Published in final edited form as:
Hippocampus. 2014 April ; 24(4): 476–492. doi:10.1002/hipo.22243.

Characterizing Context-Dependent Differential Firing Activity in the Hippocampus and Entorhinal Cortex

Michael J. Prerau[1,4], **Paul A. Lipton**[1,4], **Howard B. Eichenbaum**[1,2,4], and **Uri T. Eden**[1,3]

[1]Graduate Program in Neuroscience, Boston University, Boston, MA 02215

[2]Department of Psychology, Boston University, Boston, MA 02215

[3]Department of Mathematics and Statistics, Boston University, Boston, MA 02215

[4]Center for Memory and Brain, Boston University, Boston, MA 02215

Abstract

The rat hippocampus and entorhinal cortex have been shown to possess neurons with place fields that modulate their firing properties under different behavioral contexts. Such context-dependent changes in neural activity are commonly studied through electrophysiological experiments in which a rat performs a continuous spatial alternation task on a T-maze. Previous research has analyzed context-based differential firing during this task through the characterization of differences in the mean firing activity between left-turn and right-turn experimental trials. In this paper, we draw upon findings that demonstrate considerable trial-to-trial variability of neural activity in these regions and broaden the definition of differential firing to include context-dependent changes in the stochastic structure of the trial-to-trial rate variability. We develop qualitative and quantitative methods to characterize and compare changes in trial-to-trial variability in the CA1 region of hippocampus and in the dorsocaudal medial entorhinal cortex (dcMEC) between turn direction contexts during a spatial alternation task on a T-maze. We identify a subset of cells with context-dependent changes in firing rate variability. Additionally, we show that dcMEC populations encode turn direction uniformly throughout the T-maze stem, whereas CA1 populations encode context at major waypoints in the spatial trajectory. Our results suggest scenarios in which individual cells that sparsely provide information on turn direction might combine in the aggregate to produce a robust population encoding.

Keywords

Hippocampus; Entorhinal Cortex; Memory; Grid Cells; Place Cells

Introduction

The hippocampus (O'Keefe and Nadel, 1978) and entorhinal cortex (Mizumori et al., 1992, Quirk et al., 1992, Fyhn et al., 2004) have been shown to contain neurons that exhibit location-specific activity. Furthermore, it has been demonstrated that the hippocampus (Frank et al., 2000, Wood et al., 2000, Ferbinteanu and Shapiro, 2003, Ainge et al., 2007, Lipton et al., 2007) and entorhinal cortex (Frank et al., 2000, Lipton et al., 2007) encode information related to the disambiguation of overlapping spatial trajectories for rodents in behavioral tasks. Findings in Lipton et al. (2007) indicate that, for rats performing continuous spatial alternation on a T-maze, neurons in the dorsocaudal medial entorhinal

Corresponding Author: Michael J. Prerau, prerau@nmr.mgh.harvard.edu, Boston University Department of Math and Statistics, 111 Cummington Street, Boston, MA 02134, Tel: (646) 644-6453, Fax: (617) 353-8100.

cortex (dcMEC) more robustly differentiate overlapping trajectories from experimental trials with different turn directions, whereas hippocampal CA1 neurons more selectively encode location-related information.

Context-dependent neural activity, known also as differential firing, changes as a function of different experimental or behavioral parameters, termed *contexts*. Wood et al. (2000) and Frank (2000) initially identified certain CA1 neurons, termed “splitter cells”, that exhibited differences in mean firing activity between left-turn and right-turn contexts during continuous spatial alternation. Consequently, subsequent studies of context-dependent neural activity have focused on differences in the mean firing rate.

Spiking in the hippocampus has additionally been shown to be highly variable across experimental trials (Fenton and Muller, 1998). A biological stimulus or behavior has the potential to affect both the expected firing activity and its trial-to-trial variability. For example, in certain cortical neurons, the trial-to-trial spiking variability is significantly reduced at the onset of stimuli (Churchland et al., 2010). It is therefore important to develop analyses that effectively model the full firing rate distribution, rather than only its expected value, in order fully characterize neural coding properties.

Some previous approaches to characterizing trial-to-trial variability have been to model the history-dependent properties of neurons on small time scales with descriptive statistics (Churchland et al., 2010) or through parametric modeling (Brown et al., 2001, Brown et al., 2004, Eden et al., 2004, Truccolo et al., 2005), to model changes in network state, or “reference frame”, which occur on larger time scales (Touretzky and Redish, 1996, Redish and Touretzky, 1997, Touretzky and Muller, 2006) with doubly stochastic Poisson process models of switches between these network states (Lansky et al., 2001, Jackson and Redish, 2007), or to use a multiplicative model of stimulus specific and trial related variability components (Ventura et al., 2005).

In this paper, we broaden the definition of context-dependent activity to encompass changes in the firing rate distribution across experimental trials, in order to more completely characterize neural activity from the rat CA1 and dcMEC during a T-maze continuous spatial alternation task. We develop both qualitative and quantitative methods to visualize and characterize the stochastic structure of neural firing activity for a single set of trials, and we develop methods to identify differences between two sets of experimental trials. We use these methods to compare the spatial distribution of differential activity between CA1 and dcMEC. Finally, we identify a population of neurons in both regions that are each informative about context only on a small subset of trials.

Materials and Methods

Experimental Data and Preprocessing

Neural spiking data was acquired from the experiment described in detail in Lipton et al. (2007), in which eight Long-Evans rats were trained to perform a spatial alternation task on a modified T-maze. The rats were trained to continuously alternate between left and right turns on the T-maze, while extracellular neural activity was recorded in either CA1 or dcMEC. The rat’s spatial location on the T-maze also was recorded during the task. Generally, the term *context* is defined as a specific set of experimental or behavioral conditions. Neural activity that changes with respect to experimental context is called *context-dependent*, or is said to exhibit *differential firing*. In this paper, the two contexts of interest will be the left-turn and right-turn trials.

We used several exclusion criteria for data preprocessing. To account for differences in neural activity related to other behavioral correlates, we performed an ANOVA analysis, as in Lipton et al. (2007), on velocity and lateral stem position, with vertical stem position and turn direction as factors. For those experimental sessions in which the ANOVA found significant differences in these behaviors, we performed an ANCOVA on firing rate with the same factors. All cells with activity significantly correlated to either velocity or to lateral stem position were excluded from analysis.

To further account for lateral movement and changes in head-direction, only data from the central 90cm of the T-maze stem (the straight portion of track which precedes the choice point) were included in the analyses. Figure 1D shows the spatial trajectories from a full experimental session (gray curves) with the left-turn (red) and right-turn (blue) contexts highlighted within the 90cm region of analysis. This central portion was defined using the aggregate spatial trajectories and recorded videos of all animals, so as to construct a highly conservative range in order to maximize the number trials, independent of context, in which the animal was proceeding in a straight line in the center of the track with its head forward. Trials with highly erratic behavior within this region were excluded from analysis.

After removing cells using the aforementioned procedures, 321 cells were left to be used in this analysis: 111 cells from 10 sets of simultaneously recorded neurons from five rats with six tetrodes aimed at dorsal CA1, and 210 cells from 10 sets of simultaneously recorded neurons from three rats with 13 tetrodes aimed at dcMEC. Each of the 20 total data sets was recorded during a separate experimental session. The number of experimental trials per session ranged from 31 to 69, with an average of 46 trials per session.

It should be noted that whereas Lipton et al. (2007) only analyzed the activity from units with an overall mean firing rate >1 Hz and that were identified as having robust place fields on the T-maze stem, the current analysis included all isolated units, regardless of mean firing rate on the stem or any quantitative indicator of location-related firing. This increased sample size from 89 cells to 321 cells, and allowed us to more accurately estimate the true proportion of cells within each brain region with specific firing properties, regardless of place field prevalence or structure. Thus, the term *place field* is used in a general sense herein, defined as the firing activity as a function of position.

For each of the cells, firing rate was calculated using a kernel smoother, which convolved the spike train with 500ms hanning window (Parzen, 1962, Dayan and Abbott, 2001). The bandwidth parameter of the kernel smoother was determined using the likelihood-based framework for firing rate bandwidth estimation that is discussed in Prerau and Eden (2011). The distribution of bandwidth parameters was computed for a large cross-section of cells from both regions across many trials, which had a peak at close to 500ms, suggesting that it would be the bandwidth appropriate for the largest number of units.

ANOVA Analysis

For comparison purposes, we performed a two-way ANOVA on firing rate with turn direction and position on the stem of the T-maze as factors. Trials were grouped into left-turn trials and right-turn trials based on the direction that the animal turned at the end of the stem, and the 90cm central portion of the T-maze stem was divided into 7 equally sized spatial bins 12.85cm long and 5.7cm wide, as in Lipton et al. (2007). For each trial, the data value associated with each bin was computed as the total number of spikes fired while the animal was in that bin divided by the total amount of time spent in that bin. Using these two turn direction contexts and the seven spatial bins, a two-way ANOVA was performed on all 321 cells. A neuron was designated as exhibiting differential firing if there was significant

main effect of turn direction or if there was a significant turn direction by position interaction using a $p<0.025$ significance level to correct for multiplicity.

Analysis of Firing Rate Distribution Structure

Neural spiking activity in the hippocampus has been found to exhibit a high degree of trial-to-trial variability (Fenton and Muller, 1998). We therefore wish to employ methods of data analysis that go beyond a characterization of aggregate neural activity over many trials to capture the full trial-to-trial distribution of the firing activity. To study the statistical structure of trial-to-trial firing rate, we developed a non-parametric, data-driven description of the trial-to-trial variability. Many standard approaches, such as the ANOVA, make an implicit assumption that the data has a normal distribution, or else requires large numbers of spikes or rate estimates. These assumptions about the structure of the data are often not verified, and can lead to erroneous conclusions in subsequent higher-level analyses. Additionally, standard procedures often focus solely on the changes in expected firing rate and ignore the changes associated with other statistical features of firing rate.

Characterizing Firing Rate Distribution with Empirical Probability Surfaces—

Our analysis is based on the empirical distribution of the firing rate trajectories. Because empirical distributions are defined by the data, they do not require a priori assumptions about the stochastic structure of the data. The use of empirical distributions is therefore ideal for exploring trial-to-trial variation in neural firing.

Empirical distributions are defined by an empirical cumulative distribution function (CDF), $F(x)$ (DeGroot and Schervish, 2002), which describes the proportion of the data less than or equal to any value x. While an empirical probability density function (PDF) is not well defined, a model of the empirical density can provide a more intuitive visualization of the distribution structure of the estimated firing rates than the empirical CDF.

