

# Sequential firing codes for time in rodent mPFC

Zoran Tiganj

Department of Psychological and Brain Sciences, Center for Memory and Brain, Boston  
University, USA

Jieun Kim, Min Whan Jung

Center for Synaptic Brain Dysfunctions, Institute for Basic Science, Daejeon, Korea

Marc W. Howard

Department of Psychological and Brain Sciences, Center for Memory and Brain, Boston  
University, USA

## Abstract

We studied the firing correlates of neurons in the rodent medial PFC during performance of a temporal discrimination task. On each trial, the animal waited for a few seconds in the stem of a T-maze. The firing correlates within a trial gave us a means to assess firing on the scale of seconds. A subpopulation of units fired in a sequence consistently across trials during a circumscribed period during the delay interval. These sequentially activated “time cells” showed temporal accuracy that decreased as time passed as measured by both the width of their firing fields as well as the number of cells that fired at a particular part of the interval. In addition, most units showed gradual changes in their firing rate *across* trials. The time constants of the change in firing were distributed like a power law, with some units showing gradual changes over tens of minutes. The population of time cells showed temporal coding of decreasing temporal accuracy over the scale of a few seconds. Gradual changes across trials could reflect temporal coding over much longer scales as well.

## Introduction

A variety of brain regions have been implicated in interval timing over the scale of seconds to minutes, including the striatum (see Buhusi and Meck (2005) for a review) and medial prefrontal cortex (mPFC) (Mangels et al., 1998; Onoe et al., 2001; Kim et al., 2009). Recent evidence has shown that neural ensembles change gradually over periods of time from seconds to minutes in the mPFC (Hyman et al., 2012; Kim et al., 2013); gradual change in ensemble state could be used as a timing signal. For instance, Kim et al. (2013) recently showed that the ensemble state in the medial prefrontal cortex (mPFC) changed gradually during the delay period of a temporal discrimination task. Critically, Kim et al. (2013) found that the discriminability of the time during the delay that could be computed from

22 the ensemble similarity decreased with time elapsed. Decreasing accuracy with elapsing  
23 time is a hallmark of behavioral measures of memory and timing in both human and non-  
24 human animals (Lewis and Miall, 2009; Lejeune and Wearden, 2006; Wearden and Lejeune,  
25 2008).

26 There are many potential mechanisms that could cause a change in accuracy at the  
27 ensemble level as time elapses. For instance, a population of neurons whose firing rate  
28 changes monotonically as a function of the logarithm of the time during the delay would  
29 have this property; Kim et al. (2013) reported a population of units exhibiting just this  
30 pattern of results. However, there are other alternatives as well. For instance, several  
31 labs have reported “time cells” in the hippocampus that fire during circumscribed parts of  
32 a delay period (Pastalkova et al., 2008; Gill et al., 2011; Kraus et al., 2013; MacDonald  
33 et al., 2011). Different time cells fire at different times during the interval, enabling a  
34 population of time cells to generate a signal that could be used in interval timing. If the  
35 width of time cells’ firing fields increased with their time of peak firing (Howard et al., 2014;  
36 Kraus et al., 2013), then the population of time cells would be less able to distinguish times  
37 later in the interval. Similarly, if the density of time fields decreased as a function of time  
38 (Kraus et al., 2013), this would have the same consequence.

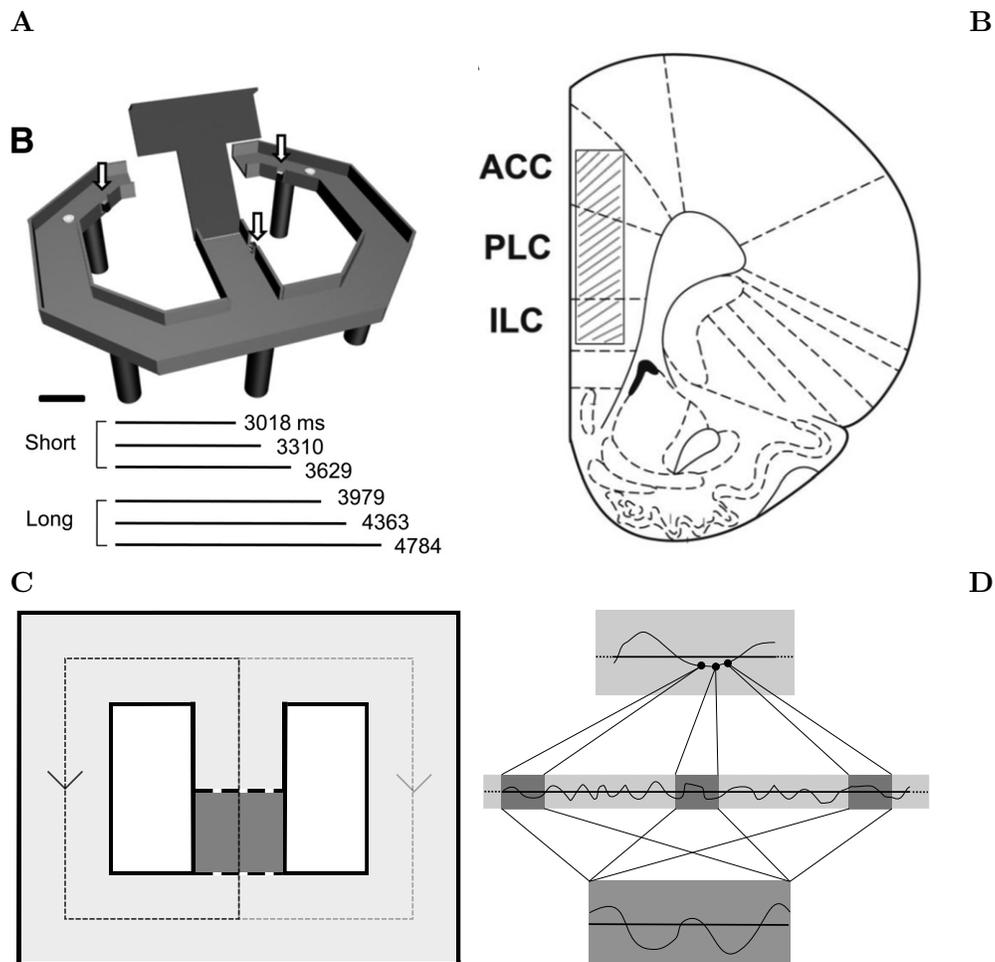
39 This paper reports the results of analyses on the data set initially reported in Kim  
40 et al. (2013). Kim et al. (2013) noted the existence of cells that fired during circumscribed  
41 periods of time during the delay interval (see for instance their Figure 3F). Here we study  
42 this phenomenon in more detail to determine if the mPFC contains a significant population  
43 of sequentially activated time cells and to determine if these cells code time in such a way  
44 that there is decreasing temporal accuracy as a function of time within the delay. In  
45 addition, we examined evidence for gradual changes of firing across scales much longer than  
46 a single trial, up to tens of minutes.

