Classifying neurons based on their activity during navigation in virtual reality

Malcolm Campbell and Mark Plitt

Introduction

How do animals navigate their environments? In mammals, part of the answer lies in a network of brain regions whose neurons signal the current location and movement of the animal, providing the basis of a "cognitive map."[1] Two of the primary brain regions in this network are the hippocampus and its input, the entorhinal cortex, which can be divided into medial and lateral zones. In the medial entorhinal cortex (MEC), a type of neuron called a "grid cell" becomes active when the animal is at regularly spaced positions that tile the environment and form the nodes of a triangular lattice (see Figure 1D for an example). Another type of cell, called a "border cell," becomes active when an animal is near to the edge of its environment. Other cells are modulated by the animal's location but in idiosyncratic patterns, and others are not modulated by position. These cells are thought to play distinct computational roles in navigation and we are often interested in analyzing their activity separately when an animal performs a navigational task.

To study these neural activity patterns and their relationship to behavior, we make electrical recordings from the brains of mice while they explore natural and virtual environments. From these electrical recordings, we isolate discrete events called action potentials (or "spikes") that come from individual neurons in the mouse's brain. Our data set therefore consists of a set of neurons; for each neuron, we record a series of discrete spike times. We also track the animal's location in its environment, and attempt to relate the spiking activity of cells to the animal's behavior.

These functional cell classes are defined by their response properties when an animal freely moves around in an environment (Figure 1C), but in order to use state of the art recording techniques and precisely control the animal's experience we often need to hold the animal's head fixed in one place and allow it to explore a virtual reality environment (VR; Figure 1E). Consequently, a tool for distinguishing these cell classes based on their activity in VR alone would be of great practical use. A previous report claimed to be able distinguish grid cells from non-grid cells with 87% true-positive rate and 13% false positive rate[2]. Here, we re-examine this claim with a larger and more balanced dataset, focusing on classifying grid vs. non-grid and grid vs. border cells, and find these numbers to be unrealistic. However, when we add features derived from the cells' responses to changes in the rate of optic flow, our classifiers separate grid and border cells at up to 80% accuracy. This success is likely due to the fact that these cell types respond to different sensory inputs. Finally, we define a new cell class which we call "drifty bursting cells" and hand-label these cells in our dataset. These cells are enriched for but not restricted to grid cells, and can be identified with 75% accuracy based on 7 hand-picked features.

Dataset

Our dataset consists of recordings from 781 MEC cells from 20 mice. Each recording consists of a series of discrete spike times from an individual neuron while the animal explored its environment. We simultaneously tracked the location of the animal throughout the session. Usually multiple neurons were recorded at once (typically 4-10 at a time). A neuron could be identified and matched between recordings based on its unique waveform signature across channels. Each neuron was recorded while the animal explored two environments: an open field environment (OF) and a virtual reality (VR) environment. The open field environment consisted of a 90cm x 90cm square box. In the virtual reality setup, the mouse's head was fixed in place while it ran on a cylindrical treadmill, with a virtual scene projected onto a hemispherical screen in front of the animal. As the mouse ran, the visual scene moved as a function of its running speed, creating the illusion of traversing a 400 cm virtual linear hallway. When it reached the end of the hallway, it received a reward (a water droplet; mouse was water deprived) and was "teleported" back to the beginning of the track for the next trial.

Previously published scores were used to classify cells as grid, border, or neither based on open field recordings[3,4]. Open field recordings were only used to classify cells; all other analyses were run on recordings during exploration of the VR environment. The dataset contains 96 grid, 97 border, and 590 other cells (2 cells were classified as both grid and border). Since many cells were recorded during multiple VR sessions, the dataset contains 1347 distinct VR recordings from 781 cells.

Features

Firing rate maps: To assess the relationship between a neuron's spiking and the animal's location along the virtual linear track, we split the track into 200 2 cm bins, counted the number of spikes that occurred in each bin and the amount of time the animal spent in that bin, and computed the "firing rate map" across the track by dividing these two numbers. A rate map is therefore a 200 dimensional feature vector, with each entry indicating the neuron's firing rate at that location in the environment. Rate maps were smoothed using a Gaussian kernel with standard deviation of 1 bin. Single trial rate maps were computed similarly but using only spikes and position data from a single trial (a trial is one full traversal of the track). Here and in subsequent analyses, only the first recording was used if multiple recordings were made from the same cell, unless otherwise noted.

Fourier transform of firing rate maps: To capture periodicity in the rate maps, which is characteristic of grid cells, we computed the magnitude of the Fourier transform of the rate map (100 features).

Stability metrics: We are interested in the stability of rate maps over time. To capture this information in a feature, we compute Pearson correlations between rate maps computed in the first and second half of a session. We also compute trial-to-trial stability as the average correlation between single trial rate maps from adjacent trials.