To visualize the density of neural firing rate for a specific region, we extract the segments of the firing rate trajectories across multiple trials that fall within that region as samples [Fig. 1A]. We then use the rate trajectory samples to compute a histogram [Fig. 1B], which is smoothed with a localized polynomial filter and then normalized [Fig. 1C]. The result is a data-based model of firing rate density, which provides a non-parametric description of how likely the neuron is to fire at any given rate within a particular fixed spatial interval. There are many well-known methods of density estimation (Turlach, 1993) that may also be appropriate, and we use this simple smoothing procedure as an initial means of visualizing the stochastic structure of neural firing rates so as to inform the development of subsequent quantitative analyses.

To capture the trial-to-trial dynamics of firing rate over regions of space or time, we compute an *empirical probability surface*. Figure 1E–G shows data from a neuron from dcMEC during repeated right-turn trials on the spatial alternation task. The empirical probability surface is constructed by computing empirical density models at 50 equal 1.8cm spatial intervals across the 90cm central portion of the T-maze stem [Fig. 1D]. Using this high spatial resolution, it is possible to estimate the probability of observing any given rate at any given position across the region.

The construction of an empirical probability surface is described in detail in APPENDIX A.

We can visualize the empirical probability surface in three dimensions [Fig. 1F] as a surface, or in two dimensions [Fig. 1G] as a color map. In the 2D visualization of the empirical probability surface [Fig. 1G, the x-axis shows the position on the T-maze, with the start box/T-maze base on the left and the choice point on the right. The rat runs from left to right. The

y-axis shows the firing rate in Hz. The color map indicates the estimated probability density at each location/rate pairing. Since sparsely firing neurons have distributions heavily weighted towards rates at 0Hz, the color map is non-linear at the higher densities, so as to allow for a clear visualization of the distribution tails.

Identifying Differential Firing Through Descriptive Statistics and Regions of Significance—Along with visualization methods, is it important to provide quantitative methods for characterizing firing rate distributions, and for comparing two firing rate distributions from different experimental contexts.

Given a single set of full firing rate trajectories (the firing rates for single trials as functions of the T-maze stem position), we used a set of descriptive statistics to characterize the firing rate distribution at each of 50 spatial bins on the T-maze stem. While computing the value of a given statistical estimate in each bin is straightforward, it is sometimes non-trivial to calculate the degree of uncertainty of the estimate at each bin. One reason for this is that the distribution of the firing rate trajectories is non-Gaussian and unknown, which makes the use of the standard error incorrect. Furthermore, we want global bounds that will reflect the confidence for a complete rate trajectory across the stem, rather than point-wise bounds that assume independence in firing rate between spatial regions. To overcome these issues, we draw upon the methods used in Fujisawa et al. (2008) for defining global confidence bounds on the mean of many firing rate trajectories, extending them to handle any descriptive statistic. Bootstrap resampling is a well-known procedure (Efron and Tibshirani, 1994, Wasserman, 2004) for computing confidence bounds for descriptive statistics based solely on the empirical data, and is used when the size of the data set is small or has unknown distributional properties.

To compute the magnitude and confidence of a given statistic across the T-maze stem from a set of N full rate trajectories observed during a particular experimental context, we sample with replacement N times from the set of trajectories. This new, resampled collection of rate trajectories is used as surrogate data from which to compute the given descriptive statistic at each of the 50 spatial bins. The set of values of the statistic from the resampled data set is called a *bootstrap sample*. We use this resampling procedure to generate a large number of bootstrap samples, in this case 500, which can be used to estimate the distributional properties of the statistic. From the collection of all the bootstrap samples, we can compute global confidence bounds (Fujisawa et al., 2008, Appendix B) at a particular confidence level α by finding the global confidence bounds that completely contain $(1-\alpha)$*100% of the bootstrap sample trajectories over all spatial locations on the stem.

The full bootstrap procedure for computing global confidence bounds is described in detail in APPENDIX B, Algorithm 1.

We next adapt this procedure to characterize differences between collections of rate trajectories observed from the two different experimental contexts. For a given statistic, two bootstrap samples are generated from the full rate trajectories observed across the T-maze stem: one from the left-turn context rate set, and one from the right-turn context rate set. By taking the difference of these two bootstrap samples, we compute an estimate of the difference statistic between the left-turn and right-turn contexts. We perform this difference procedure a large number (500) of times—at each iteration computing the difference statistic for the pair of bootstrap samples. From these difference trajectories, we form an estimate of the distribution on the difference statistic. As in the single statistic case, we then compute global confidence bounds on the difference statistic distribution at a particular confidence level α. If the global confidence bounds contain zero, we can be confident at level α that there is no difference between the value of the statistic for left-turn and right-turn contexts.

If the global confidence bounds exclude zero, then we can say that the statistic differs between the left-turn and right-turn contexts at that location with confidence α.

We define a *region of significance* for a given sample statistic at confidence level α as a contiguous positional interval during which the global confidence bounds of the difference statistic exclude zero. Thus, regions of significance define the set of locations on the T-maze stem at which there is a context-dependent difference for a given statistic. Regions of significance are useful because they provide a measure of significance of the context-dependent difference in firing activity for any statistic, and indicate where these differences occur spatially.

The full procedure involved in computing regions of significance is described in detail in APPENDIX B, Algorithm 2.

Similarly, we can construct hypothesis tests for comparing the distribution of firing activity between contexts based on a resampling technique. We test the null hypothesis that the activity from all trials come from the same firing rate distribution, regardless of context, by pooling trials from both contexts and randomly assigning a context to each. We then compare the observed trajectory for a selected statistic to the global bounds computed from the distribution of trajectories generated by this resampling procedure, and reject the null hypothesis if the observed trajectory ever leaves these bounds.

Determining Salient Statistics in the Characterization of Differential Firing—In order to capture the trial-to-trial variability described in previous studies, we performed an analysis of three common descriptive statistics for each of the 321 cells studied,—mean, variance, and 95th percentile (p95)—at all 50 spatial intervals of the T-maze stem.

Across all of the cells, we computed the frequency with which a region of significance occurred for each of the three statistics to assess how informative a given statistic can be in describing differential firing in these neural systems. We flagged a cell as being a candidate for differential firing if any regions of significance existed for one or more of the statistics applied to the cell. It is important to note that the goal of this analysis was not to establish a rigorous hypothesis testing regime for differential firing given regions of significance from multiple statistics, but rather to develop an exploratory measure for characterizing and comparing two sets of related spike trains, which we apply to understanding the differences in the statistical structure of neural firing between experimental conditions.

The Spatial Distribution of Differential Firing—Regions of significance describe the magnitude and significance of the differences in a given statistic between two contexts as well as the location on the T-maze stem at which these differences occur. We can therefore use the location of the regions of significance to gain an understanding of the spatial distribution of the location at which differential firing occurs on the stem. For the set of 111 CA1 cells and the 210 dcMEC cells, we calculated the regions of significance for the mean and for the 95th percentile using the aforementioned procedures. For each of 25 equal 3.6cm intervals into which we divide the T-maze stem, we compute the proportion of cells in CA1 and dcMEC with regions of significance for the mean and 95th percentile that contain that bin. A larger bin size is used in this analysis, as the size of the data set is not sufficient to calculate the proportion with reasonable confidence at higher resolutions.

To determine whether the spatial distribution of differential firing differed between regions, we performed a maximum likelihood ratio test on the proportion of cells with a region of significance in each spatial bin. The null hypothesis for the test specified that the probability that a bin is part of a region of significance is equal across both brain regions. The likelihood

ratio statistic in this case is −2 times the log of a ratio of Bernoulli likelihoods. For likelihood in the numerator, the probability of being in a region of significance depends only on the spatial bin, while for the likelihood in the denominator, the probability depends on both spatial bin and brain region. Under the null hypothesis, this statistic is asymptotically chi-squared distributed with number of degrees of freedom equal to the number of spatial bins.

Finally, for each brain region, we constructed an additional maximum likelihood ratio test to examine the null hypothesis that the probability of any bin being part of a region of significance is constant over the entire stem. In this case, the likelihood in the numerator uses a constant probability for each bin, while the likelihood in the denominator has a probability that can change between spatial bins. Under the null hypothesis, this test statistic is also asymptotically chi-squared distributed with number of degrees of freedom equal to the number of spatial bins minus 1.

Results

We analyzed spiking data from cells in the CA1 region of hippocampus and the dorsocaudal entorhinal cortex while a rat performed a continuous spatial alternation task on a T-maze. We began by constructing empirical probability surfaces as described in the methods to visualize the distribution of firing rates over trials for movements in each direction. For each cell, we then computed regions of significant differential firing between contexts based on the mean firing rate, the variance from trial-to-trial, and the 95^{th} percentile (p95) of firing rate.

Qualitative Characterization of Differential Firing Using Empirical Rate Data

Characterizing Differential Firing Using Empirical Probability Surfaces—Figure 2 shows an example of an empirical probability surface from a cell from dcMEC during the spatial alternation task. This cell exhibited significant differential firing by a two way ANOVA based on two contexts (left vs. right trials) by 7 spatial bins, and is characteristic of the “splitter cells” identified in (Wood et al., 2000). The most elementary method of visualizing differential firing is to plot a raster of spikes as a function of position and trial [Fig. 2A]. Figure 2B shows estimated firing rate trajectories over multiple trials for both left-turn (left panel) and right-turn (right panel) contexts. From these rate trajectories, the aggregate place field is computed as the mean firing rate across trials as a function of position [Fig. 2B, black curves]. It should be noted that since the spike rasters and rate trajectories are plotted as a function of position, trials in which the rat does not proceed monotonically towards the choice point could show rate traces that double back. Additionally, due to the temporal continuity constraint imposed by the kernel smoother, if a cell fired just prior to the rat entering the center of the stem, the rate trace could carry over into the region. Thus, there can be trials with non-zero initial rates even if there are no spikes shown.

By calculating empirical probability surfaces for the left-turn trial [Fig. 2C, left panel] and right-turn trial [Fig. 2C, right panel] contexts, we can visualize the full structure of the firing rate distribution from trial to trial, including the difference in the mean firing rate identified by the ANOVA analysis. The empirical probability surface also provides insight into firing variability, modal structure, and other stochastic features. In this example, the firing rate on the first half of the stem is more likely to be high on right turns, but there is substantial variability in the firing rate from one trial to the next.

Qualitative Modes of Differential Firing—To explore how neurons in CA1 and dcMEC encode contextual information, we computed empirical probability surfaces for both

the left-turn and right-turn contexts for each of the 321 cells. Figure 3 shows left-turn (left panels) and right-turn (right panels) context empirical probability surfaces for four neurons, each representative of different qualitative forms of differential firing.

