5688), 1258–64.
- Gothard, K. M., Skaggs, W. E., Moore, K. M., & McNaughton, B. L. (1996). Binding of hippocampal CA1 neural activity to multiple reference frames in a landmark-based navigation task. *Journal of Neuroscience*, 16(2), 823–35.
- Hardcastle, K., Ganguli, S., & Giocomo, L. M. (2015). Environmental boundaries as an error correction mechanism for grid cells. *Neuron*, 86(3), 827–839.
- Howard, M. W., Luzardo, A., & Tiganj, Z. (2019). Evidence accumulation in a laplace decision space. *Computational Brain and Behavior*.

- Howard, M. W., MacDonald, C. J., Tiganj, Z., Shankar, K. H., Du, Q., Hasselmo, M. E., & Eichenbaum, H. (2014). A unified mathematical framework for coding time, space, and sequences in the hippocampal region. *Journal of Neuroscience*, 34(13), 4692-707. doi: 10.1523/JNEUROSCI.5808-12.2014
- Howard, M. W., Shankar, K. H., & Tiganj, Z. (2015). Efficient neural computation in the Laplace domain. In *Cognitive computation: Integrating neural and symbolic computation*.
- Jin, D. Z., Fujii, N., & Graybiel, A. M. (2009). Neural representation of time in cortico-basal ganglia circuits. *Proceedings of the National Academy of Sciences*, 106(45), 19156-19161.
- Jutras, M. J., & Buffalo, E. A. (2010). Recognition memory signals in the macaque hippocampus. *Proceedings of the National Academy of Sciences*, 107(1), 401-406.
- Killian, N. J., Jutras, M. J., & Buffalo, E. A. (2012). A map of visual space in the primate entorhinal cortex. *Nature*, 491(7426), 761-4. doi: 10.1038/nature11587
- Killian, N. J., Potter, S. M., & Buffalo, E. A. (2015). Saccade direction encoding in the primate entorhinal cortex during visual exploration. *Proceedings of the National Academy of Sciences USA*, 112(51), 15743-8. doi: 10.1073/pnas.1417059112
- Kraus, B. J., Brandon, M. P., Robinson, R. J., Connerney, M. A., Hasselmo, M. E., & Eichenbaum, H. (2015). During running in place, grid cells integrate elapsed time and distance run. *Neuron*, 88(3), 578-589.
- Kraus, B. J., Robinson, R. J., 2nd, White, J. A., Eichenbaum, H., & Hasselmo, M. E. (2013). Hippocampal "time cells": time versus path integration. *Neuron*, 78(6), 1090-101. doi: 10.1016/j.neuron.2013.04.015
- Leitner, F. C., Melzer, S., Lütcke, H., Pinna, R., Seeburg, P. H., Helmchen, F., & Monyer, H. (2016). Spatially segregated feedforward and feedback neurons support differential odor processing in the lateral entorhinal cortex. *Nature Neuroscience*.
- Lever, C., Burton, S., Jeewajee, A., O'Keefe, J., & Burgess, N. (2009). Boundary vector cells in the subiculum of the hippocampal formation. *Journal of Neuroscience*, 29(31), 9771-7.
- Liu, Y., Tiganj, Z., Hasselmo, M. E., & Howard, M. W. (2019). A neural microcircuit model for a scalable scale-invariant representation of time. *Hippocampus*.
- MacDonald, C. J., Carrow, S., Place, R., & Eichenbaum, H. (2013). Distinct hippocampal time cell sequences represent odor memories immobilized rats. *Journal of Neuroscience*, 33(36), 14607-14616.
- MacDonald, C. J., Lepage, K. Q., Eden, U. T., & Eichenbaum, H. (2011). Hippocampal "time cells" bridge the gap in memory for discontinuous events. *Neuron*, 71(4), 737-749.
- Mau, W., Sullivan, D. W., Kinsky, N. R., Hasselmo, M. E., Howard, M. W., & Eichenbaum, H. (2018). The same hippocampal CA1 population simultaneously codes temporal information over multiple timescales. *Current Biology*, 28, 1499-1508.
- Meister, M. L., & Buffalo, E. A. (2018). Neurons in primate entorhinal cortex represent gaze position in multiple spatial reference frames. *Journal of Neuroscience*, 38(10), 2430-2441.
- Mello, G. B., Soares, S., & Paton, J. J. (2015). A scalable population code for time in the striatum. *Current Biology*, 25(9), 1113-1122.
- Meyer, T., & Rust, N. C. (2018). Single-exposure visual memory judgments are reflected in inferotemporal cortex. *Elife*, 7. doi: 10.7554/eLife.32259
- Milner, B. (1959). The memory defect in bilateral hippocampal lesions. *Psychiatric research reports*, 11, 43.
- Momennejad, I., & Howard, M. W. (2018). Predicting the future with multi-scale successor representations. *bioRxiv*, 449470.
- Morey, R. D. (2008). Confidence intervals from normalized data: A correction to Cousineau (2005). *Tutorial in Quantitative Methods for Psychology*, 4(2), 61-64.
- Naya, Y., & Suzuki, W. (2011). Integrating what and when across the primate medial temporal lobe. *Science*, 333(6043), 773-776.