Gain manipulations: In a subset of our experiments, we performed what we call "gain manipulations," where we altered the gain of the relationship between the rotation of the running wheel and the translation of the visual scene (Figure 2). This changes the amount of optical flow experienced by the mouse as it runs and effectively lengthens (gain decreases) or shortens (gain increases) the track. The logic behind these experiments was that they could allow us to separate the influence of locomotion from visual input on the firing patterns of cells. Gain manipulations were run in blocks, with 5- or 10-trial blocks of changed gain—either decreased by 0.5 or increased by 1.5—interleaved with longer blocks of baseline gain. Similarly to the stability metrics, we generated rate maps during baseline and gain-changed periods and computed cross-correlations between them, with lags of up to +/- 50 bins (101 features). If cells were recorded on multiple gain manipulation sessions, their cross-correlation maps from multiple sessions were averaged together. Cells were only included in this analysis if they had mean firing rate greater than 0.2 Hz in both the baseline and gain-changed periods and passed a stability score (z-score vs. shuffled data > 1). These criteria were included to focus on cells that had coherent spatial responses to the gain manipulations.

Burstiness: We noticed that a subset of cells displayed a particular firing pattern in VR where they had short bursts of high activity that were tied to spatial locations on the track, but which drifted slowly over time or abruptly jumped to a new location. We tried to capture this behavior using 7 hand-crafted features: mean firing rate (r); median inter-spike interval (i); "burstiness" (1/(r*i)); stability (s); trial-to-trial stability (t); stability ratio (t/s); and field size (f). The logic behind the burstiness metric is that cells with most of their spikes clustered close together will have high burstiness. We computed the stability ratio to capture cells that had high local stability (short time scale) but low stability at longer time scales. We also noticed that these cells tended to have small fields, and so computed the field size (f) as the first point in the autocorrelation function of the firing rate map that fell below 0.2.

Soft dynamic time-warping: To further characterize the drifting cells described above, we performed Soft Dynamic Time Warping (Soft DTW) to align firing fields across trials[5]. DTW is a popular dynamic programming procedure to stretch and align time series[6] often used to align the timing of spoken words when analyzing speech. In our case, we hope to align the firing fields of individual cells as they drift across trials. Soft DTW is an updated version of the procedure that uses a differentiable softmin cost to take an expectation over all possible alignments. To compute an "average" aligned firing rate map, we applied bary center averaging[7]. This iterative procedure finds the firing rate map that minimizes the squared distance to all of the individual trial firing rate maps. We then computed the Euclidean distance from each trial's firing rate map to the bary centered average, and gave this distance a sign based on whether the center of mass of the soft alignment matrix was below or above unity (an overall shift to the left or the right, respectively, of the firing fields). We then normalized this vector to take on values between 0 and 1. If a cell's firing fields are drifting coherently in the same direction over time, then we should see that this normalized distance will be monotonic function of trial number. Under the assumption that grid cells show this continuous slow drift and other cells do not, we fit a line to the normalized distance as a function of trial number using least squares regression. This also allowed us to compare cells for which there were differing numbers of trials. From this analysis we extracted several features that could be used for cell type classification: the bary-centered average firing rate map, the slope and intercept of the linear regression, and the r^2 value derived from the regression.

Methods

Throughout the report, we trained four classifiers on our data: 1) logistic regression with ℓ_2 normalization; 2) linear Support Vector Machines (SVM); 3) SVM with a Gaussian kernel; 4) Gaussian Discriminant Analysis (GDA). Logistic regression fits a sigmoid function to the data, maximizing the log-likelihood subject to ℓ_2 normalization:

$$\Sigma_{i=1}^{m} y^{(i)} \log(h_\theta(x^{(i)})) + (1 - y^{(i)}) \log(1 - h_\theta(x^{(i)})) + \lambda ||\theta||_2^2 \qquad h_\theta(x) = 1/(1 + e^{-\theta^T x})$$

SVM finds the decision plane and associated support vectors by solving the following minimization problem:

$$\min_{\gamma,w,b} \ \frac{1}{2}||w||^2 + C\sum_{i=1}^{m}\xi_i \qquad \text{s.t.} \ \ y^{(i)}(w^T x^{(i)} + b) \geq 1 - \xi_i, \ \ i = 1, \ldots, m$$
$$\xi_i \geq 0, \ \ i = 1, \ldots, m.$$

By examining the dual problem, we can replace $w^T x$ with $\sum_{i=1}^{m} \alpha_i y^{(i)} {x^{(i)}}^T x$. Replacing ${x^{(i)}}^T x$ with $K(x^{(i)}, x)$, where K is the desired kernel (in this case, a Gaussian), yields the kernelized SVM. GDA fits a separate Gaussian to the distribution of the data for each of the two classes, yielding the parameters:

$$\mu_0 = \frac{\sum_{i=1}^{m} 1\{y^{(i)} = 0\} x^{(i)}}{\sum_{i=1}^{m} 1\{y^{(i)} = 0\}} \qquad \mu_1 = \frac{\sum_{i=1}^{m} 1\{y^{(i)} = 1\} x^{(i)}}{\sum_{i=1}^{m} 1\{y^{(i)} = 1\}} \qquad \Sigma = \frac{1}{m}\sum_{i=1}^{m}(x^{(i)} - \mu_{y^{(i)}})(x^{(i)} - \mu_{y^{(i)}})^T$$

Here, μ_0 is the mean of the distribution when $y = 0$, and μ_1 is the mean of the distribution when $y = 1$. Both distributions share the same covariance matrix Σ in the model. Because of the limited size of our dataset (96 grid cells and 97 border cells), we used leave-one-out cross validation. For comparisons with unequal numbers of cells, we randomly down-sampled the larger class to match the size of the smaller class.