Perhaps the most pronounced form of differential firing can be seen in characteristic "splitter" cells, like the neuron from Figure 2, which exhibit one-sided preferential firing. These cells, like those first identified in Wood et al. (2000), fire predominantly during one context and have significantly reduced or no firing in the other context. A neuron from dcMEC exhibiting one-sided preferential firing is shown in Figure 3. The empirical probability surface for the left-turn context [Fig. 3A, left panel] has a major band of high probability density at 0Hz throughout the entire stem, indicating that this neuron rarely fires before the rat takes a left-turn. The empirical probability surface for the right-turn context [Fig. 3A, right panel] has a well-defined place field with an enveloping variance that peaks just around the middle of the stem. This empirical probability surface still has a strong density band around 0Hz, indicating that while there are right-turn trials at which this neuron fires strongly, there are also right-turn trials at which this neuron does not fire.

Other cells exhibited contextually-dependent field-shifting, in which a neuron has clear place fields in both contexts, but the field is scaled in magnitude or translated in position in one context relative to the other. The neuron from CA1 shown in Figure 3B has well defined place fields for both turn contexts with both positional and scale differences between them. The left-turn surface peaks [Fig. 3B, left panel] around 40cm into the stem region, with the location of the place field ranging from 0 to 60cm. The right-turn surface [Fig. 3B, left panel] is shifted up the stem towards the start box, peaking around 20cm. The place field also is much broader, with a tail stretching throughout the entire stem. Thus, this cell fires later and for a shorter time on left-turn trials than on right-turn trials. Because of its very low spatial resolution, the ANOVA procedure failed to reject the null hypothesis that there was no difference between left-turn and right-turn contexts for this cell.

An additional class of cells was observed with similar mean field structures across contexts but differing trial-to-trial variation structure. Figure 3C shows a cell from CA1 exhibiting a context-modulated trial-to-trial variability. This neuron has an empirical probability surface with a high variability for the left-turn context [Fig. 3C, left panel], but has an empirical probability surface with significantly less variability for the right-turn context [Fig. 3C, right panel].

Figure 3D shows differential firing in an interneuron, which persistently fires everywhere on the stem. These cells are commonly excluded from analyses of differential firing, yet this cell shows marked contextual differences in field structure. The left-turn context has an arcing field with a single pronounced peak around 35cm into the stem region [Fig. 3D, left panel], whereas the right-turn context has a multimodal structure, with a small arc that peaks around 15cm and descends in a larger arc that troughs around 60cm [Fig. 3D, right panel].

Statistical Analysis of Differential Firing in dcMEC and CA1

The results of the qualitative analyses suggest that there are a wide variety of ways in which neurons may exhibit context-dependent changes in neural activity. Only a relatively small number of cells appeared to be clear-cut splitters or were easy to categorize visually. We observed, however, a large number of cells that exhibited differential firing involving multi-faceted changes in the trial-to-trial firing rate distribution structure. To quantify these changes, we computed the significance of differences in the firing rate distribution based on three sample statistics—the mean, variance, and 95^{th} percentile—for each of the 321 cells using the bootstrap procedure described in APPENDIX B.

Figure 4 illustrates the bootstrap procedure for a CA1 neuron with a context dependent change in trial-to-trial variability. Figure 4A and B show the empirical probability surfaces for left-turn and right-turn contexts, respectively. We computed bootstrap estimates of the firing rate mean [Fig. 4C, top panel] as well as the firing rate variance [Fig. 4D, top panel] for the left-turn (light blue curves) and right-turn (light red curves) contexts across the stem. The 95% global confidence intervals for the sample statistics are shown by the dark red and blue curves. We then computed the bootstrap estimates of the difference statistics (light green curves) for the firing rate mean [Fig. 4C, bottom panel] and for the firing rate variance [Fig. 4D, bottom panel] as a function of position. From the distribution of the difference statistic, we calculated 95% global confidence bounds (dark green curves), from which we could compute regions of significance.

Figure 4c shows that the values of the left and right estimates of the firing rate mean (top panel) are nearly identical, with overlapping confidence bounds at every point on the stem. The between context similarity in mean is reflected by the bounds of the difference statistic (bottom panel) always including zero (black line). These results indicate that at every point in the T-maze stem, we cannot exclude the possibility that the mean firing rate is the same in the left and right contexts with 95% confidence.

Figure 4D shows that, unlike the mean, the distributions of the firing rate variance as a function of position look quite different for each context. The variance for the left-turn context [Fig. 4D, top panel, blue curve] initially rises from the start box, then stays relatively constant through the middle of the stem, in contrast with the right-turn context variance [Fig. 4D, top panel, red], which rises to a much higher peak and then slowly declines towards the choice point. The distribution of difference statistic [Fig. 4D, bottom panel, green region] reveals that the 95% confidence intervals are significantly different from zero (black line) for a region in the center of the stem [Fig. 4A, B, and D, dashed vertical bars]. These differences in variance suggest that this cell has a place field that retains its overall structure across most all of the trials, and there is an increased variability for a subset of trials, all of which occur during the same behavioral context.

We performed this analysis for the mean, variance, and 95^{th} percentile on every cell in our dataset, and compared the results of the bootstrap analysis with the standard ANOVA, which we summarize in Figure 5. The ANOVA found 27% (56/210) and 12% (13/111) of the cells significant for dcMEC and CA1, respectively. As in the analysis in Lipton et al. (2007), the ANOVA found a significantly higher proportion of cells with differential firing in dcMEC than in CA1 (z-test, $p<.001$). For each of the three bootstrapped statistics, we found that there were many cells that showed statistically significant differences in firing that were not detected by the ANOVA. For both dcMEC and CA1, the 95^{th} percentile was the most salient statistic (dcMEC: 45% (95/210), CA1: 32% (36/111)), the proportion of which was also significantly different between both regions (z-test, $p<.05$). This was followed by the variance (dcMEC: 36% (75/210), CA1: 28% (31/111)), followed by the mean (dcMEC: 29% (61/210), CA1: 24% (27/111)), both of which did not show significant differences between regions. These results are surprising, as previous studies of differential firing focus exclusively on finding context-dependent differences in mean firing rate.

Of the cells determined to show differential firing by the bootstrap procedure in both regions, many of them were not found to be significant by the ANOVA. In dcMEC, 36% (51/95), 52% (39/75), and 53% (22/61), of the cells found by the bootstrap 95^{th} percentile, variance, and mean, respectively, were not found by the ANOVA. In CA1, 75% (27/36), 71% (22/31), and 70% (19/27), of the cells found the mean, variance, and 95^{th} percentile, respectively, were not found by the ANOVA (left panel, light gray stacked bars). Conversely, there were only 5 cells in CA1 and 4 in dcMEC found by the ANOVA alone.

The fact that 95th percentile and variance provide substantial power for identifying context dependent firing indicates that differential firing may be occurring in other ways besides a shift in the mean firing rate. Both statistics provide information about the spread of a distribution, and thus indicate that context-dependent differences in firing may be manifested through changes in the structure of the trial-to-trial variability.

The Spatial Distribution of Mean Rate Differential Firing on the T-Maze Stem

By computing regions of significance for a given statistic, we can express both the confidence of a statistical difference, as well the spatial location on the T-maze stem at which it occurred.

For the mean, variance, and 95th percentile statistics, we computed a maximum likelihood ratio test using the null hypothesis that the spatial distribution of proportions was the same for each brain region. We found statistically significant differences in the spatial distributions for all three statistics (mean: $p<.01$, variance: $p<.05$, p95: $p<.01$). This strongly indicates that the spatial distributions of differential firing in dcMEC and CA1 are different across dimensions of the firing rate distribution other than the mean.

Figure 6 shows the spatial distribution of the locations in which cells in dcMEC (left panels) and CA1 (right panels) exhibited significant differences between left-turn and right-turn trial contexts for the mean, and 95th percentile statistics. The proportions of cells with regions of significance at each region are mapped to T-maze schematics, with increasing brightness corresponding a larger proportion of cells with differential firing at that location.

To examine whether contextual firing differences occur uniformly across the stem, we performed another maximum likelihood ratio test for each brain region separately, using the null hypothesis that of the proportion of cells with differential firing was uniform across the entire stem. For dcMEC, the p-values for each statistic were close to 1 (mean: $p=.99$, variance: $p=.95$, p95: $p=.99$) indicating that there is not sufficient evidence to reject the null hypothesis of uniformity across position. For CA1, the likelihood ratio statistic did not achieve significance for the mean and variance statistics (mean: $p=.36$, variance: $p=.25$). For the 95th percentile, however, the likelihood ratio was able to reject the null ($p<.01$), suggesting that the spatial distribution for the trials with the highest firing rates is non-uniform in CA1.

By examining the spatial distributions of differential firing for the 95th percentile [Fig. 6B], it is evident that the proportions in dcMEC appear to be roughly the same throughout the entirety of the T-maze. In CA1, however, the spatial distribution of differential firing appears to have the highest proportion of firing at the very beginning of the stem, followed an area of decreased differential firing in the center portion of the T-maze stem.

These results imply that dcMEC may encode context uniformly, independent of spatial landmarks, whereas CA1 may encode context near important landmarks, such as the beginning of the stem.

Context-Modulated Variability

Through our analyses, we found many cells that fired identically during both contexts, except for a few trials with elevated firing rates that only occurred during a single context. To explore this phenomenon, we examined our largest population, consisting of 37 simultaneously recorded cells from dcMEC. For each cell, we performed a hypothesis test with a null hypothesis that the firing rate distribution is equivalent in both contexts, by computing the difference in the 95th percentile statistic for the firing rate of the observed data, and comparing to the distribution under the null hypothesis obtained by randomizing

the left and right context labels. This analysis identified a subset of 8 cells (22% of the population), which exhibited context-dependent differences in the 95th percentile of the firing rate distribution. To further account for behavioral anomalies that might produce high firing rates, all trials in which the rat had erratic behavior or slowed down below a linear velocity of 50px/s (~30cm/s) for 0.5s or more were removed from this analysis.

Figure 7A shows the firing rates for these 8 cells during the left-turn (left panels) and right-turn (right panels) contexts along with the corresponding p-values from the hypothesis test. For many of these significant cells, the firing rate distributions are very similar across both contexts, except for a small subset of trials that have much higher firing rates than all the others. To explore how these trials might function to sparsely encode context in the population, we used the maximum firing rate of the context with the lower rates as a threshold to detect high firing rate trials (black curves) in the context with the higher rates. Figure 7B illustrates for which trials the firing rates were above threshold (black) and below threshold (white) for each cell. All excluded trials are shown in gray.

While these cells averaged 5.5 (min: 2, max: 8) trials above threshold each, the subpopulation as a whole covered 63% (23/37) of the total observed trials. For each cell pair, we computed the correlation coefficient between the binary sequences, indicating whether the cell was above threshold for each trial. We found that 72% (20/28) of the cell pairs were not significantly correlated, and that no two cells had correlation coefficients above 0.54. This shows that different neurons in the ensemble are informative about context in different subsets of the trials.