## 47 Methods

### 48 *Recordings and behavioral procedure*

49 The details of the behavioral task are described in Kim et al. (2013). On each pass  
50 through the maze (Figure 1A), the animal waited for a period of time in front of a T-junction  
51 (dark shaded area in Figure 1C). To obtain a water reward, the animal had to navigate to  
52 one goal when a short time interval ( $< 3.75$  s) was presented, and navigate to the opposite  
53 goal when a long time interval ( $> 3.75$  s) was presented. In this study, we analyzed only  
54 the data recorded during the delay intervals since in that period there were no behavioral  
55 demands on the animals. Recordings were made using tetrodes implanted in mPFC of three  
56 rats (Figure 1B).

57 A total of 993 well isolated single units were recorded. Of these, we eliminated 160  
58 units with mean firing rate  $< 1$  Hz during the waiting intervals. Additionally, in order to  
59 restrict our attention to units with spike waveforms that were stable over the recording  
60 session we eliminated 10 units with a difference of more than 10% in amplitude from the  
61 first to the last 5 min of each session. A total of 723 units contributed to the subsequent  
62 analyses.



*Figure 1.* **A.** The maze contained a drawbridge that required animals to wait at a particular location on each trial. If the delay was short, the animal was rewarded for turning one direction at the T; a long delay required a turn in the other direction. **B.** Schematic of recording locations (shaded regions). The diagram is a coronal section view of the brain (2.7 mm anterior to bregma). **A,** **B** reproduced from Kim et al. (2013). **C.** Temporal bisection task. The animal had to wait for one of six different delay intervals in a limited space (dark gray shaded area), and then navigate through either left or right path (dark gray and light gray dashed lines respectively), depending on the duration of the presented time interval. **D.** The across- and within-trial analysis. The schematic in the middle displays a snapshot of the timeline, represented with a full black line with dots at each end. Dark gray shaded areas are the delay intervals, and light gray shaded areas are the time when the animal was moving through the rest of the maze. The black line is a cartoon example of the firing rate from one cell. All the analysis was done on the neural activity recorded during the delay intervals. In the across-trial analysis (top plot) each delay interval is represented with the mean firing rate, while in the within-trial analysis (bottom plot) neural activity was averaged across all delay intervals.

63 *Analysis across time scales*

64 We examined the firing during delay intervals across two very different time scales  
65 (Figure 1D). First, we considered the firing of neurons as a function of time within the  
66 delay period. For this analysis we considered only the longest delay interval (almost 5 s).  
67 Second, we examined changes in firing from one delay period to the next. Because each  
68 delay period was separated by on average 20 s ( $20 \pm 14$  s) as the animal traversed back to  
69 the waiting location, this analysis allowed us to compare changes in firing over much longer  
70 time scales. We analyzed the first 164 trials in each recording session, meaning that we  
71 could assess changes in firing up to tens of minutes ( $164 \times 20$  s is more than 50 minutes).

72 *Classification of time cells*

73 Kim et al. (2013) reported a population of units that started firing prior to the initia-  
74 tion of the delay and decreased their firing as the delay proceeded and another population of  
75 units that increased their firing monotonically during the delay interval. Both groups could  
76 be responding to some event that preceded the delay interval or they could be predicting an  
77 event that follows the delay interval. In these analyses we restricted our attention to units  
78 that both started and stopped firing within the delay interval on trials in which the animal  
79 completed the task successfully. We first processed the data by smoothing the spike train  
80 recorded on each trial with a Gaussian-shaped window with 200 ms standard deviation. We  
81 then averaged the smoothed activity across correct trials. To be classified as a time cell,  
82 units had to satisfy several criteria. First, units had to exhibit an average firing frequency of  
83 at least 4 Hz over the delay interval and fire at least one spike in at least 15 different trials.  
84 Second, to identify units that showed variability in firing during the delay we required there  
85 be at least one time point in the delay interval where the unit's averaged firing rate was no  
86 more than 40% of its peak firing rate in the interval. Finally, we required that the unit's  
87 firing rate 400 ms before and after the averaged delay interval did not exceed the peak firing  
88 rate observed during the interval. The last criterion was set to avoid including cells which  
89 firing rate had a general tendency of growth or decay, even outside the delay interval.

90 *Quantifying the time scale of across-trial fluctuations*

91 To quantify long range gradual changes in neural activity we constructed a measure of  
92 the duration of each units' autocorrelation across trials. For each unit we took the average  
93 firing rates in the delay intervals of the first 164 trials in the recording session. We then  
94 computed the autocorrelation function of this time series. We defined the "time constant"  
95 of the unit as the time at which the autocorrelation function of the actual data fell within  
96 the first standard deviation of the autocorrelations of a surrogate data set constructed from  
97 1000 independent shuffles of the firing rates. This measure can produce time constants as  
98 small as zero trials for a unit that is not autocorrelated. Under most circumstances, the  
99 method cannot yield time constants longer than 82 trials. In reporting time constants, we  
100 multiply the number of trials by the average time of a trial (20 s) to give an intuitive sense  
101 of the scale of the autocorrelation.