- O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map. preliminary evidence from unit activity in the freely-moving rat. *Brain Research*, 34(1), 171-175.
- Pastalkova, E., Itskov, V., Amarasingham, A., & Buzsaki, G. (2008). Internally generated cell assembly sequences in the rat hippocampus. *Science*, 321(5894), 1322-7.
- Pilkiw, M., Insel, N., Cui, Y., Finney, C., Morrissey, M. D., & Takehara-Nishiuchi, K. (2017). Phasic and tonic neuron ensemble codes for stimulus-environment conjunctions in the lateral entorhinal cortex. *Elife*, 6, e28611.
- Ratcliff, R., & Murdock, B. B. (1976). Retrieval processes in recognition memory. *Psychological Review*, 83(3), 190-214.
- Salz, D. M., Tiganj, Z., Khasnabish, S., Kohley, A., Sheehan, D., Howard, M. W., & Eichenbaum, H. (2016). Time cells in hippocampal area CA3. *Journal of Neuroscience*, 36, 7476-7484.
- Shankar, K. H., & Howard, M. W. (2012). A scale-invariant internal representation of time. *Neural Computation*, 24(1), 134-193.
- Shankar, K. H., & Howard, M. W. (2013). Optimally fuzzy temporal memory. *Journal of Machine Learning Research*, 14, 3753-3780.
- Solstad, T., Boccara, C. N., Kropff, E., Moser, M. B., & Moser, E. I. (2008). Representation of geometric borders in the entorhinal cortex. *Science*, 322(5909), 1865-8.
- Squire, L. R., Stark, C. E., & Clark, R. E. (2004). The medial temporal lobe. *Annual Review of Neuroscience*, 27, 279-306.
- Takeuchi, D., Hirabayashi, T., Tamura, K., & Miyashita, Y. (2011, Mar). Reversal of interlaminar signal between sensory and memory processing in monkey temporal cortex. *Science*, 331(6023), 1443-7. doi: 10.1126/science.1199967
- Terada, S., Sakurai, Y., Nakahara, H., & Fujisawa, S. (2017). Temporal and rate coding for discrete event sequences in the hippocampus. *Neuron*, 94, 1-15.
- Tiganj, Z., Cromer, J. A., Roy, J. E., Miller, E. K., & Howard, M. W. (2018). Compressed timeline of recent experience in monkey LPFC. *Journal of Cognitive Neuroscience*, 30, 935-950.
- Tiganj, Z., Kim, J., Jung, M. W., & Howard, M. W. (2017). Sequential firing codes for time in rodent mPFC. *Cerebral Cortex*, 27, 5663-5671.
- Tiganj, Z., Shankar, K. H., & Howard, M. W. (2017). Scale invariant value computation for reinforcement learning in continuous time. In *AAAI 2017 spring symposium series - science of intelligence: Computational principles of natural and artificial intelligence*.
- Tsao, A., Sugar, J., Lu, L., Wang, C., Knierim, J. J., Moser, M.-B., & Moser, E. I. (2018). Integrating time from experience in the lateral entorhinal cortex. *Nature*, 561, 57-62. Retrieved from <https://doi.org/10.1038/s41586-018-0459-6> doi: 10.1038/s41586-018-0459-6
- Tulving, E. (1983). *Elements of episodic memory*. New York: Oxford.
- Wang, J., Narain, D., Hosseini, E. A., & Jazayeri, M. (2018). Flexible timing by temporal scaling of cortical responses. *Nature neuroscience*, 21(1), 102.
- Wilson, M. A., & McNaughton, B. L. (1993). Dynamics of the hippocampal ensemble code for space. *Science*, 261, 1055-8.
- Xiang, J.-Z., & Brown, M. (1998). Differential neuronal encoding of novelty, familiarity and recency in regions of the anterior temporal lobe. *Neuropharmacology*, 37(4-5), 657-676.

Methods

Subjects, training, and surgery

Two, male rhesus macaques (*Macaca mulatta*), 10 and 11 years old, and weighing 13.8 and 16.7 kg respectively, were trained to sit in a primate chair (Crist Instrument Company, Inc., Hagerstown, MD) and to release a touch-bar for fruit slurry reward delivered through a tube. The monkeys were trained to perform various tasks by releasing the touch-bar at appropriate times relative to visual stimuli presented on a screen. Magnetic resonance images of each monkey's head were made both before and after surgery to plan and confirm implant placement. Separate surgeries were performed to implant a head post, then months later, a recording chamber, and finally a craniotomy within the chamber. All experiments were performed in accordance with protocols approved by the Emory University and University of Washington Institutional Animal Care and Use Committees.

Electrophysiology

Each recording session, a laminar electrode array (AXIAL array with 13 channels, FHC, Inc.) mounted on a microdrive (FHC, Inc.) was slowly lowered into the brain through the craniotomy. Magnetic resonant images along with the neural signal were used to guide the penetration. Spikes and local field potentials were recorded using hardware and software from Blackrock, Inc., and neural data were sampled at 30 kHz. A 500 Hz high-pass filter was applied, as well as an electric line cancellation at 60 Hz. In some recording sessions, a channel without any spiking activity was used as a reference electrode in order to subtract artifact noise (e.g., reward delivery, movement of the monkey). Spikes were sorted offline into distinct clusters using principal components analysis (Offline Sorter, Plexon, Inc.). Sorted clusters were then processed further by custom code in MATLAB to eliminate any data where minimum inter-spike interval was less than 0.001 s, and to identify any missed changes in signal (e.g., decreased amplitude in the waveform of interest, a new waveform appearing), using a raster and plots of waveforms across the session for each neuron. When change in signal was identified, appropriate cuts were made to exclude compromised spike data from before or after a change point. 455 potential single neurons originally isolated in Offline Sorter were reduced to 357 single neurons. To further ensure recording location within the entorhinal cortex and identify from which cortical layers neurons were recorded, we examined each session's data for the stereotypical, electrophysiological signature produced across entorhinal cortical layers at the onset of saccadic eye movement (Killian et al., 2012; Killian, Potter, & Buffalo, 2015; Meister & Buffalo, 2018). One recording session, which other electrode placement metrics suggest was conducted above the entorhinal cortex within the hippocampus, lacked this electrophysiological signature and was excluded from further analysis (8 single neurons were excluded from being categorized as entorhinal cells). No recording sessions showed the current source density electrophysiological signature of adjacent perirhinal cortex (Takeuchi, Hirabayashi, Tamura, & Miyashita, 2011) at stimulus onset.