Discussion

In this paper, we broaden the definition of context-dependent differential firing to encompass changes in the structure of the trial-to-trial firing rate distribution. We develop quantitative and qualitative methodologies to characterize and compare turn-direction selectivity in CA1 and in dcMEC. We identified cells in both regions in which firing variability structure and firing rate distribution tails were modulated by behavioral context in a statistically significant manner, many of which were not be identified by previous mean-based analyses. We found a greater prevalence of cells with differences in the probability distribution of high rate trials in dcMEC over CA1. We also found that the spatial distribution of the representation of turn direction differs significantly between the two regions, such that turn-direction is encoded primarily at the beginning of the T-maze stem in CA1, and uniformly across the stem in dcMEC. In cells that had context-modulated firing rate distributions, we showed that, for each cell, a small subset of the highest firing rates reliably occurred during one context. Furthermore, we showed that in the population, these high firing rates were well-distributed across experimental trials.

Differential Firing in the Hippocampus and Entorhinal Cortex Can Involve Context-Dependent Changes in The Firing Rate Distribution

We have illustrated that the stochastic structure of the trial-to-trial variability plays an important role in context-dependent differential firing. Churchland et al. (2010) found that the trial-to-trial variability in the spiking activity of cortical cells was significantly reduced at the onset of stimuli. These results indicate that certain cells within the hippocampus and entorhinal cortex also have trial-to-trial variability that is modulated by behavioral context.

Several potential sources of this variability have been suggested. Spiking is governed by intrinsic history dependent factors, such as refractoriness, which render it non-Poisson (Pfeiffer and Kiang, 1965, Tuckwell, 1988, Fenton and Muller, 1998, Koch and Segev, 1998, Shadlen and Newsome, 1998, Brown et al., 2004, Truccolo et al., 2005). Thus,

estimates of instantaneous firing rate are improved by modeling conditional intensity as a function of history dependency (Brown et al., 2001, Brown et al., 2004, Eden et al., 2004, Truccolo et al., 2005). Another theory suggests that trial-to-trial variability is a result of changes in network state or "reference frame" (Touretzky and Redish, 1996, Redish and Touretzky, 1997, Touretzky and Muller, 2006). Two-state doubly stochastic Poisson process models of reference frame switching (Lansky et al., 2001, Jackson and Redish, 2007) have been shown to reduce firing rate variability.

While some of the observed "excess variability" (Fenton and Muller, 1998) is expected due to model misspecification from the Poisson assumption, it is clear that this super-Poisson trial-to-trial variability surpasses history dependent effects (Muller, 1996). Hence, computing the full empirical distribution of the firing rate naturally extends the two-state model, while providing a more complete description of variability.

The Spatial Distribution of Context-dependent Differential Firing on the T-Maze Stem

We found that the population proportion of cells with context-modulated variability differs significantly between CA1 and dcMEC. Whereas cells in dcMEC were equally likely to show context dependence at any location along the track, cells in CA1 were more likely to be context dependent at locations at which turning decisions are initiated such as the start and end of the stem. This is in line with Lipton et al. (2007), which showed a high spatial specificity in CA1 and a strong encoding of turn direction in dcMEC. The spatial localization of turn direction encoding in CA1 also parallels other findings in which a large proportion of CA1 cells have been shown to fire at particular times/regions related to critical events in the task (Pastalkova et al., 2008, Ferbinteanu and Shapiro, 2003, Robitsek et. al, 2012).

While the hippocampus is not necessary for a rat to learn or perform a T-maze continuous alternation task, it is vital when a delay is introduced between trials (Ainge et al., 2007). Kim and Frank (2009) reported that rats with hippocampal lesions either did not learn or could not rapidly learn a W-maze continuous spatial alternation task. Together, these findings suggest the turn-direction dependent differential firing observed in the hippocampus may involved in maintaining task-related information over long periods of time, or related to the learning of more difficult tasks. The contrasting spatial distributions further suggest different roles for CA1 and dcMEC in these tasks. It is possible that CA1 may be reinforcing the animal's decisions at major way-points of the trajectory, while dcMEC helps to maintain the overall trajectory goal throughout the entire path (Hasselmo and Eichenbaum, 2005).

Differences in Firing Rate Variability Could Indicate Sparse Encoding of Context

While it is not possible to detect changes in variance based on the firing rate for a single trial, our results suggest that cells with context-dependent changes in firing rate distribution, like those described in Figures 3C, 4, and 7, could potentially be encoding context in a small set of trials with relatively high firing rates. Thus, these cells with context-dependent variability can also be thought of having transient trial-to-trial differences in field structure, which occur exclusively during a single behavioral context. These cells provide contextual information on only a subset of trials, and thus only exhibit a sparse context-dependence. This contrasts with the notion of fixed context-dependent differences in firing activity (Frank et al., 2000, Wood et al., 2000, Ferbinteanu and Shapiro, 2003, Ainge et al., 2007, Lipton et al., 2007). Upon further examination, the classic "splitter cells" that we identified with an ANOVA procedure are also sparsely context dependent. Rather than fire consistently in one context and not in the other, units like Cell 3 and Cell 6 in Figure 7, have

numerous trials with no spiking activity even during the context with the high mean firing rate.

The pairwise correlation analysis shown in Figure 7B suggests that a population of cells that sparsely encode context over a small subset of trials together would robustly provide contextual information throughout the entire experimental session. These cells may be recruited for short time spans based on cognitive demand, which may account for the way-point based spatial distribution of differential firing in CA1.

The idea of sparse encoding fits well within the multiple-map hypothesis (Touretzky and Redish, 1996, Redish and Touretzky, 1997, Touretzky and Muller, 2006), in which the variability could stem from context-specific maps in which there is only a partial remapping. Alternatively, sparse encoding could arise from interactions between cells that robustly encode context and those that encode position. These cells maintain a clear field structure across all trials, but produce the highest firing rates during only one context (e.g. Cells 2, 3, 4, 5, 7 and 8 in Figure 7).

There were more cells in dcMEC with significant context-dependent differences in the 95^{th} percentile than in CA1. These results suggest that tail-based differential firing is more prevalent in dcMEC than in CA1. This higher level processing is consistent with the confluence of regional inputs to the entorhinal cortex. The higher proportion of cells differentially firing in dcMEC is also consistent with the previous analysis in (Lipton et al., 2007), which we confirm with our ANOVA analysis. A wide range of proportions have been reported for the percentage of turn-direction selective cells from CA1 with well-defined place fields on the central stem, ranging from 33% (16/48) (Lipton et al., 2007) to 68% (53/78) (Lee et al., 2006). This large range may be attributable to differences in exclusion criteria, training methods, maze dimensions, or other experimental techniques. Moreover, such differences might also arise from different choices of large spatial bin sizes required in the ANOVA procedures to make the assumption of normality, as well as from the small number of cells used in these analyses. However, beyond the ability to analyze the proportion of cells within regions within and across regions that exhibit differential firing with greater statistical power and spatial resolution than the ANOVA, the most significant benefit of the new approach is that it provides us with the ability to characterize features of the neural coding of context that involve changes in the trial-to-trial variability, which would have otherwise gone unidentified.

Future Directions

In this paper we used an analysis of trial-to-trial variability to gain insight into context-dependent differential firing in the hippocampus and entorhinal cortex. In doing so, we have developed techniques, which can be generalized to characterize and compare the trial-to-trial variability for any spatiotemporally evolving collection of neural data, for any number of contexts.

Further study is needed to fully understand sparse encoding activity in the hippocampus and entorhinal cortex. In future work, decoding algorithms to predict context from ensemble activity may be used to examine the information content within the population versus individual cells. Additionally, improvements can be made to the methods, including the use of more sophisticated density estimation techniques for visualization (Turlach, 1993), a history dependent conditional intensity model of rate, or the development of state space models of the empirical probability surface (Frank et al., 2002, Czanner et al., 2008) to estimate the proper spatiotemporal continuity constraints.

Moreover, the questions brought to light through these new methods beg further investigation. The question of the temporal correlation of the tail-based differential activity could be assessed by developing statistical models of differential firing that explicitly model trial-to-trial history dependence. Additionally, the behavioral contributions to spatial differences in differential firing could be explored by computing firing rate in a statistical framework in which the effects of behavioral trajectory are explicitly modeled. The empirical probability analyses could then be performed on the component of the firing rate that was related only to the forward motion on the center part of the stem. Finally, future modeling and data analyses may allow us to speculate as to how downstream brain areas might use activity from populations of sparsely-encoding neurons to efficiently process and relay information about external stimuli and planned behaviors.

Acknowledgments

We thank M. Hasselmo, J. White, R. J. Robitsek, and C. MacDonald for helpful discussions and comments.

Grants

This work was supported by the NIMH grant MH71702 and by the National Science Foundation under grant number IIS-0643995.

References

Ainge JA, van der Meer MAA, Langston RF, Wood ER. Exploring the role of context-dependent hippocampal activity in spatial alternation behavior. Hippocampus. 2007; 17:988–1002. [PubMed: 17554771]

Brown EN, Kass RE, Mitra PP. Multiple neural spike train data analysis: state-of-the-art and future challenges. Nat Neurosci. 2004; 7:456–461. [PubMed: 15114358]

Brown EN, Nguyen DP, Frank LM, Wilson MA, Solo V. An analysis of neural receptive field plasticity by point process adaptive filtering. Proceedings of the National Academy of Sciences of the United States of America. 2001; 98:12261–12266. [PubMed: 11593043]

Churchland MM, Yu BM, Cunningham JP, Sugrue LP, Cohen MR, Corrado GS, Newsome WT, Clark AM, Hosseini P, Scott BB, Bradley DC, Smith MA, Kohn A, Movshon JA, Armstrong KM, Moore T, Chang SW, Snyder LH, Lisberger SG, Priebe NJ, Finn IM, Ferster D, Ryu SI, Santhanam G, Sahani M, Shenoy KV. Stimulus onset quenches neural variability: a widespread cortical phenomenon. Nat Neurosci. 2010; 13:369–378. [PubMed: 20173745]

Czanner G, Eden UT, Wirth S, Yanike M, Suzuki WA, Brown EN. Analysis of between-trial and within-trial neural spiking dynamics. Journal of Neurophysiology. 2008; 99:2672–2693. [PubMed: 18216233]

Dayan, P.; Abbott, LF. Theoretical neuroscience : computational and mathematical modeling of neural systems. Cambridge, Mass: MIT Press; 2001.

DeGroot, MH.; Schervish, MJ. Probability and statistics. Boston: Addison-Wesley; 2002.