102 *Estimating distributions using maximum likelihood*

103 Analyses of the within-trial activation generated distributions of the time point at  
 104 which units were maximally active. Across-trial analyses generated distributions of the  
 105 time constants across units. In order to characterize the form of these distributions, we fit  
 106 various models to the distribution. Given a value  $x$ , we computed the likelihood  $P(x|\theta)$   
 107 that value  $x$  given a model parameterized by  $\theta$ . For each model and each parameterization,  
 108 we estimated the joint probability of all of the values by taking the sum of the logarithm  
 109 of the likelihoods. Given that models we considered were either zero parameters (uniform  
 110 distribution) or one-parameter (exponential and power law distribution) we found the max-  
 111 imum likelihood estimate of the parameter by simply sweeping through all possible values of  
 112 the parameter. Models with different numbers of parameters were compared using standard  
 113 methods (AIC and BIC). To estimate a confidence interval on the parameter around the  
 114 best-fitting value  $\theta_o$ , we estimated the values  $\theta_-$  and  $\theta_+$  such that

$$\frac{\int_{\theta_-}^{\theta_o} P(\mathbf{x}, \theta') d\theta'}{\int_{-\infty}^{\theta_o} P(\mathbf{x}, \theta') d\theta'} = \frac{\int_{\theta_o}^{\theta_+} P(\mathbf{x}, \theta') d\theta'}{\int_{\theta_o}^{\infty} P(\mathbf{x}, \theta') d\theta'} = 0.95,$$

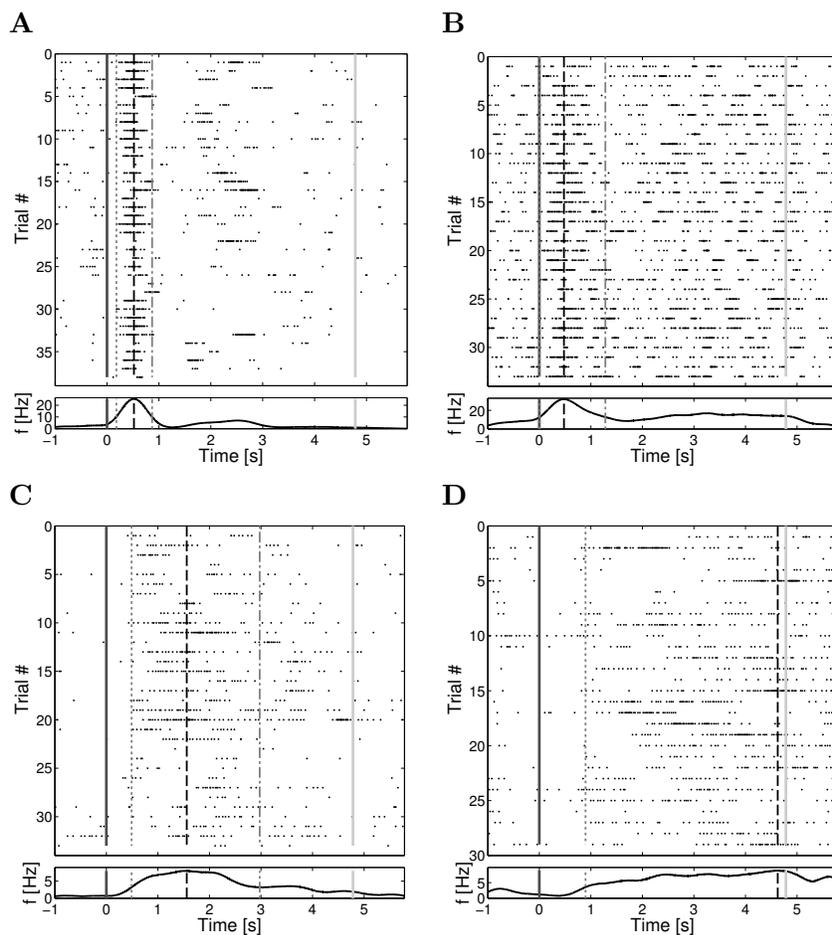
115 where  $\mathbf{x}$  is the entire set of values in the experimental data. The range between  $\theta_-$  and  $\theta_+$   
 116 thus contains 95% of the probability mass of the distribution.

117 **Results**118 *Temporal coding on the order of seconds*

119 From the within-trial analysis we identified a subpopulation of sequentially acti-  
 120 vated units that fired at a consistent, circumscribed time during delay trials (Figure 2).  
 121 These mPFC units appear to have firing correlates that resemble time cells observed  
 122 in the hippocampus (Kraus et al., 2013; Gill et al., 2011; MacDonald et al., 2011;  
 123 Pastalkova et al., 2008; Modi et al., 2014). A total of 122/723 units were classified as  
 124 time cells.

125 First, we note informally that the population of time cells decreased in its temporal  
 126 accuracy as time during the interval proceeds. Figure 3A shows the ensemble similarity  
 127 (cosine of the normalized firing rate vectors) of the population of time cells between all  
 128 pairs of time points during the delay period. This finding replicates the conclusions of Kim  
 129 et al. (2013) but restricting attention to the population of time cells. Further analyses  
 130 revealed two causes for the decrease in temporal accuracy. These can be read off from  
 131 Figure 3B, which shows the temporal profile of all 122 units classified as time cells, sorted  
 132 by their median spike time.