Experimental design and behavioral task

For all recordings, the monkey was seated in a dark room, head fixed and positioned so that the center of the screen (54.1 cm × 29.9 cm LCD screen, 120 Hz refresh rate, 1280

$\times 720$ pixels, BenQ America Corp., Costa Mesa, CA) was aligned with his neutral gaze position and 60 cm away from the plane of the his eyes (equating to 25 screen pixels per degree of visual angle, or $1^\circ/\text{cm}$). Stimulus presentation was controlled by a PC running Cortex software (National Institute of Mental Health, Bethesda, MD). Gaze location was monitored at 240 Hz with an infrared eye-tracking system (I-SCAN, Inc., Woburn, MA). Gaze location was calibrated before and during each recording session with calibration trials in which the monkey held a touch-sensitive bar while fixating a small (0.5°) gray square presented at various locations on the screen. The square turned yellow after a brief delay chosen uniformly from the interval from 0.40 s to 0.75 s. The monkey was required to release the bar in response to the color change for delivery of the fruit slurry reward. The subtlety of the color change forced the monkey to fixate the location of the small square to correctly perform those trials, therefore allowing calibration of gaze position to the displayed stimuli. Specifically, the gain and offset of the recorded gaze position were adjusted so that gaze position matched the position of the fixated stimulus. Throughout the session, intermittent calibration trials enabled continual monitoring of the quality of gaze position data and correction of any drift.

Before each image presentation, a crosshair ($0.3^\circ \times 0.3^\circ$) appeared in one of eighteen possible screen locations. Once gaze position registered within a $3^\circ \times 3^\circ$ window around the crosshair and was maintained within that spatial window for between 0.50 and 0.75 s (chosen uniformly), the image was presented. Images were large, complex images downloaded from the public photo-sharing website, Flickr (www.flickr.com). If necessary, images were re-sized by the experimenter for stimulus presentation (sized $30^\circ \times 15^\circ$ for Monkey WR and $30^\circ \times 25^\circ$ for Monkey MP). Monkeys freely viewed the image, and then the image vanished after gaze position had registered within the image frame for a cumulative 5 seconds. No food reward was given during image viewing trials. Each image presentation was followed by three calibration trials.

Image stimuli were unique to each session, and each image was presented twice within a session about 20 to 40 minutes apart. Images were presented in a block design so that novel and previously-seen images were presented throughout the session. Within a trial block, novel images (30 or 60) would first be shown, and then presented again in pseudorandom order. After completing a block of trials, a new block of trials would begin. In the first 16 sessions, a three-block design of 60 image presentations (30 novel) per block was used, with a total maximum of 180 image presentations per session. In the rest of the sessions ($n = 41$), there were a total maximum of 240 image presentations across two trial blocks (120 image presentations of which 60 were novel within each trial block).

Analysis of Neural Firing Fields

In order to determine temporal firing fields, spikes were analyzed using a custom maximum likelihood estimation script run in MATLAB 2016a. We calculated model fits across all trials available for each particular neuron. Fits of nested models were compared using a likelihood ratio test. In the present paper, we considered three models: a constant firing model, a model adding a Gaussian time term, and an ex-Gaussian model for which the time term was given by the convolution of a Gaussian and an exponential time term. The constant model,

$$M_{\text{const}}(t; a_o) = a_o \quad (1)$$

consisted of a single parameter a_o that predicted the constant probability of a spike at each time t .

The ex-Gaussian model describes the temporally-modulation of the firing field as the convolution of a Gaussian function with an exponentially decaying function:

$$M_{\text{ex-gauss}}(t; a_o, a_1, \sigma, \mu, \tau) = a_o + a_1 \int_{-\infty}^{\infty} e^{-\frac{(t-\mu)^2}{2\sigma^2}} e^{-\frac{t}{\tau}} dt \quad (2)$$

The ex-Gaussian distribution has been used extensively in studies of human response time data for many years (e.g., Ratcliff & Murdock, 1976). Here we can see that in the limit as $\tau \rightarrow 0$ the exponential function becomes a delta function and the convolution in Eq. 2 reduces to a Gaussian function. Similarly, in the limit as $\sigma \rightarrow 0$ the Gaussian function becomes a delta function and the convolution becomes an exponential function starting at μ . As such, this model is able to describe a range of peak firing times as well as varying degrees of skew (Fig. 1).

Two terms, a_o and a_1 describe the contributions of the constant and time-modulated terms. Three parameters describing the shape of the temporally-modulated term (Figure S1). μ and σ describe the mean and standard deviation of the Gaussian distribution, which estimates the time that a neuron's response maximally deviates from baseline, and the variability in that response time respectively. τ measures the time constant of the exponential decay, and captures the time that a neuron has returned 63% of the way back to baseline.