Eden UT, Frank LM, Barbieri R, Solo V, Brown EN. Dynamic analysis of neural encoding by point process adaptive filtering. Neural Computation. 2004; 16:971–998. [PubMed: 15070506]

Efron B, Gong G. A Leisurely Look at the Bootstrap, the Jackknife, and Cross-Validation. Am Stat. 1983; 37:36–48.

Efron, B.; Tibshirani, R. An introduction to the bootstrap. New York: Chapman & Hall; 1994.

Fenton AA, Muller RU. Place cell discharge is extremely variable during individual passes of the rat through the firing field. Proceedings of the National Academy of Sciences of the United States of America. 1998; 95:3182–3187. [PubMed: 9501237]

Ferbinteanu J, Shapiro ML. Prospective and retrospective memory coding in the hippocampus. Neuron. 2003; 40:1227–1239. [PubMed: 14687555]

Frank LM, Brown EN, Wilson M. Trajectory encoding in the hippocampus and entorhinal cortex. Neuron. 2000; 27:169–178. [PubMed: 10939340]

Frank LM, Eden UT, Solo V, Wilson MA, Brown EN. Contrasting patterns of receptive field plasticity in the hippocampus and the entorhinal cortex: An adaptive filtering approach. Journal of Neuroscience. 2002; 22:3817–3830. [PubMed: 11978857]

Fujisawa S, Amarasingham A, Harrison MT, Buzsaki G. Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. Nat Neurosci. 2008; 11:823–833. [PubMed: 18516033]

Fyhn M, Molden S, Witter MP, Moser EI, Moser MB. Spatial representation in the entorhinal cortex. Science. 2004; 305:1258–1264. [PubMed: 15333832]

Hasselmo ME, Eichenbaum HB. Hippocampal mechanisms for the context-dependent retrieval of episodes. Neural Networks. 2005; 18:1172–1190. [PubMed: 16263240]

Jackson J, Redish AD. Network dynamics of hippocampal cell-assemblies resemble multiple spatial maps within single tasks. Hippocampus. 2007; 17:1209–1229. [PubMed: 17764083]

Kim SM, Frank LM. Hippocampal lesions impair rapid learning of a continuous spatial alternation task. PLoS One. 2009; 4:e5494. [PubMed: 19424438]

Koch, C.; Segev, I. Methods in neuronal modeling : from ions to networks. Cambridge, Mass: MIT Press; 1998.

Lansky P, Fenton AA, Vaillant J. The overdispersion in activity of place cells. Neurocomputing. 2001; 38:1393–1399.

Lee I, Griffin AL, Zilli EA, Eichenbaum H, Hasselmo ME. Gradual translocation of spatial correlates of neuronal firing in the hippocampus toward prospective reward locations. Neuron. 2006; 51:639–650. [PubMed: 16950161]

Lipton PA, White JA, Eichenbaum H. Disambiguation of overlapping experiences by neurons in the medial entorhinal cortex. Journal of Neuroscience. 2007; 27:5787–5795. [PubMed: 17522322]

Mizumori SJY, Ward KE, Lavoie AM. Medial Septal Modulation of Entorhinal Single Unit-Activity in Anesthetized and Freely Moving Rats. Brain Research. 1992; 570:188–197. [PubMed: 1617411]

Muller R. A quarter of a century of place cells. Neuron. 1996; 17:813–822. [PubMed: 8938115]

O'Keefe, J.; Nadel, L. The hippocampus as a cognitive map. Oxford, New York: Clarendon Press; Oxford University Press; 1978.

Parzen E. On Estimation of a Probability Density Function and Mode. The Annals of Mathematical Statistics. 1962; 33:1065–1076.

Pastalkova E, Itskov V, Amarasingham A, Buzsaki G. Internally generated cell assembly sequences in the rat hippocampus. Science. 2008; 321:1322–1327. [PubMed: 18772431]

Pfeiffer RR, Kiang NY. Spike Discharge Patterns of Spontaneous and Continuously Stimulated Activity in the Cochlear Nucleus of Anesthetized Cats. Biophysical Journal. 1965; 5:301–316. [PubMed: 19431335]

Prerau MJ, Eden UT. A General Likelihood Framework for Characterizing the Time Course of Neural Activity. Neural Computation. 2011

Quirk GJ, Muller RU, Kubie JL, Ranck JB. The Positional Firing Properties of Medial Entorhinal Neurons - Description and Comparison with Hippocampal Place Cells. Journal of Neuroscience. 1992; 12:1945–1963. [PubMed: 1578279]

Redish AD, Touretzky DS. Cognitive maps beyond the hippocampus. Hippocampus. 1997; 7:15–35. [PubMed: 9138665]

Rosenthal, J.; Gilliam, DS. Mathematical systems theory in biology, communications, computation, and finance. New York: Springer; 2003.

Robitsek J, White J, Eichenbaum H. Place cell activation predicts subsequent memory. Behavioural Brain Research. 2013; 254:65–72. [PubMed: 23295394]

Shadlen MN, Newsome WT. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. Journal of Neuroscience. 1998; 18:3870–3896. [PubMed: 9570816]

Sun, W.; Yuan, Y-x. Optimization theory and methods : nonlinear programming. New York: Springer; 2006.

Touretzky DS, Muller RU. Place field dissociation and multiple maps in hippocampus. Neurocomputing. 2006; 69:1260–1263.

Touretzky DS, Redish AD. Theory of rodent navigation based on interacting representations of space. Hippocampus. 1996; 6:247–270. [PubMed: 8841825]

Truccolo W, Eden UT, Fellows MR, Donoghue JP, Brown EN. A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. Journal of Neurophysiology. 2005; 93:1074–1089. [PubMed: 15356183]

Tuckwell, HC. Introduction to theoretical neurobiology. Cambridge ; New York: Cambridge University Press; 1988.

Turlach, BA. Institut de Statistique, Voie du Roman Pays 34, B-1348 Louvain-la-Neuve, Belgium Discussion Paper 9317. 1993. Bandwidth selection in kernel density estimation: A review.

Ventura V, Cai C, Kass RE. Trial-to-trial variability and its effect on time-varying dependency between two neurons. Journal of Neurophysiology. 2005; 94:2928–2939. [PubMed: 16160096]

Wasserman, L. All of statistics : a concise course in statistical inference. New York: Springer; 2004.

Wood ER, Dudchenko PA, Robitsek RJ, Eichenbaum H. Hippocampal neurons encode information about different types of memory episodes occurring in the same location. Neuron. 2000; 27:623–633. [PubMed: 11055443]

Appendix A: Empirical Probability Surfaces

This section outlines the details of the computations in involved in creating empirical probability surfaces. We define an empirical probability surface as a spatiotemporally evolving probability density model of the data.

When the empirical model of density is based on a normalized histogram, then computationally, the basis of an empirical probability surface is nothing more than a normalized two-dimensional histogram with position and firing rate as variables.

We perform our analysis in discrete time, which we define as

$$t_i = i\Delta t \quad \text{(A0)}$$

where Δt is the sampling frequency.

We define our data as the set of simultaneously occurring position and rate observations

$$\begin{aligned} x_i &= X(t_i) \\ r_i &= R(t_i) \end{aligned} \quad \text{(A0)}$$

where $X(t)$ and $R(t)$ are the animal's linear position and the continuous time firing rate estimate for the given neuron, respectively.

The empirical probability surface *EP* is defined in terms of intervals in the position space with size $\Delta\chi$, and of intervals in the rate space with size $\Delta\rho$ as

$$EP(r,x) \propto \left[\sum_{i=1}^{N_\rho} I(r,x,r_i,x_i)\right] + \gamma(r) \quad \text{(A0)}$$

where *I* is an indicator function such that

$$I(r,x,\rho_i)=\begin{cases}1, & \text{if } r_i \in [r\Delta\rho,(r+1)\Delta\rho)\ \&x_i \in [x\Delta\chi,(x+1)\Delta\chi)\\ 0, & \text{otherwise}\end{cases} \quad \text{(A0)}$$

To avoid the possibility of having portions of the distribution with zero probability, a gamma distribution with a peak at 0 Hz is added to the empirical density model and renormalized. The form of the gamma probability density function (DeGroot and Schervish, 2002) is

$$\gamma(x)=\frac{1}{b^a\Gamma(a)}x^{a-1}e^{-\frac{x}{b}} \quad \text{(A0)}$$

The added gamma distribution ensures that there is a positive value at every point of the empirical distribution. Thus, computational and interpretational problems associated with probabilities of 0 are avoided. Building the empirical distribution on top of the gamma also enables the calculation of a distribution even if no spikes are observed.

In our computations, we divided the 90cm long central portion of T-maze stem into 50 equal position intervals so that $\Delta\chi$ =1.8cm, and was a linearized representation of a region 1.8cm long and 5.7cm wide. For the intervals in the rate space, we set $\Delta\rho=\frac{R_{\max}}{N_{rbin}}$, where $R_{\max}$ is the maximum firing rate observed in the stem during the duration of the experiment, and N_{rbin} is the number of rate-space bins used—in this case, N_{rbin} = 30. Thus, the basis for our empirical probability surface was a normalized 30×50 two-dimensional histogram of firing rate over the animal's linear position on the stem and rate frequency. For the prior gamma distribution in equation (A0), we used parameters a = 2, b = 0.3 to create a distribution with a long tail and a mode close to 0Hz.

For visualization purposes and to compute a smooth model of the empirical density from the histogram, the empirical probability surface was smoothed before it was normalized. Each of the 50 columns in the histogram corresponding to position was smoothed using a Savitzky-Golay smoothing filter (Rosenthal and Gilliam, 2003) with a span of 7 and a degree of 2, computed using the MATLAB *smooth* function. Then the entire surface was convolved with a 3×3 2D Gaussian. We then used a two-dimensional cubic interpolation (Sun and Yuan, 2006) to up-sample the resolution 300×300 using the MATLAB *interp2* function. Finally the entire surface was normalized so that every position column of the empirical probability surface integrated to one.

The purpose of the empirical probability surfaces in this paper is to provide initial insight into the complexity of the underlying densities and to motivate exploratory qualitative analyses. The smoothing parameters were chosen to produce a simple to compute, reasonable, smoothed estimate of the empirical density at each position interval. The question of how to properly smooth a set of evolving densities is beyond the scope of this work, and will likely combine sophisticated methods of density estimation with adaptive smoothing in two dimensions.

Appendix B: A Bootstrap Procedure for Generating Regions of Significance

In this section, we develop a bootstrap procedure to compute and compare based on empirical firing data from small numbers of experimental trials.

The bootstrap (Efron and Gong, 1983) is a class of the Monte Carlo method that can be useful in calculating the distribution of the value of a statistic from data with a small population size and unknown or complex underlying distribution. This technique is highly applicable to the data recorded in electrophysiological behavior experiments, where there may only be a few trials in a given context, which possess an unknown structure.