133 *The width of firing fields increased with the passage of time.* First, note that the width  
 134 of the central ridge in Figure 3B increases as one moves from the left of the plot to the right  
 135 of the plot. This suggests that the units that have elevated firing rate earlier in the delay  
 136 interval tend to have narrower time fields than the units that fire later in the delay interval.  
 137 This impression was confirmed by analyses of the across-units relationship between the time  
 138 of the peak firing rate and widths of the time fields across units. The width was defined as  
 139 the time that the activity in the averaged delay interval is above the 40% of its peak firing



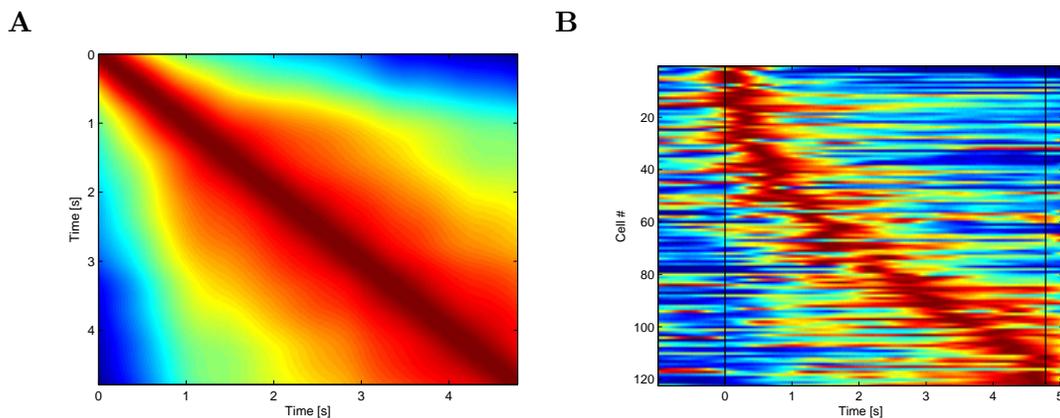
*Figure 2.* Examples of mPFC time cells that fired consistently across trials during a time window within the delay interval. Each of the four columns (A-D) displays activity of a single cell. The cells are ordered such that width of the time field and the peak time increase progressively from the first to the fourth cell. The top row shows raster plots and the bottom row shows the averaged trial activity. Dark gray and light gray lines mark the start and the end of delay intervals respectively. Gray dotted and dash-dotted lines mark the start and the end of the time fields respectively. Black dashed lines mark the time of the peak firing rate. The activity of the unit in D did not decrease to the threshold level after reaching the peak so only start of the time fields is marked.

140 rate in the interval. We found weak but significant correlation between the width and the  
141 peak time (Pearson’s correlation 0.34,  $p < .001$ ).

142 *Later times are represented by fewer cells than earlier times.* Second, the population  
143 of cells covers the entire delay interval, but not evenly. The number of cells with peak  
144 firing later in the interval is smaller than the number of cells with peak firing earlier in the  
145 interval. This can be seen from the fact that the central ridge does not follow a straight  
146 line, as would have been expected of a uniform distribution of peak times, but flattens as  
147 the interval proceeds. To quantify this, we examined the distribution of the peak times.  
148 We found the distribution was much more likely assuming a power law distribution than  
149 a uniform distribution ( $\Delta\text{AIC}=30$ ,  $\Delta\text{BIC}=33$ ) and much more likely with a power law  
150 distribution than an exponential distribution ( $\Delta\text{LL} = 7$ ), meaning that the likelihood of  
151 the data given the best-fitting power law distribution was about 1000 times greater than  
152 the likelihood of the data given the best-fitting exponential distribution. The best fitting  
153 value for the exponent of the power law was  $-.41$ . The 95% confidence interval did not  
154 overlap with zero ( $-.37$  to  $-.44$ ). This does not provide strong evidence that the “true”  
155 distribution is in fact power law rather than some other function with a long tail, but it  
156 does compellingly reject the uniform distribution, meaning that more units had time fields  
157 early in the delay than later in the delay.

158 *mPFC time cells and ramping cells convey comparable amount of temporal informa-*  
159 *tion.* We quantified how well the mPFC neuronal ensemble kept track of the elapse of time.  
160 The longest time interval (4784 ms) was divided into 10 equal-duration bins and the order  
161 of the middle eight bins was decoded based on neural activity within each bin using linear  
162 discriminant analysis (Kim et al., 2013). We compared the results on different populations  
163 of cells: all 722 cells (Figure 4A), all 122 time cells (Figure 4B) and 122 ramping cells  
164 (selected randomly from a total of 228 cells that exhibit ramping firing rate, Figure 4C).  
165 The number of selected ramping cells that were also time cells was 66. The mean error in  
166 the prediction of elapsed time was similar for all three populations. This suggests that pop-  
167 ulations of time cells and ramping cells can convey roughly the same amount of information  
168 about the elapse of time.

169 *Neither of these findings were an artifact of trial averaging.* To confirm that the  
170 properties seen in Figure 3 were not simply an averaging artifact, we repeated the analyses,  
171 but rather than taking the average smoothed firing rate as input, we took the average of  
172 the product of the smoothed firing rate on adjacent trials. In these alternate analyses, only  
173 temporally-specific firing that is consistent from one trial to the next contributes to the  
174 description of each unit’s time field. The findings were qualitatively similar to those from  
175 Figure 3. Again there was a significant correlation between time of peak firing and the  
176 width of the time field (Pearson’s correlation 0.41,  $p < .001$ ). As before, the distribution  
177 of time fields was better fit by a power law distribution than by a uniform distribution  
178 ( $\Delta\text{AIC}=20$ ,  $\Delta\text{BIC}=17$ ) and better fit by a power law than by an exponential distribution  
179 ( $\Delta\text{LL}=8$ ). The best fitting value for the exponent of the power law was  $-.39$ , close to the  
180 value ( $-.41$ ) found for the actual data. As in the actual data, the 95% confidence interval  
181 did not overlap with zero ( $-.34$  to  $-.43$ ).



*Figure 3.* mPFC Time fields show decreasing temporal accuracy for events further in the past. **A.** Ensemble similarity given through a cosine of the angle between normalized firing rate population vectors. The angle is computed at all pairs of time points during the delay period. The bins along the diagonal are necessarily one (warmest color). The similarity spreads out indicating that the representation changes more slowly later in the delay period than it does earlier in the delay period. **B.** Each row on the heatplot displays the firing rate (normalized to 1) for one time cell. White corresponds to high firing rate, while black corresponds to low firing rate. Vertical black lines mark the start and the end of the delay interval. The cells are sorted with respect to the median of the spike time in the delay interval. There are two features related to temporal accuracy that can be seen from examination of this figure. First, time fields later in the delay are more broad than time fields earlier in the delay. This can be seen as the widening of the central ridge as the peak moves to the right. In addition the peak times of the time cells were not evenly distributed across the delay, with later time periods represented by fewer cells than early time periods. This can be seen in the curvature of the central ridge; a uniform distribution of time fields would manifest as a straight line.