To estimate parameters of Eq. 2 numerically we used an explicit form for the solution of the convolution in Eq. 2:

$$M_{\text{ex-gauss}}(t; a_o, a_1, \sigma, \mu, \tau) = a_o + \frac{a_1}{2} e^{\frac{1}{2\tau}(2\mu + \frac{\sigma^2}{\tau} - 2t)} \text{erfc}\left(\frac{\mu + \frac{\sigma^2}{\tau} - t}{\sqrt{2}\sigma}\right) \quad (3)$$

where erfc is the complementary error function. μ was allowed to take values between 0 and 5 s. τ was allowed to take values between 0 and 20 s. σ was allowed to take values between 0 and 1 s. Likelihood was estimated in a 5.5 s long window, that included the 0.5 s prior to presentation of the image and the 5 s after presentation of the image. Kendall's τ correlation coefficients and Bayes Factors were calculated using JASP statistical software.

To validate the three time-varying parameters of the ex-Gaussian model, we compared the fits of this model to a simpler Gaussian model with only two time-varying parameters, μ and σ .

$$M_{\text{gauss}}(t; a_o, a_1, \sigma, \mu) = a_o + a_1 e^{-\frac{(t-\mu)^2}{2\sigma^2}} \quad (4)$$

Note that the Gaussian model is nested within the ex-Gaussian model Eq. 2, with $\tau = 0$. We evaluated these models for each neuron *via* a likelihood ratio test and counted the number of neurons that were better fit by the ex-Gaussian model at the 0.05 level, Bonferonni-corrected by the total number of 349 neurons (this is arguably conservative), changed their firing by at least 1 Hz, and reached a firing rate of at least 3 Hz. Of the 128 neurons that were better fit by the ex-Gaussian time-varying model, 49 showed a reliably better fit from the three-parameter ex-Gaussian than the two-parameter Gaussian temporal term. This is very different from chance and we conclude that the population was better described by the ex-Gaussian model.

Linear Discriminant Analysis (LDA)

An LDA classifier was used to decode time since onset of the image from the population using information from all neurons (subject to the restrictions described below).

LDA Implementation. Even and odd trials were used for training and testing respectively. The number of available trials varied for each neuron. To mitigate any problems from this, several steps were taken. First, four neurons with less than 30 trials each were entirely excluded from this analysis. Neurons with less than 200 trials were bootstrapped to 200 trials, while neurons with more than 200 trials were randomly down-sampled. Time was discretized into 0.25 s bins. For each bin of each trial, the firing rate was calculated across neurons. To avoid errors due to a singular covariance matrix, a small amount of uniform random noise (between 0 and 0.25×10^{-13} Hz) was added to the firing rate in each time bin. The averaged firing rate of each time bin for each training trial across all neurons made up an element of the training data. The averaged firing rate of each time bin for each testing trial across all neurons made up an element of the testing data. LDA was implemented using the MATLAB function “classify.” This function takes in the training data, testing data, labels for the training data, and a selection of the method of estimation for the covariance matrix (the option “linear” was used) and returns a posterior distribution across bins for each test trial.

Estimating the duration of temporal coding. To verify that the temporal information could be decoded across the time interval, the LDA was repeated for successively fewer bins, removing the earliest time bin with each repetition. If time since presentation of the image can be decoded above chance using only information after time T , one can conclude that the population contained temporal information about time at least a time T after presentation of the image. For each repetition the decoder was tested by training it on permuted data set generated by randomly permuting the time labels of the time bins. We compared the absolute error of the actual data to the distribution generated from 1,000 permutations. The classifier’s performance was considered better than chance if fewer than 10/1,000 permutations gave a better result than the unpermuted data. The decoder stopped performing better than chance when time bins below 2 s were omitted. Thus, we can conservatively conclude that entorhinal neurons carried information about the time of image presentation for at least 1.75 s.

Population similarity

We constructed population vectors to evaluate the degree to which the population of entorhinal neurons was sensitive to the identity of the visual image. For each repeated image, we created two population vectors, one corresponding to the first presentation and one corresponding to the second presentation. Each vector was created from the mean firing activity of all the available neurons during the 5 s of free viewing. Mean firing rates were normalized by each neuron’s maximum average firing rate so that firing rates ranged from 0 to 1. Only blocks where all images were presented twice were considered. In order to control for different block lengths between sessions, only the first 30 images presented in each block were used. Similarity was measured by the cosine similarity of the two population vectors. We compared the cosine similarity of two presentations of the same image to the

first presentation of one image and the second presentation of a different image. As a control we instead compared the population vector from the repetition of an image to the adjacent near-neighbors of the original image presentation. Within session error bars were calculated using the method outlined in Morey (2008).

Supplementary Information

To illustrate how the analyses in this paper illustrate properties of different temporal coding schemes, Figures 1 and S3 include predictions from two different schemes for coding temporal information. One coding scheme is given by a population of ideal exponentially-decaying temporal context cells. The other coding scheme is given by a population of ideal sequentially activated time cells. These ideal populations are based on theoretical work (e.g., Shankar & Howard, 2013).

The firing rate of the i th ideal temporal context cells a time t after the presentation of the image is given by

$$F_i(t > 0) = e^{-s_i t} \quad (\text{S1})$$

The rate constants s_i controls how fast ideal cell i decays. It can be shown that at time t after the image, the population codes the Laplace transform of the time in the past at which the image was presented (Shankar & Howard, 2012). The values of s_i were chosen such that s_i/s_{i+1} was constant for all i . Equation S1 was used to generate the populations in Figure 1d and on the right of Figure S3.