For a given cell, we first calculate the firing rate estimate as a function of position for every trial in the experimental session. For computational efficiency, we downsample each trial rate trajectory by taking its mean firing rate in 50 equal 1.8cm positional intervals. We next sample with replacement from all the trials to get a bootstrap sample *S*, a collection of trial-to-trial rate trajectories across T-maze position. A given sample statistic (e.g. mean, variance, 95th percentile) is then computed on *S* at each of the 50 positional intervals. This procedure is repeated by generating some large number N_B bootstrap samples. In this case, $N_B = 500$. For each bootstrap sample we calculate and store $\hat{\varphi^i}$, the values of the given sample statistic at each positional interval for the i^{th} bootstrap sample. We define $\hat{\Psi} = \{\hat{\varphi^1}, \ldots, \hat{\varphi^{N_B}}\}$ as the set of estimated values of the sample statistic from every iteration. We can use $\hat{\Psi}$ to estimate a sampling distribution on the given sample statistic on the original data, from which confidence bounds can be constructed.

To construct confidence bounds of a level α around the statistical estimate, we can start by taking the $\frac{\alpha}{2}$ and $100\left(1-\frac{\alpha}{2}\right)$ percentiles of $\hat{\Psi}$. As there is likely to be a positional dependency, we employ a method similar to that used in Fujisawa et al. (2008) and create global confidence bounds. This is done by finding the point-wise percentiles of $\hat{\Psi}$ within which $(1-\alpha)N_B$ of the bootstrapped statistical estimates $\hat{\varphi}$ are completely contained at every interval.

In summary, we estimate the sampling distribution of a given sample statistic over a set of positional intervals from a set of firing rate estimate trajectories ρ from *N* experimental trials, we use the following procedure:

Algorithm 1

Estimating Statistics on Empirical Firing Rate Data Using the Bootstrap

1 Repeat some large number N_B times:

- **a.** Generate a bootstrap sample *S*, by sampling with replacement *N* times from the trial-to-trial firing rate trajectory data set ρ
- **b.** Calculate $\hat{\varphi^i}$, the sample statistic on *S* at each positional interval for the i^{th} iteration
- **c.** Store the value of $\hat{\varphi^i}$ in $\hat{\Psi}$

2 Take the median of all N_B values of $\hat{\Psi}$ computed in step 1 as an estimate of the sample statistic on the original data

3 Calculate global bounds for a given confidence level α by finding an integer *a* such that the $\frac{a}{2}$ and $100\left(1-\frac{a}{2}\right)$ point-wise percentiles of $\hat{\Psi}$ contain $(1-\alpha)N_B$ of the $\hat{\varphi}$ computed in step 1

The bootstrap can also be used to estimate whether two sets of firing rates differ with respect to given statistic at a given confidence level. For each individual bootstrap iteration, we generate bootstrap samples from both of the data sets individually, calculate a given statistic for each set, and compute the difference of the two resultant values. If we store this difference for many iterations of the bootstrap, we can construct an estimate of the distribution of the difference statistic for the two sets. We then compute global bounds on

the distribution of the difference statistic at a given confidence in the same manner as for the individual statistic. If these bounds contain zero, then we conclude that there is no difference in the value of the statistic at that confidence level. If zero falls outside those bounds, then we conclude that there is significant difference in the value of the statistic for the two sets at that confidence level.

To determine whether a given sample statistic differs with confidence a from two collections trial-to-trial of firing rate trajectories across position ρ_1 and ρ_2, with N_1 and N_2 measurements respectively, we perform the following procedure:

Algorithm 2

Procedure for Confidence in the Difference Statistics for Empirical Firing Rate Data Using the Bootstrap

1 Repeat some large number N_B times:

- **a.** Generate a bootstrap sample S_1 by sampling with replacement N_1 times from the firing rate trajectory data from ρ_1
- **b.** Generate a bootstrap sample S_2 by sampling with replacement N_2 times from the firing rate trajectory data from ρ_2
- **c.** Calculate $\widehat{\phi}_1^i$ and $\widehat{\phi}_2^i$, the sample estimates of the given statistic at each positional interval from the data in S_1 and S_2, respectively, for the i^{th} iteration
- **d.** Calculate $\widehat{\phi}_{\text{diff}}^i = \widehat{\phi}_1^i - \widehat{\phi}_2^i$
- **e.** Store the value of $\widehat{\phi}_{\text{diff}}^i$ in $\hat{\Psi}_{diff}$

2 Calculate global bounds for a given confidence level a by finding an integer a such that the $\frac{a}{2}$ and $100\left(1-\frac{a}{2}\right)$ point-wise percentiles of $\hat{\Psi}_{diff}$ contain $(1-a)N_B$ of the $\widehat{\phi}_{\text{diff}}^i$ computed in step 1

3 If zero falls outside the bounds computed in step 2, then conclude with confidence a that the sample statistic differs for ρ_1 and ρ_2

We apply these techniques to the experimental neural data by defining a *region of significance* for a given sample statistic at confidence a as a contiguous positional interval during which the global confidence bounds of the difference statistic exclude zero. Thus, a region of significance indicates a positional interval on the T-maze stem where the neural firing for two different contexts differs with respect to a particular statistic. A neuron may possess more than one region of significance for a given statistic.

Characterizing contextual differences in neural firing by calculating regions of significance is useful because it provides information on the locations at which the differences in firing are occurring. Rather than perform a single hypothesis test, which can only accept or reject a specified null hypothesis, regions of significance are able to incorporate spatial information into our understanding of evolving distributional differences beyond a binary conclusion. Moreover, using regions of significance allows us to perform higher resolution analyses than conventional methods. With single p-value tests such as the ANOVA, the chance for a spurious significant finding increases as the resolution of the analysis increases. Using global confidence bounds enables us to use greater spatial precision without sacrificing accuracy.

Perhaps the most important feature of bootstrap analysis on an empirical probability surface is that regions of significance can be calculated for any arbitrary statistic. Standard hypothesis tests are statistic and distribution specific, and a new test must be found to fit

different distributions or to analyze different statistics. This bootstrapping procedure is able to easily accommodate any statistic for any distribution structure without requiring any major alteration in the computation other than the implementation of the statistical calculation itself. Consequently, this makes it easy to explore the structure of distributions with a battery of multiple different statistical measures, enabling us to determine which statistics are the most informative for a given system.

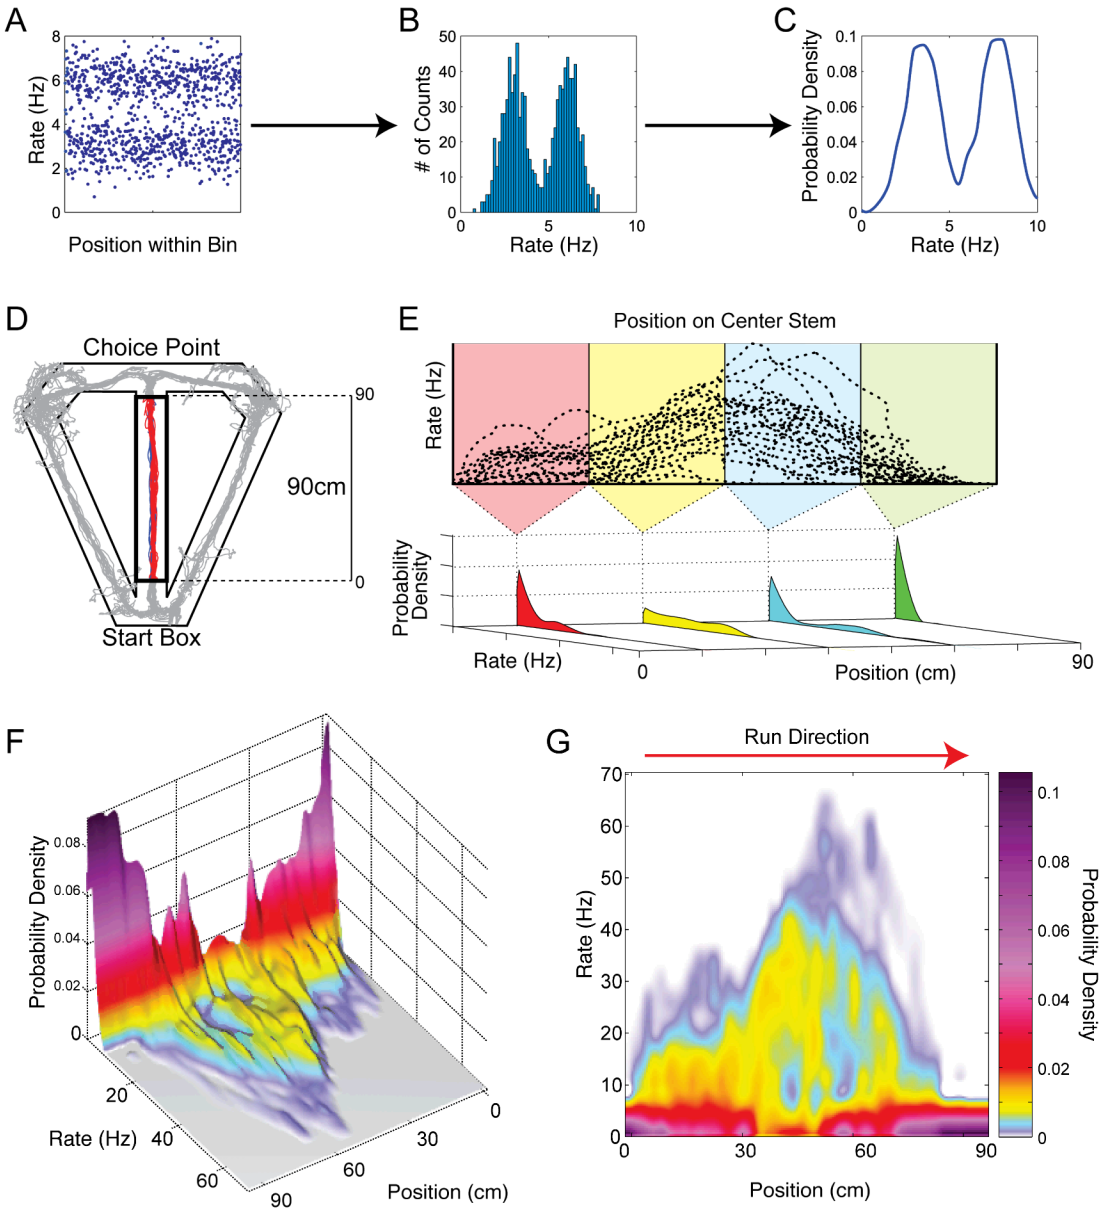

Figure 1.