182 *Time fields could not be accounted for by observed behavioral correlates.* It is possible  
 183 that units that fire during circumscribed periods of time do so not because of time *per se*,  
 184 but because of some behavioral state that happens to occur at the same time during each  
 185 trial. For instance, perhaps the animal adopts a strategy of walking very slowly from one  
 186 side of the maze to the other at a constant velocity; the animal's location at the time that  
 187 the interval ends serves as a proxy for time since the interval began.

188 To determine whether the time cell findings were solely due to behavioral correlates,  
 189 we repeated the analyses considering only the units that did not show a significant behavioral  
 190 correlate. The behavioral parameters we had available were position along the x axis,  
 191 position along the y axis and movement speed. We divided each longest delay interval into  
 192 50 bins and computed the mean firing rate for each bin for all the intervals. Firing rate  
 193 of 48 out of 122 time cells was significantly correlated with at least one of the behavioral  
 194 parameters (Pearson's correlation coefficient with  $p < .01$ ). Instead of doing the analysis  
 195 on all 122 time cells we used only 74 behaviorally uncorrelated cells. The findings were  
 196 qualitatively similar to the results found for all 122 units classified as time cells. Even with  
 197 relatively low number of cells the time of peak firing and the width of the time field were  
 198 still correlated (Pearson's correlation 0.27,  $p = .018$ ). The distribution of time fields was  
 199 better fit by a power law distribution than by a uniform distribution ( $\Delta AIC=7$ ,  $\Delta BIC=5$ )

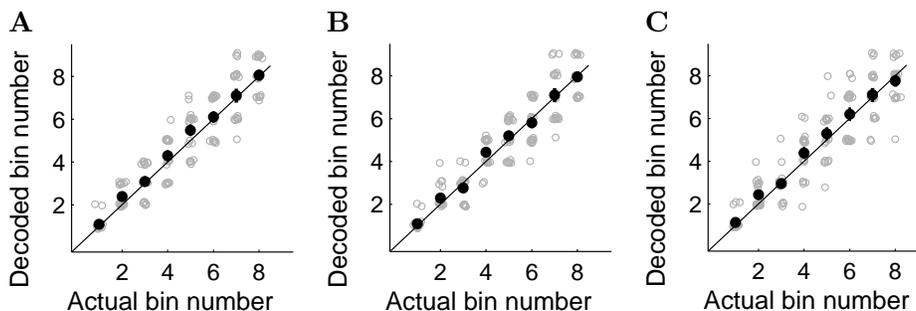


Figure 4. Population of mPFC time cells carried similar amount of temporal information as a same-size population of ramping cells. Decoded bin number versus actual bin number. Open gray circles denote the trial-by-trial decoding results for each bin. Filled black circles and error bars denote their means and SEM across trials. **A.** Temporal decoding based on all 723 reported units. Mean error: 0.71 bins. **B.** Temporal decoding based on all 122 time cells Mean error: 0.59 bins. **C.** Temporal decoding based on the randomly chosen 122 ramping cells. Mean error: 0.70 bins.

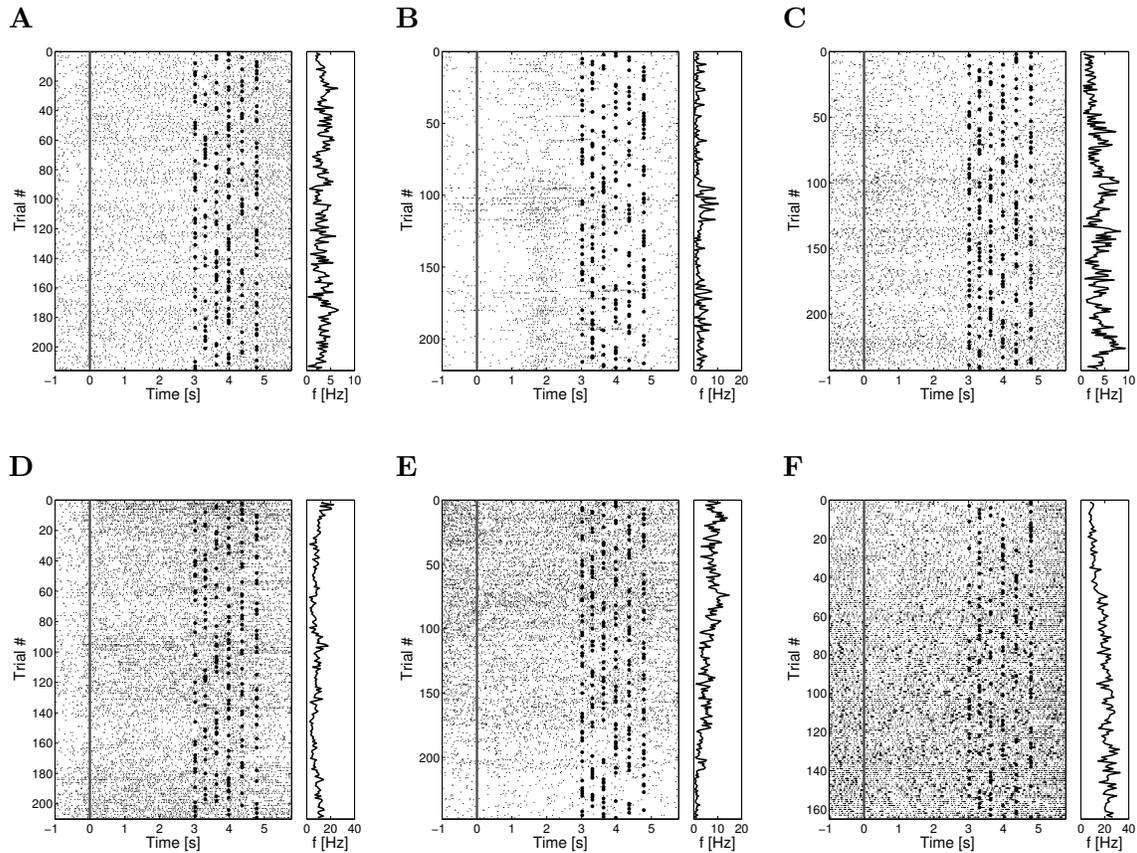
200 and slightly better fit by a power law than by an exponential distribution ( $\Delta LL=1.5$ ). The  
 201 best fitting value for the exponent of the power law was  $-0.29$ .