The population of ideal sequentially activated time cells can each be described by values of s that correspond to those of the ideal temporal context cells. However, rather than rising immediately at time $t = 0$ and then decaying exponentially, the firing rate of these ideal time cells are given by

$$\tilde{f}(t > 0) \propto s_i (s_i t)^k e^{-s_i t} \quad (\text{S2})$$

where k is a constant that was set to 4 in these figures. The properties of this population can be better understood by considering their activity as a function of $\tau_i^* \equiv k/s_i$, which has the units of time rather than rate. Then,

$$\tilde{f}(t > 0) \propto \frac{1}{\tau_i^*} \left(\frac{t}{\tau_i^*} \right)^k e^{-\frac{kt}{\tau_i^*}} \quad (\text{S3})$$

Here we can see the firing rate of cell i is peaked a characteristic time after the image is presented. The time-dependence of the firing rate in Eq. S3 is the product of a growing polynomial term, $\left(\frac{t}{\tau_i^*} \right)^k$ and a decaying exponential term, $e^{-\frac{kt}{\tau_i^*}}$. This product is zero at time zero (because of the polynomial term) and zero at times much bigger than τ_i^* (because of the exponentially-decaying term). The product peaks somewhere in between that depends on the value of τ_i^* (it can be shown that this expression peaks precisely at τ_i^* Shankar & Howard, 2012).

The choice of s_i was chosen as in the ideal temporal context cells. One can understand these two populations as encoding the Laplace transform of the delta function describing the time since the onset of the stimulus and an scale-invariant estimate of that function, respectively (Shankar & Howard, 2012).

Relationship of μ to τ for ideal temporal context cells and for time cells. Ideal cells generated by these two coding schemes, Eq. S1 and Eq. S3, will show different properties in their relationship between parameters μ , σ and τ estimated from the ex-Gaussian model

(see Methods, Eq. 2). It is easy to convince oneself from Eq. S1 that cells generated from the exponentially-decaying term will have systematic changes in their value of τ_i (τ_i is just $1/s_i$) but there should not be a relationship between τ_i and μ_i or between τ_i and σ_i . We would expect μ and σ to result from the delay associated with the time it takes for the image to cause activation in the population of ideal temporal context cells. That is, ideal temporal context cells would be expected to have about the same value of μ that is independent of τ and small values of σ . In contrast, for a population of ideal sequentially activated time cells because there are a range of peak times, we would expect to see a range of μ_i . Moreover, because all of the cells have the same shape only rescaled with the peak time τ_i^* we would expect to see that μ_i correlates with both τ_i and σ_i .

Decoding of time from ideal temporal context cells and ideal time cells is possible and results in similar properties

To illustrate the ability of both temporal decoding schemes to carry information about time, we applied the LDA used on the data to two populations of simulated spiking data generated using the ideal equations shown above.

To generate the simulated spiking data the following procedure was used. The ideal equation was treated as a normalized firing rate, where the maximum value (across all time constants) of the equation was normalized to 1. For each simulated millisecond, the probability of firing was set using this normalized simulated firing rate, a maximum firing rate selected for consistency with empirical spiking data, and a constant background firing rate. The probability of a spike within each millisecond is equal to the maximum firing rate multiplied by the normalized simulated firing rate plus a constant background firing rate. 1000 trials were simulated, 500 trials were used for training and 500 trials for testing.

Exponentially-decaying temporal context cells can be used to decode temporal information. The scale over which each neuron contributes maximally to decoding should be on the order of its time constant. The simulated exponentially decaying cells were constructed using Eq. S1 with a spectrum of τ^* , ranging from 0.05 s to 40 s spaced geometrically with a ratio of approximately 1.1 between successive time constants (for a total of 70 simulated cells). For exponentially decaying cells, their time constant is equal to τ^*/k and k is equal to 4, so these cells have time constants ranging from 0.0125 s to 10 s. The simulated spiking data was then generated from this using a max firing rate of 20 Hz and a constant background firing rate of 20 Hz. This range of time constants means that the smallest time constant is less than 1/10 the duration of the bin size, and that the largest time constant is over twice the total duration being decoded. This reduces the possibility of any edge effects. Because there are fewer cells with slow time constants (because of the geometric spacing) decoding accuracy should go down with the passage of time. Figure S3 b and d shows results of the same LDA decoder used on the empirical data when applied to this simulated population of temporal context cells. As can be seen from the figure, temporal information decreased and error increased as a function of time. A linear regression of decoding error as a function of time showed a reliable slope, 0.13 ± 0.02 with $p < 0.001$, $R^2 = 0.73$, as well as intercept 0.12 ± 0.05 with $p < 0.05$; $df = 16$. Thus the temporal information decreased as time elapsed.

We also applied the LDA to a population of ideal time cells. These time cells were