A method for the calculation of the smoothed empirical probability surface for a cell from dcMEC. To calculate an empirical distribution, rates from a spatial window (A) are used to calculate a histogram (B), which is then normalized by the number of observations and smoothed to create estimate an empirical probability density function (C). A 90cm center section of the T-maze stem with consistent left (blue) and right (red) turn run trajectories (D) is broken up into multiple windows and the empirical probability surface is computed for each window (E). The final smoothed high-resolution empirical probability surface is shown in 3-D (E) and 2-D (F), with the probability density represented by the height/color of the surface.

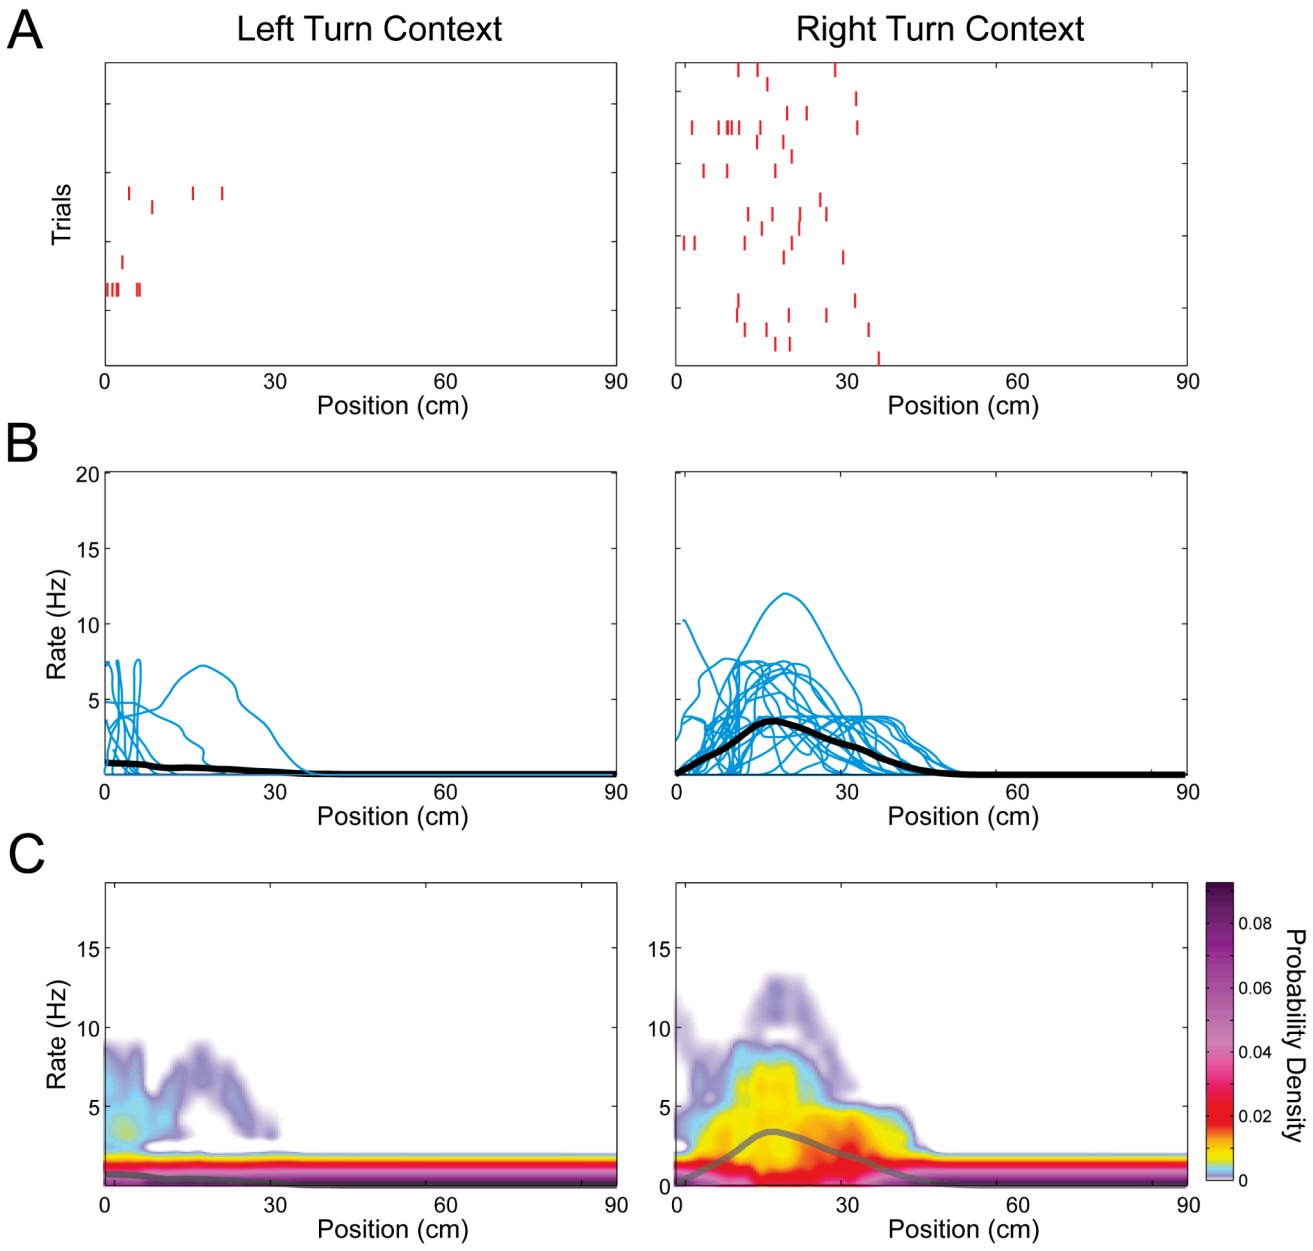

Figure 2.

Comparative representations of differential firing from a cell in the rat dcMEC during a spatial alternation task on a T-maze. The neural activity from this cell on the stem of the T-maze before left-turns (left panels) and before right-turns (right panels) is shown using the raw spike rasters as a function of position (A), the mean firing rate (B, black curves) taken from individual firing rate trajectories (C blue curves), and using empirical probability surfaces (C). The mean firing rate is superimposed on the empirical probability surfaces for comparison.

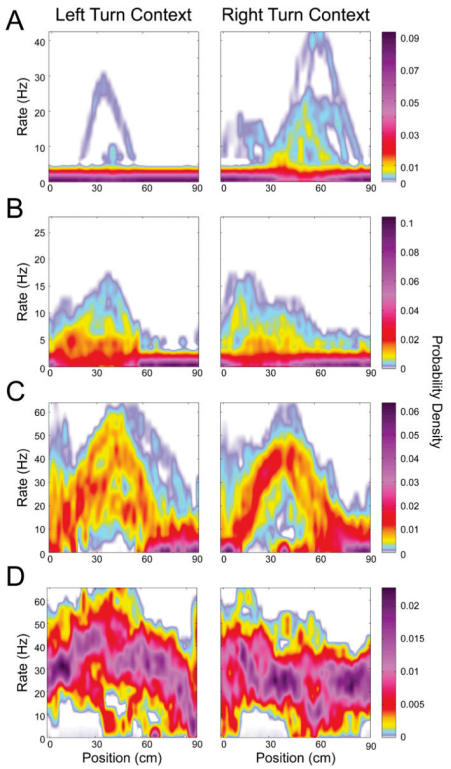

Figure 3.
Examples of ways in which of differential firing encoding paradigms was observed in both CA1 and MEC. Shown are illustrative examples of cells exhibiting preferential one-sided (classic splitter) (A), field shifting (B), context-modulated variability (C), and differential interneuron (D) firing patterns.

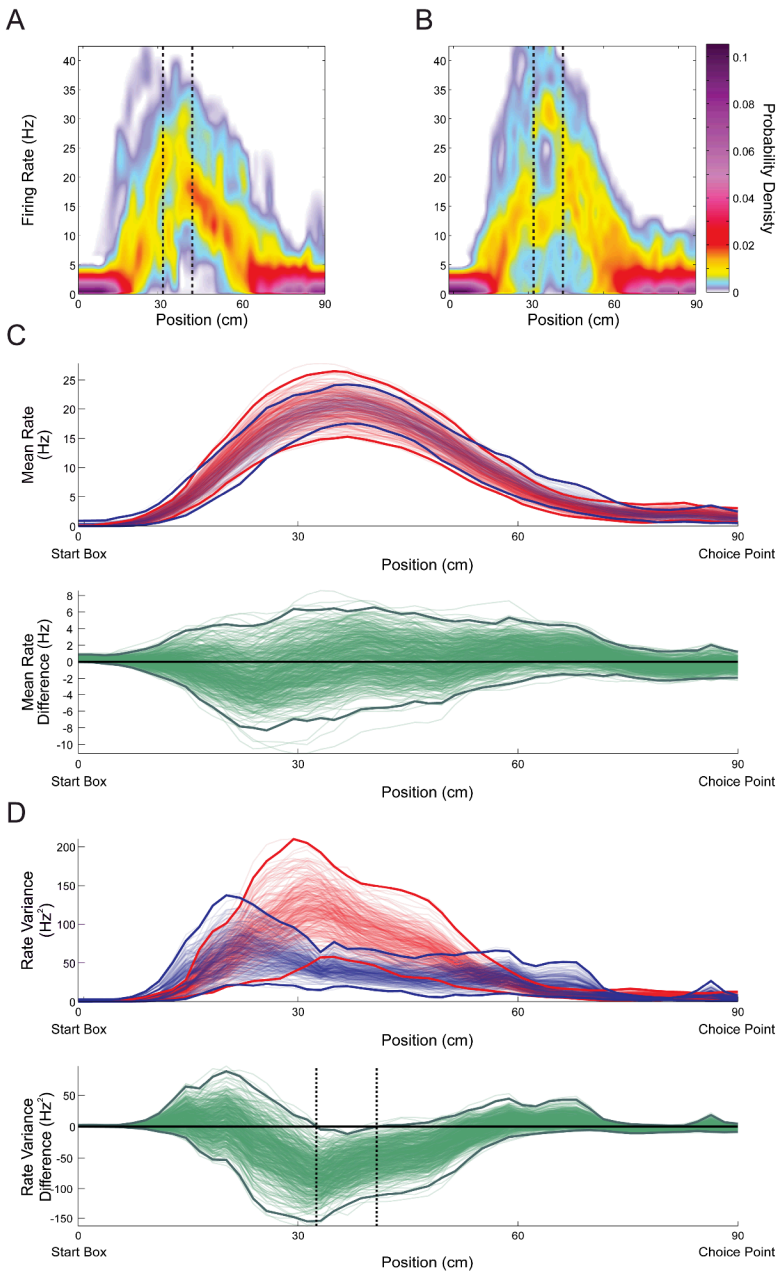

Figure 4.