202 *Temporal variability in firing across minutes*

203 In addition to the reliable changes in the firing of time cells on the scale of seconds  
 204 within the delay interval, we also observed gradual changes in the firing properties of many  
 205 units that changed slowly across trials. Figure 5 shows representative examples. Note that  
 206 some units increased their firing transiently; others decreased or increased over the entire  
 207 session. Almost all of the units showed some evidence of autocorrelation across trials. Out  
 208 of 723 units, 561 showed a time constant of at least one trial. Somewhat reminiscent of  
 209 the distribution of peak times of the time cells, many more units had short time constants  
 210 than a long time constants. The distribution of time constants across units was described  
 211 well by a power law distribution (Figure 6). The power law fit was much more likely than  
 212 uniform ( $\Delta AIC > 1000$ ,  $\Delta BIC > 1000$ ) and exponential fit ( $\Delta LL = 119$ ). The exponent of  
 213 the best fitting power law distribution was  $-1.76$  with the 95% confidence interval defined  
 214 with exponents  $-1.65$  and  $-1.88$ .

215 *Across-trial variability was observed in a population that overlapped with within-trial*  
 216 *temporal coding.* Some units exhibited both within and across-trial gradual changes of the  
 217 firing rate. The distribution of across-trial time constants for cells classified as time cells  
 218 did not differ reliably from the distribution of across-trial time constants of all units (K-S  
 219 test statistic 0.0579).

220 *Across-trial variability could not be attributed to behavioral correlates.* We tested  
 221 whether the gradual changes in the neural activity are caused by any of the available  
 222 behavioral correlates. As in the earlier analysis on the time cells, behavioral correlates were  
 223 position along the x axis, position along the y axis and movement. Two pieces of evidence  
 224 argue against the hypothesis that the long time constants we observed were attributable to  
 225 behavioral correlates.



*Figure 5.* Examples of units that gradually changed their firing rate across trials. Each raster plot is aligned on the start of the waiting period of each trial (gray line). The end of the interval is marked by a large black dot. The plot on the right shows firing rate during the delay period as a function of trial number. The start of each trial was separated by approximately 20 s. The time constants of the six units were A: 280 s; B: 340 s; C: 380 s; D: 440 s; E: 740 s; F: 940 s.

226 First, the measured behavioral correlates were autocorrelated over much shorter time  
 227 scales than the neural data. Neural changes were quantified through a time constant derived  
 228 from the autocorrelation function of firing rate. Therefore, we computed an analogous  
 229 measure for the behavioral data. Distributions of the time constants were, for each of  
 230 the three behavioral correlates significantly different than the distribution coming from the  
 231 neural data (K-S test,  $p < 0.001$ ). Behavioral time constants were on average about five  
 232 times shorter than neural time constants.

233 Second, if behavior was causing the autocorrelation observed in the units, because  
 234 behavior is the same for all units recorded in the same session, we would expect to see  
 235 units from the same session to have time constants that are correlated with one another. In  
 236 contrast, if behavior was not a major factor in causing across-trial changes in firing, then  
 237 units from the same session would have the same statistics as units recorded from different  
 238 sessions. This hypothesized correlation in time constants should manifest as a change in the  
 239 distribution across sessions of mean time constants for units within a same session. To test

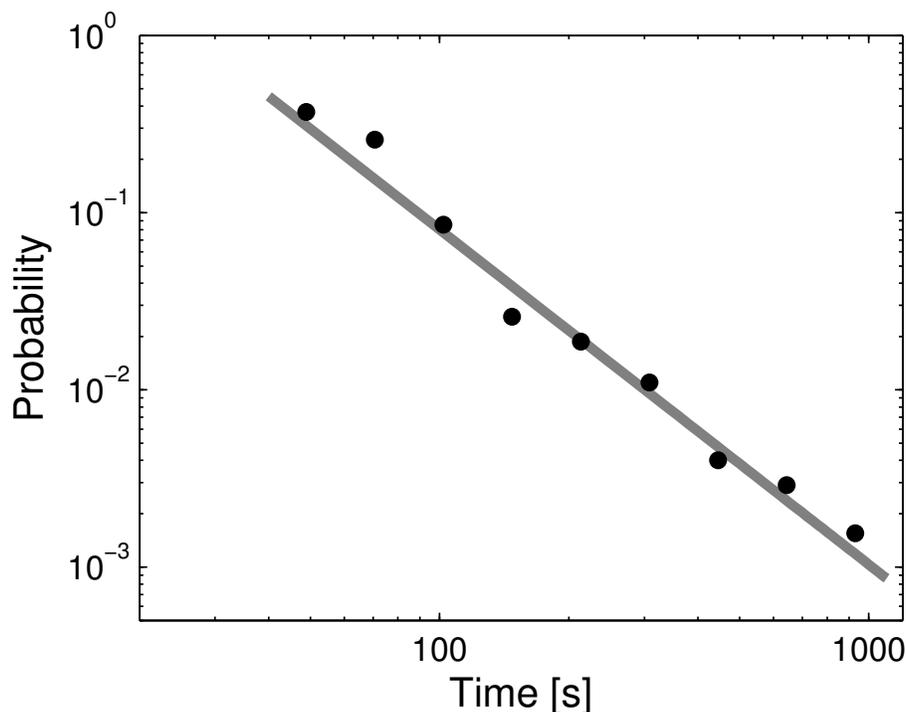


Figure 6. The distribution of time constants across units approximates a power law distribution. For each unit, a time constant of across-trial firing was estimated from its autocorrelation (see text for details). The time constant measured in number of trials was then multiplied by the average time between trials (20 s) in order to provide a sense of the scale of the fluctuations. The black dots show the probability density function of the data on log-log paper. The gray line gives the maximum likelihood power law fit. The exponent of the power law is -1.76.