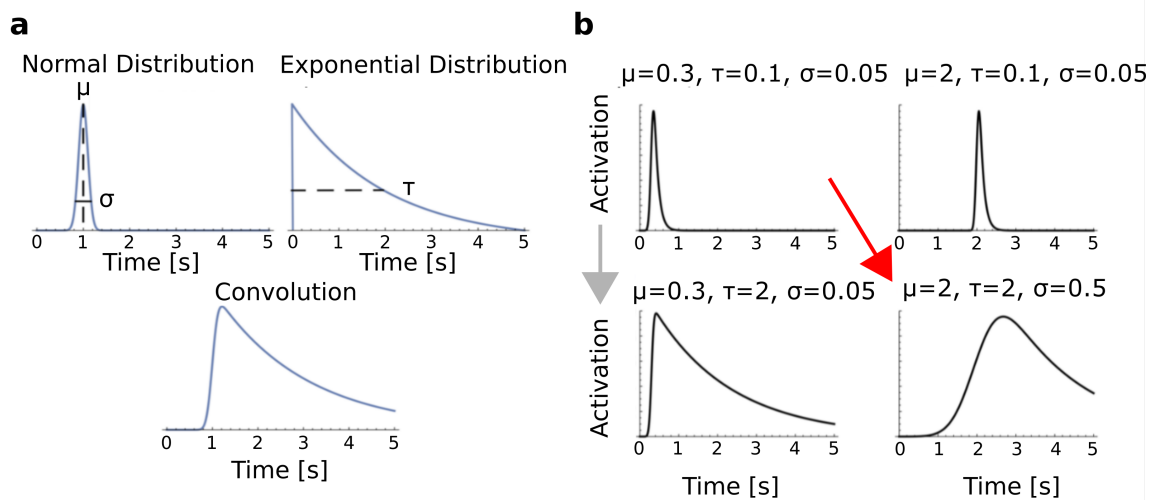


Figure S1. : Modeling neuron responses using an ex-Gaussian distribution. **a**, The ex-Gaussian model of a neuron's response (bottom) is formed by convolving a Gaussian distribution (with parameters μ and σ) with an exponential distribution (with parameter τ) (top). **b**, In the ex-Gaussian model, an increase in μ shifts the distribution to the right, an increase in σ widens the central peak of the distribution, and an increase in τ lengthens its decay rate. The gray and red arrows correspond to the parameter changes expected in the process of fitting the response of temporal context cells and sequentially activated time cells respectively.

constructed using Eq. S2, with values of τ^* ranging from 0.05 s to 40 s, spaced geometrically with a ratio of approximately 1.1 (for a total of 70 simulated cells). In the case of time cells, their peak time is precisely equal to τ^* , so their peaks range from 0.05 s to 40 s. The simulated spiking data was then generated from this using a max firing rate of 40 Hz and a constant background firing rate of 1 Hz. The width of the receptive fields expands with the peak time and because there are fewer neurons with peak times later in the delay, the decoding accuracy of this population of time cells should also go down with the passage of time. Figure S3a and c shows the results of the LDA applied to this set of simulated time cells. As can be seen from the figure, error increased as a function of time. A linear regression of decoding error as a function of time showed a reliable slope, $0.14 \pm .03$ with $p < 0.001$, $R^2 = 0.57$, as well as intercept 0.22 ± 0.08 with $p < 0.05$; $df = 16$. Thus the temporal information decreases as time elapses.

Despite the fact that these two populations have different forms of temporal responsiveness, they both code information about time with similar properties. This is a natural consequence of the fact that the time cell population is just a linear transformation of the temporal context cell population (Shankar & Howard, 2012, 2013).

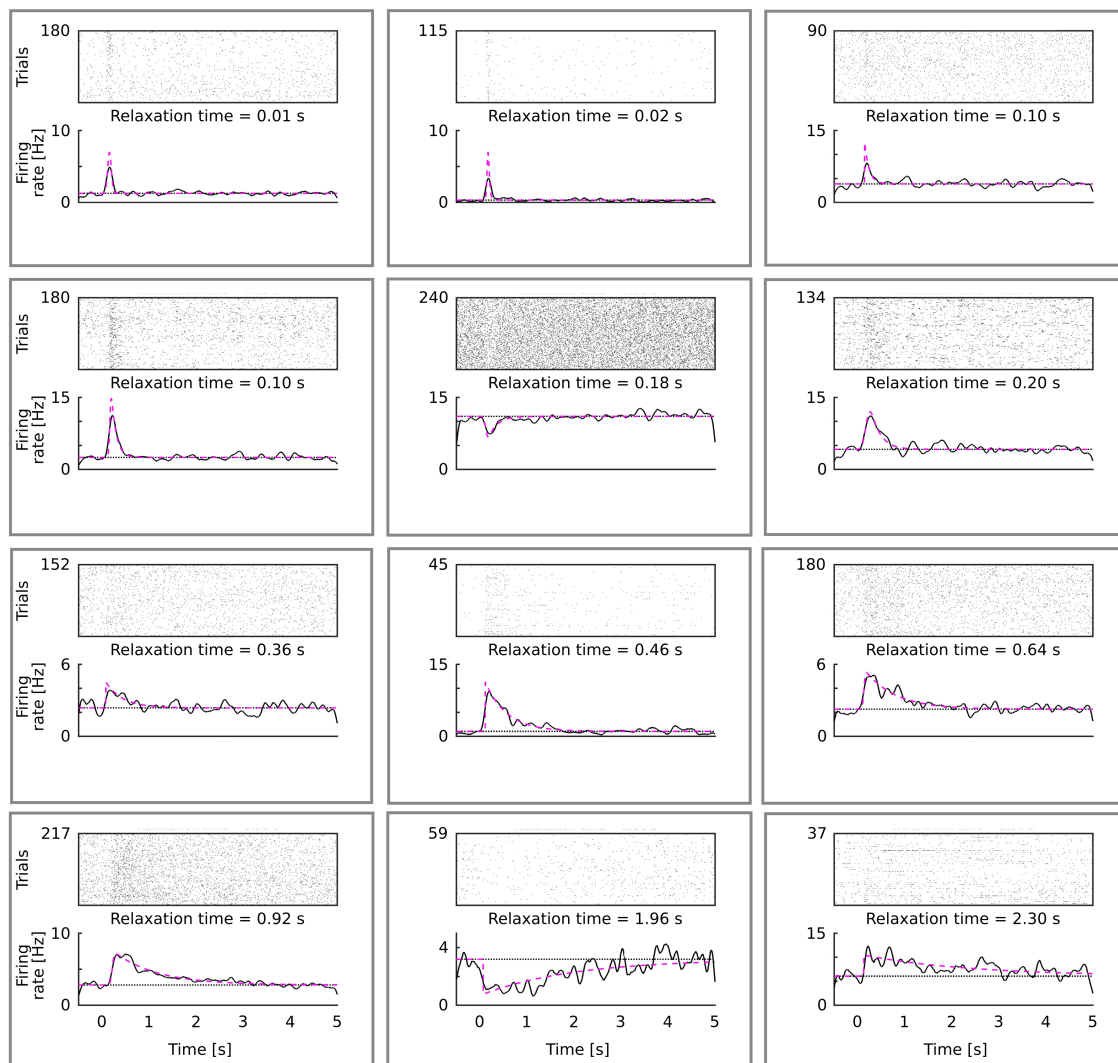


Figure S2. : Additional examples of temporal context cells. Format is as in Figure 2a.