Regions of significance for a cell from the rat CA1 showing differential firing distributions during a spatial alternation task on a T-maze. The empirical probability surfaces for the neural activity on the stem preceding left-turns (A) and right-turns (B) are shown. We calculate the regions of significance for the trial-to-trial mean (C) as well as the firing rate variance (D). In (C) and (D), the top panels show the bootstrapped left-turn (light blue curves) and right-turn (light red curves) statistical estimates with the 95% global confidence bounds (dark curves). The bottom panels show the bootstrapped difference statistic (light green curves) with 95% global confidence bounds (dark green curves) along with the zero line (black line). This analysis finds regions of significance (A, B, and D vertical dotted bars) for the variance but not the mean for this cell.

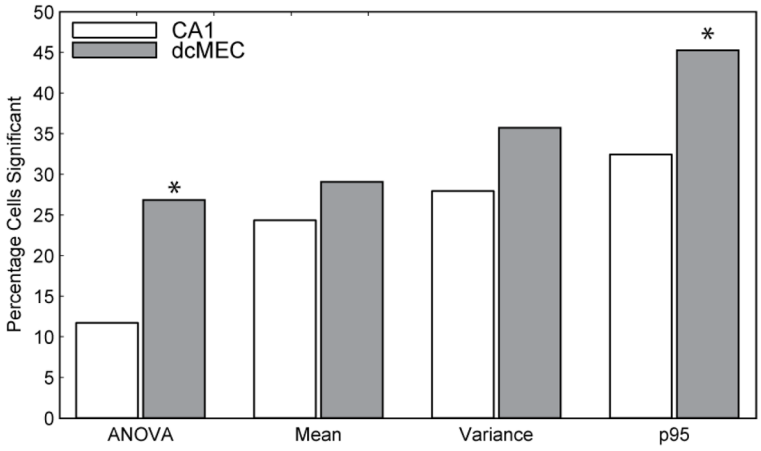

Figure 5.

The percentage of differentially firing cells as detected by the ANOVA and by bootstrap statistics for CA1 (white bars) and dcMEC (gray bars) regions. There was a significantly greater the proportion of cells in dcMEC than in CA1 that were significant for the 95th percentile (z-test, $p<.05$), as well as for the ANOVA (z-test, $p<.01$).

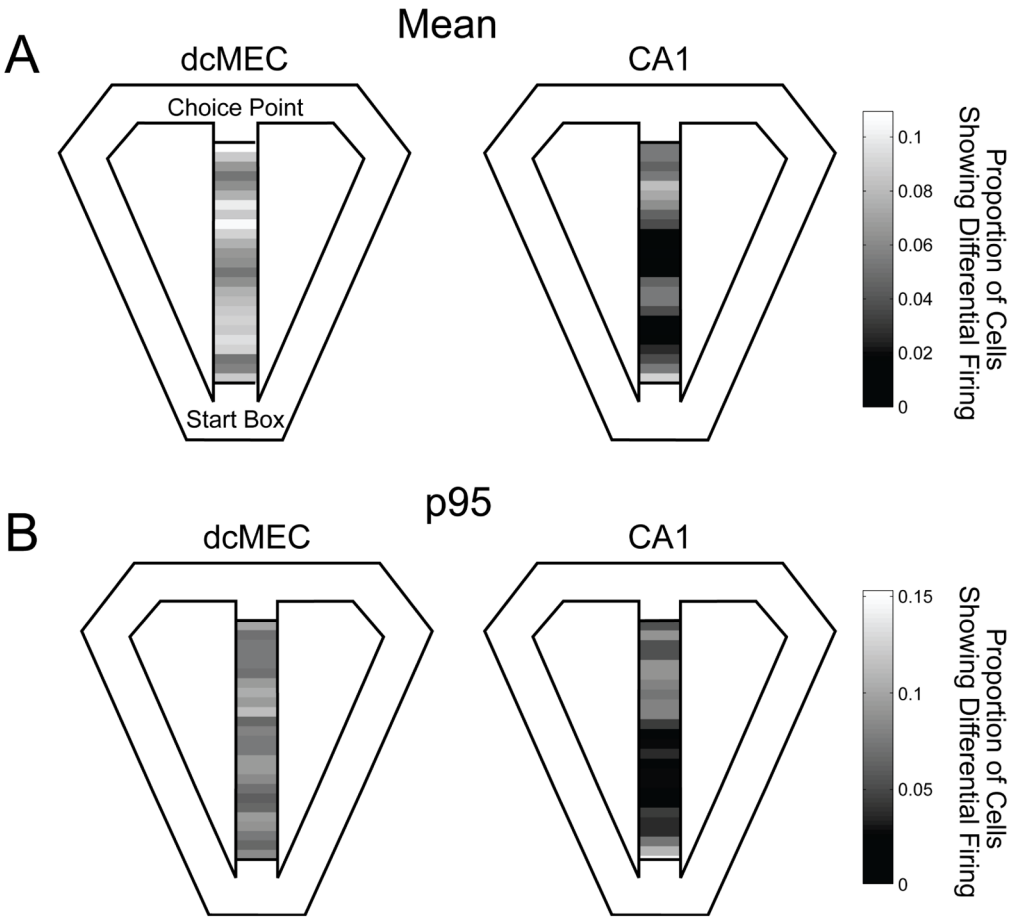

Figure 6.

The spatial distribution of differential firing across the T-maze stem. The proportion of cells in dcMEC (left schematic) and cells in CA1 (right schematic) exhibiting differential firing in the mean (a) and 95th percentile (b) statistics are shown as a function of linear position, mapped to T-maze schematics, with brightness corresponding to proportion magnitude. For mean, variance, and p95 statistics, there was a significant difference between proportions as a function of brain region. For the p95, the spatial distribution for CA1 was significantly different from uniform, showing a region of reduced differential firing in the middle of the stem.

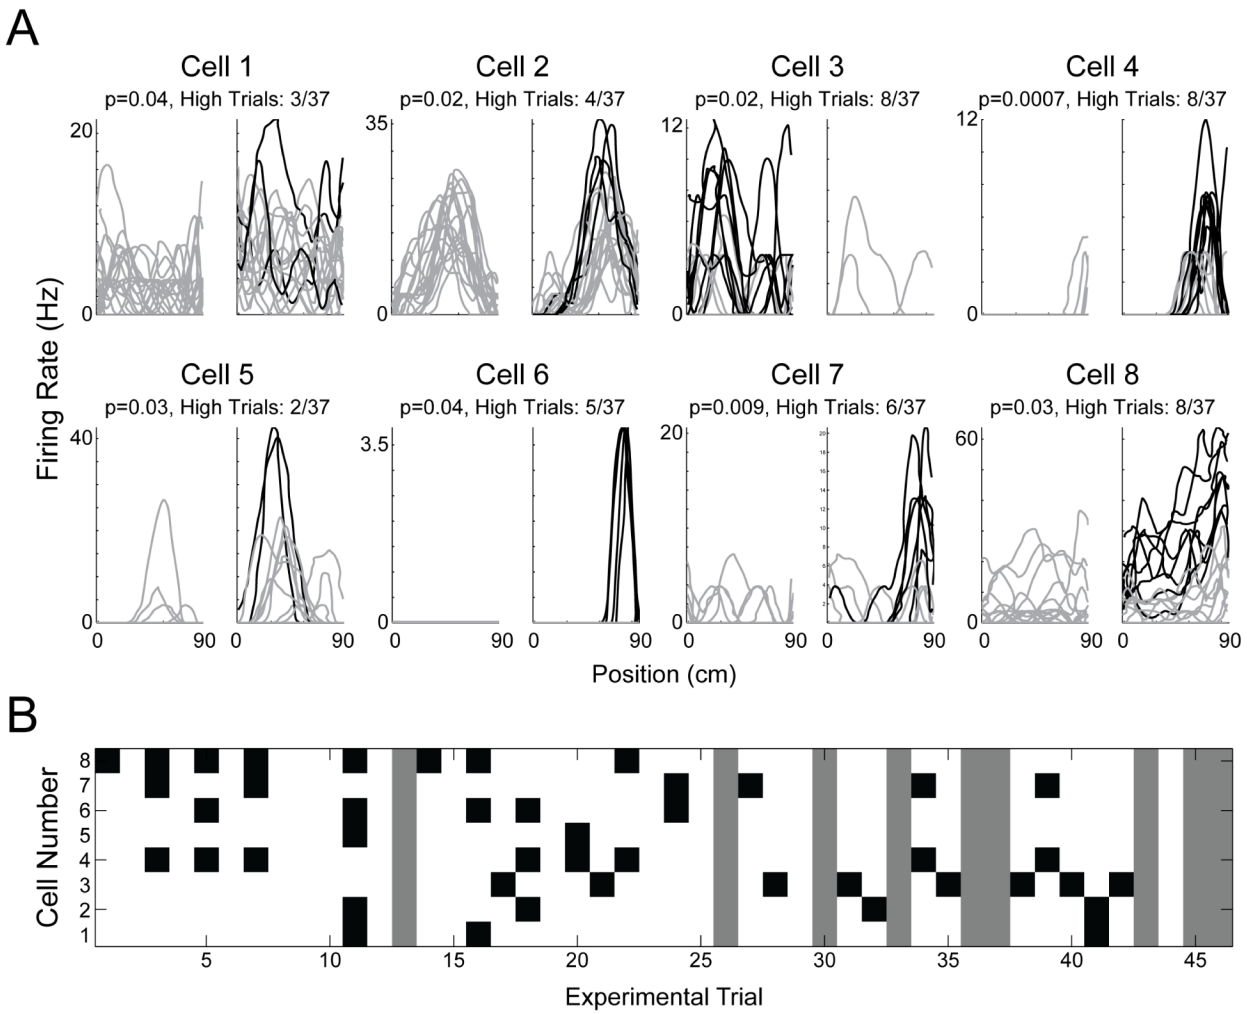

Figure 7.

Sparse individual-cell encoding of context is robust across the population. Presented is the subpopulation of 8 units (out of a population of 37 simultaneously recorded cells from dcMEC) that had statistically significant context-dependent differences in the 95th percentile (a). The computed rates for the neural activity on the stem during a spatial alternation task on the T-maze are shown for the left-turn (left panels) and right-turn (right panels) contexts. The trials above a threshold defined by the maximum firing rate of the context with less firing are accentuated in black. The p-values for the context-dependent differences 95th percentile are reported. The firing rates are displayed in a matrix (b) as a function of cell (rows) and experimental trial (columns). For each cell, the trials above threshold are marked in black, and excluded trials are marked in gray. A pairwise correlation analysis on this matrix shows that the majority of the cell-to-cell interactions are not significantly correlated, suggesting that these trials are well-distributed throughout the experiment, and not due to large behavioral anomalies.