240 this hypothesis, we computed F statistics from the time constants of all 722 units, treating  
 241 the session identity as a categorical variable. Since the time constants are not normally  
 242 distributed, to evaluate whether there is significant correlation between the time constants  
 243 and the sessions identities we shuffled the unit identity with respect to recording sessions  
 244 for 1000 time and computed F statistics for each shuffle. Rank of the observed data within  
 245 the shuffled data was 627, suggesting that units that were recorded in a same session were  
 246 not more likely to have a particular time constant.

247

### Discussion

248 This study shows that mPFC contains sequentially activated time cells, similar to  
 249 those previously reported in the hippocampus. The time fields of these units spanned the  
 250 entire 5 s delay interval, but with temporal accuracy that decreased as the delay elapsed.  
 251 The width of the time fields increased with temporal distance from the onset of the delay  
 252 period and distribution of the firing rate peaks strongly deviated from the uniform such  
 253 that more units represented time periods early in the delay rather than later in the delay.  
 254 Additionally, neurons in mPFC exhibited gradual changes in firing across trials spanning up  
 255 to at least tens of minutes. The number of units that exhibited a particular time constant

256 decreased as a power law function of the duration. Taken together, these results suggest  
257 that mPFC could be used for timing over a variety of time scales from a few hundred  
258 milliseconds up to tens of minutes.

259 *Could these findings be recording artifacts*

260 The results in this paper are consistent with, but do not uniquely specify, the hypoth-  
261 esis that firing of mPFC neurons maintain a temporal memory over a variety of time scales.  
262 One alternate possibility is that the temporally modulated firing reflect some other factor  
263 that also changes over time. Temporally-correlated behavior is one candidate; recording  
264 artifacts are another.

265 The behavioral measures that were measured in this experiment (x-position, y-  
266 position and running speed) were not sufficient to account for either the within-trial or  
267 the across-trial temporal modulation. However, this does not exclude the possibility that  
268 there are other behavioral factors that were not measured. For instance, it is possible that  
269 some animal's might have engaged in some subtle behavioral strategy within each trial, such  
270 as shifting weight or some pattern of whisking, that was not measured. Over the course  
271 of the session, we would expect the animals to get progressively less thirsty, or for body  
272 temperature to change due to exertion. However, we saw across-trial changes across a range  
273 of time scales, and cells that both increased and decreased their firing. As a result it is not  
274 likely that a single behavioral correlate could cause the gradual change across time scales.

275 There are a number of factors that could result in artifactual changes in spike-sorting  
276 over time on the scale of time within a trial and also across trials. For instance, when a  
277 neuron fires repeated action potentials over hundreds of milliseconds, the waveform might  
278 change. Alternatively, tetrodes might shift gradually over the recording session. We reduced  
279 the possibility that the results are influenced by recording artifacts by eliminating 10 units  
280 which average spike waveforms significantly changed during the recording, but there is no  
281 way to know with certainty that the results are not attributable to some recording artifact.  
282 However, similar findings have been observed with calcium imaging in the hippocampus,  
283 which would not be subject to the same set of recording artifacts. Modi et al. (2014) found  
284 time cells that fire during a circumscribed part of the delay period of a trace conditioning  
285 experiment. Ziv et al. (2013) showed that the hippocampal representation of place on a  
286 simple linear track changed gradually across days.

287 *Relationship to temporally-modulated firing in the hippocampus*

288 This paper reports that mPFC contained sequentially activated time cells with de-  
289 creasing temporal accuracy and cells that changed their firing gradually over long periods  
290 of time. Both of these phenomena have previously reported in the hippocampus. For  
291 instance, several studies have found evidence for hippocampal cells that fire during cir-  
292 cumscribed periods of time within a delay interval (Gill et al., 2011; Kraus et al., 2013;  
293 MacDonald et al., 2011; MacDonald et al., 2013; Modi et al., 2014; Naya and Suzuki, 2011;  
294 Pastalkova et al., 2008). Some of these studies have found evidence for decreasing temporal  
295 accuracy as a function of delay, due to spread in time field width (Howard et al., 2014;  
296 Kraus et al., 2013) or due to a non-uniform distribution of time field locations (Kraus  
297 et al., 2013). In addition, gradual changes in firing across minutes have been ob-  
298 served in the human (Howard et al., 2012) and rat hippocampus (Mankin et al., 2012;

299 Manns et al., 2007). However, these studies have characterized gradual change at the  
300 population level; it is not yet clear whether the hippocampus also shows a power law dis-  
301 tribution of time constants like we observed in the mPFC and, if so, whether the exponent  
302 corresponds.

303 It is also not clear in either the mPFC or the hippocampus whether the gradually-  
304 changing firing carries meaningful information about past events or not. This could be  
305 established (and recording artifacts definitively ruled out) if an experiment were to demon-  
306 strate control over gradually changing firing. For instance, the unit in Figure 5E decreases  
307 its firing around trial 80 and then decays gradually over about 50 trials, extending a few  
308 hundred seconds. Even if we were able to identify some unusual event that occurred around  
309 trial 90, this would not demonstrate causal control over the cell's firing. In order to do  
310 so, we would have to present the hypothetical stimulus multiple times, separated by a few  
311 hundred seconds and show that the stimulus consistently causes the same profile of firing.  
312 Examining recordings from monkeys, Bernacchia et al. (2011) showed that gradual changes  
313 in the firing of neurons in a variety of regions, including prefrontal cortex, reflected the  
314 history of reward, so it is at least possible in principle for the brain to maintain information  
315 about some past events over long periods of time.