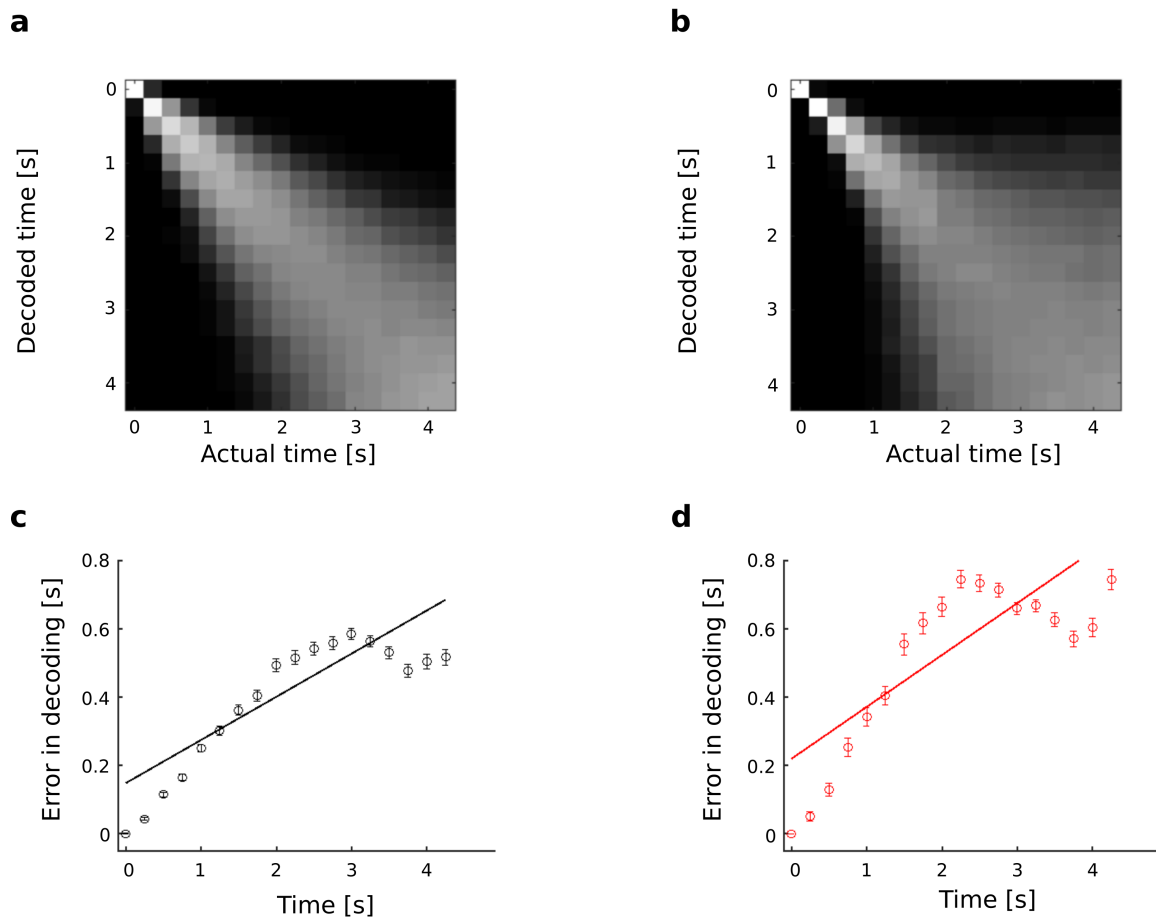


Figure S3. : Time can be decoded from both ideal time cells and ideal exponentially decaying temporal context cells. Simulated noisy time cells (**a, c**) and exponentially decaying temporal context cells (**b, d**). Time is binned in 0.25 s bins. A linear decoder was trained on odd trials and tested on even trials. **a, b** The log of the average posterior probabilities of the classify function are averaged across trials for each time bin. Perfect decoding would manifest as a bright white diagonal. The posterior distribution of the classifier shows increasing uncertainty in the decoded time for both populations. **c, d** Averaged absolute value of decoding error. The decoding error goes up with time for both populations, as shown by the fitted regression line (black line).