#### 316 *Concluding remarks*

317 Previous work has shown that neural ensembles in the rodent mPFC code for time  
318 with decreasing temporal accuracy (Kim et al., 2013) and change gradually over long periods  
319 of time (Hyman et al., 2012). This paper extends these findings in two ways. First, a  
320 subpopulation of units in the mPFC fired like sequentially activated time cells, firing for  
321 circumscribed periods of time during the delay of an interval discrimination task. These  
322 time cells exhibited decreasing temporal accuracy in two ways. First, time cells that fired  
323 later in the delay interval had wider temporal receptive fields than time cells that fired  
324 earlier in the delay. Second, the distribution of time fields was not uniform. More cells  
325 had time fields earlier in the delay period than later in the delay period. In addition to  
326 these findings regarding firing correlates while timing delays on the order of seconds, we also  
327 observed gradual changes in firing rate over time scales up to a thousand seconds (Hyman  
328 et al., 2012). The gradual change across the population was attributable to units that  
329 showed autocorrelation at different time scales. Most units showed at least some significant  
330 autocorrelation across trials, which were separated by on average 20 s. A few units showed  
331 autocorrelations across the entire session, lasting tens of minutes. The distribution of time  
332 constants across units was well-described by a power law distribution. Taken together,  
333 these findings are consistent with the hypothesis that the mPFC is part of a system that  
334 represents time with decreasing accuracy over a range of time scales from a few hundred  
335 milliseconds up to thousands of seconds.

#### 336 References

- 337 Bernacchia A, Seo H, Lee D, Wang XJ (2011) A reservoir of time constants for memory  
338 traces in cortical neurons. *Nature Neuroscience* 14:366–72.
- 339 Buhusi CV, Meck WH (2005) What makes us tick? Functional and neural mechanisms of  
340 interval timing. *Nature Reviews Neuroscience* 6:755–65.

- 341 Gill PR, Mizumori SJY, Smith DM (2011) Hippocampal episode fields develop with learn-  
342 ing. *Hippocampus* 21:1240–9.
- 343 Howard MW, Viskontas IV, Shankar KH, Fried I (2012) Ensembles of human mtl neurons  
344 “jump back in time” in response to a repeated stimulus. *Hippocampus* 22:1833–1847.
- 345 Howard MW, MacDonald CJ, Tiganj Z, Shankar KH, Du Q, Hasselmo ME, Eichenbaum  
346 H (2014) A unified mathematical framework for coding time, space, and sequences in the  
347 hippocampal region. *Journal of Neuroscience* 34:4692–707.
- 348 Hyman JM, Ma L, Balaguer-Ballester E, Durstewitz D, Seamans JK (2012) Contextual  
349 encoding by ensembles of medial prefrontal cortex neurons. *Proceedings of the National*  
350 *Academy of Sciences USA* 109:5086–91.
- 351 Kim J, Jung AH, Byun J, Jo S, Jung MW (2009) Inactivation of medial prefrontal cortex  
352 impairs time interval discrimination in rats. *Frontiers in Behavioral Neuroscience* 3.
- 353 Kim J, Ghim JW, Lee JH, Jung MW (2013) Neural correlates of interval timing in rodent  
354 prefrontal cortex. *Journal of Neuroscience* 33:13834–47.
- 355 Kraus BJ, Robinson n RJ, White JA, Eichenbaum H, Hasselmo ME (2013) Hippocampal  
356 “time cells”: time versus path integration. *Neuron* 78:1090–101.
- 357 Lejeune H, Wearden JH (2006) Scalar properties in animal timing: conformity and viola-  
358 tions. *Quarterly Journal of Experimental Psychology* 59:1875–908.
- 359 Lewis PA, Miall RC (2009) The precision of temporal judgement: milliseconds, many  
360 minutes, and beyond. *Philosophical Transactions of the Royal Society London B: Biological*  
361 *Sciences* 364:1897–905.
- 362 MacDonald CJ, Carrow S, Place R, Eichenbaum H (2013) Distinct hippocampal  
363 time cell sequences represent odor memories immobilized rats. *Journal of Neuro-*  
364 *science* 33:14607–14616.
- 365 MacDonald CJ, Lepage KQ, Eden UT, Eichenbaum H (2011) Hippocampal “time cells”  
366 bridge the gap in memory for discontinuous events. *Neuron* 71:737–749.
- 367 Mangels JA, Ivry RB, Shimizu N (1998) Dissociable contributions of the prefrontal and  
368 neocerebellar cortex to time perception. *Cognitive Brain Research* 7.
- 369 Mankin EA, Sparks FT, Slayyeh B, Sutherland RJ, Leutgeb S, Leutgeb JK (2012) Neuronal  
370 code for extended time in the hippocampus. *Proceedings of the National Academy of*  
371 *Sciences* 109:19462–7.
- 372 Manns JR, Howard MW, Eichenbaum HB (2007) Gradual changes in hippocampal activity  
373 support remembering the order of events. *Neuron* 56:530–540.
- 374 Modi NM, Ashesh DK, Bhalla SU (2014) CA1 cell activity sequences emerge after reor-  
375 ganization of network correlation structure during associative learning. *eLife* 3.

- 376 Naya Y, Suzuki W (2011) Integrating what and when across the primate medial temporal  
377 lobe. *Science* 333:773–776.
- 378 Onoe H, M. K, Onoe K, Takechi H, Tsukada H, Watanabe Y (2001) Cortical networks  
379 recruited for time perception: A monkey positron emission tomography (pet) study. *Neu-*  
380 *roImage* 13:37 – 45.
- 381 Pastalkova E, Itskov V, Amarasingham A, Buzsaki G (2008) Internally generated cell  
382 assembly sequences in the rat hippocampus. *Science* 321:1322–7.
- 383 Wearden JH, Lejeune H (2008) Scalar properties in human timing: conformity and viola-  
384 tions. *Quarterly Journal of Experimental Psychology* 61:569–87.
- 385 Ziv Y, Burns LD, Cocker ED, Hamel EO, Ghosh KK, Kitch LJ, El Gamal A, Schnitzer  
386 MJ (2013) Long-term dynamics of cal hippocampal place codes. *Nature Neuro-*  
387 *science* 16:264–6.