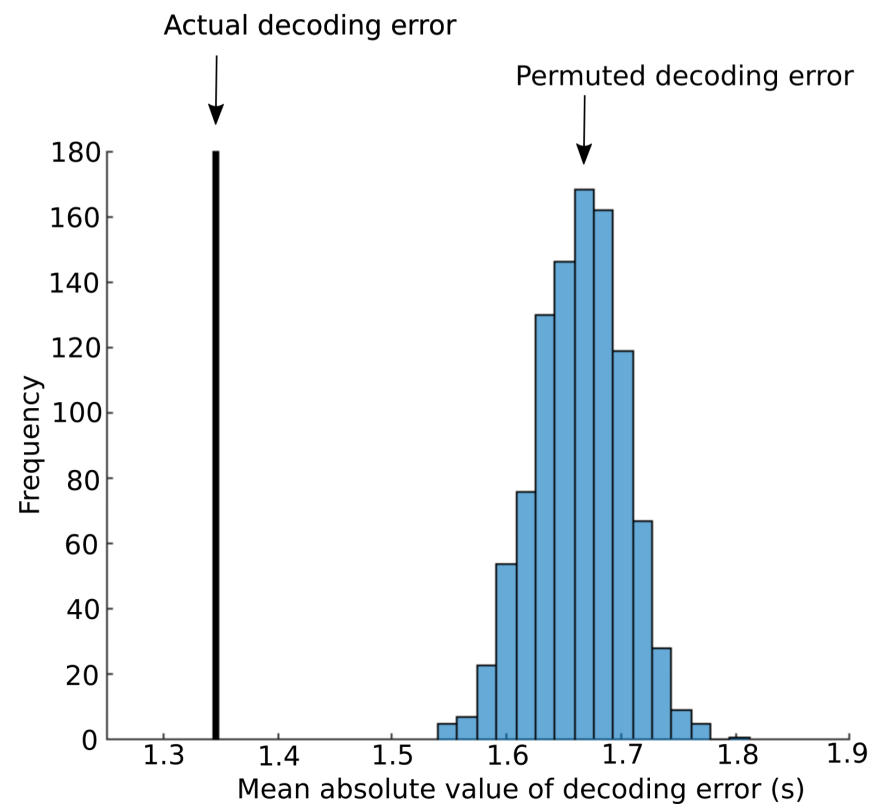


Figure S4. : **Time decoded from entorhinal firing showed a much lower decoding error than time decoded from permuted data.** The histogram shows the distribution of mean absolute value of decoding error for the permuted data. In this analysis, the training data were the same as the data used in the actual decoder, except the true time labels were randomly permuted. The mean absolute value of decoding error for the original data (1.35 s) is marked by a vertical line. This value is more than seven standard deviations from the mean of decoding error of the permuted data.

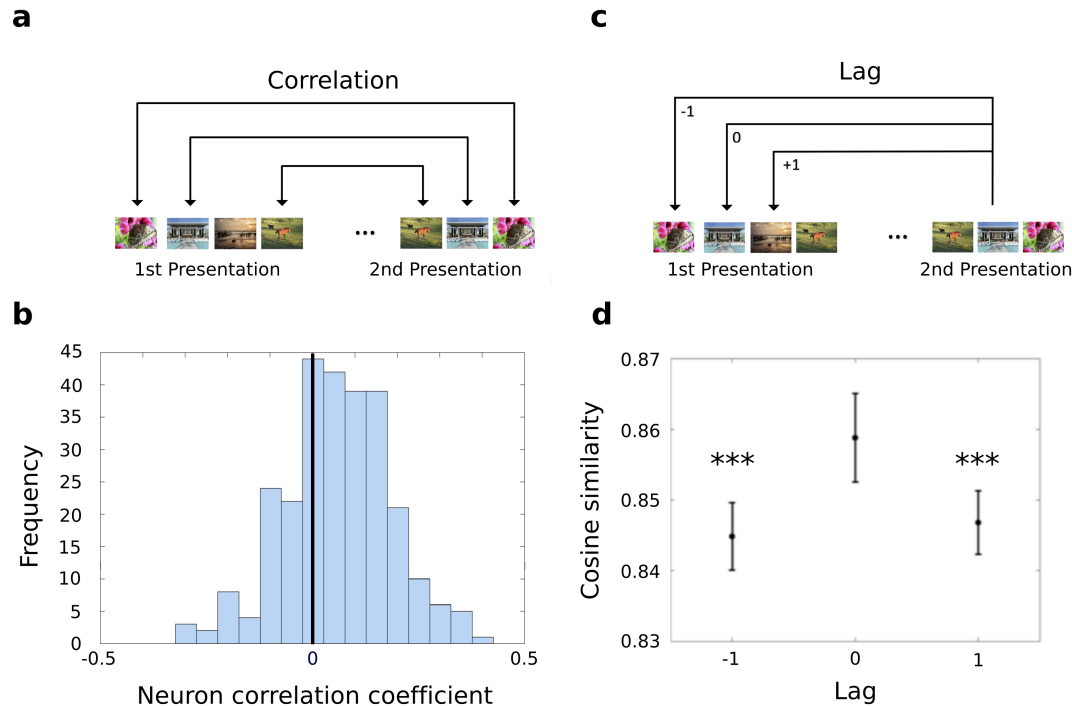


Figure S5. : The entorhinal ensemble carries information about stimulus identity.

a, We measured the firing rate of each unit to first and second presentations of each image. For each neuron that was recorded for an entire block of images ($n = 270$), we measured the correlation in firing rate across images using Kendall's τ . **b**, The distribution of Kendall's τ for all entorhinal neurons is shown. This distribution is reliably different from zero ($p < 0.001$, see text for details) indicating that entorhinal firing was sensitive to image identity. **c**, Schematic of a population similarity analysis. We measured the similarity of population response to the second presentation of each image to the first presentation of the same image (lag 0). As controls we also computed the similarity between the population response to the second presentation of an image and the responses to the images neighboring the first presentation of that image. Lag -1 refers to the similarity to the immediate predecessor of the image; lag $+1$ refers to the similarity to the immediate successor of the image. **d**, Cosine similarity of the population response to the second presentation of an image to the original presentation of an image (lag 0) and the images neighboring the original presentation of the image (lags -1 and $+1$) over 64 blocks of first and second presentations. The similarity to lag 0 is greater than the similarity to either of the neighboring images. *** indicates significance at the $p < 0.001$ level. Error bars correspond to the 95 % confidence interval of mean cosine similarity calculated across sessions.