Results

Baseline classification of grid vs. border and grid vs. non-grid cells using features of the rate maps: We first tested whether a set of baseline features derived from the firing rate maps could distinguish grid from non-grid cells or grid from border cells. These features were: the mean firing rate (1 x 1), the firing rate map (200 x 1), the magnitude of the FFT of the firing rate map (100 x 1), the average cross correlation between adjacent trials at lags of +/- 10 bins (21 x 1), and the location of the peak cross correlation (1 x 1). We started with these 323 features as similar features were used in a prior publication claiming to separate grid from non-grid cells[2]. We used large regularization parameters because the number of features exceeded the number of examples (ridge regularization for logistic regression and linear SVM, $\lambda = 100$). All four classifiers performed poorly on our dataset (Table 1), with test accuracy never exceeding 59% for either grid vs. non-grid or grid vs. border comparisons. Errors tended to be balanced between false positives and false negatives. This confirms the visual impression that many non-grid cells resemble grid cells in VR (Figure 1F).

Data augmentation: We wanted our classifiers to generalize firing rate features over locations on the track, such that if a cell is classified as a grid cell with high confidence, any shifted or reflected version of this cell would also be confidently classified as a grid cell. To achieve this, we performed data augmentation, rotating our data in increments of 10 bins, and also reflecting the original and rotated data. We treated each rotated and/or reflected copy of the data as its own new example. The augmented dataset therefore consisted of 3840 "grid cells," 3880 "border cells," and 23,600 "other cells." Even with data augmentation, the four classifiers performed very poorly, overfitting the data even with large amounts of regularization (Table 2).

Gain manipulations separate grid and border cells: Having failed to separate grid from non-grid or grid from border cells using features derived from the rate map, we next turned to data from gain manipulation experiments (Figure 2). The response of cells to gain manipulations will depend on whether their spatial firing is driven by visual or locomotor cues. If they are driven by visual cues, their firing rate patterns will not change during gain manipulations. Alternatively, if they are driven by locomotor cues—the number of steps taken since the beginning of the track—their firing patterns will be changed by gain manipulations. Therefore, if cell types are driven by different sensory and motor cues, they could be separated using their response to gain manipulations.

As explained above, the feature used for gain manipulation comparisons was the 101 x 1 vector of the Pearson cross correlation between the firing rate maps generated in the A (baseline) and B (gain change; either 0.5 or 1.5) periods. The feature vector was of size 101 x 1 because we only considered lags of up to 50 bins. Using this feature, we could separate grid and border cells with 73.9% accuracy for gain decrease sessions (0.5 gain; Table 3) and 81.4% accuracy for gain increase sessions (1.5 gain; Table 3). Logistic regression, linear SVM, and RBF-SVM performed similarly, but GDA performed very poorly (<54% accuracy in both cases). This shows that the cross-correlation features are not well fit by a mixture of two Gaussian distributions.

For logistic regression and linear SVM, training accuracies were similar to test accuracies, indicating that we are not over-fitting our data. We also generated a learning curve where we down-sampled the data, refit the models and recomputed training and test accuracies (Figure 2C). Averaging over 10 iterations of this process gave learning curves that indicated were are nearing the limit of performance using this feature, especially with gain increase features.

A new cell class, "drifty bursting cells:" Having had limited success separating grid from non-grid and grid from border cells, we looked to see if there was structure in our data that was not captured by the grid and border categories. Rather than perform unsupervised clustering, we looked through our dataset and noticed a response type that we call "drifty bursting" behavior (Figure 3A). This response type consists of tight clusters of spiking activity that are clearly related to the animal's location on the track, but whose location drifts over time, disappears, or bifurcates between discrete locations. These responses are interesting because they reflect spatial structure that is not tied to the location of visual landmarks. The activity could instead be shaped by the network architecture of the brain region, i.e. the structure of the connectivity between cells[8]. We found these responses only in a subset of mice (8 out of 20), and hypothesized that they might exist only in a specific subregion of the MEC. In this section, the dataset consisted of 1,244 recording blocks from these 8 mice (from 358 unique cells, 688 recording sessions; multiple recording blocks were recorded in one session). Here, we considered each recording block as a separate example, since activity patterns varied across recordings. We hand-labeled the dataset and identified 375 drifty bursters (30.1%). The drifty burster population was enriched for grid cells: 30.3% of drifty bursters were grid cells, compared to only 15.1% of other cells. We attempted to capture the drifty bursting phenotype using 7 hand-crafted features (see above). The four classifiers performed with up to 78% accuracy separating drifty bursters from other cells (down-sampled to match class sizes), with GDA performing better than the other classifiers (Table 4). Unlike in the gain manipulation data, therefore, these features are reasonably well described as two Gaussian distributions, one for drifty-bursters and one for other cells. PCA on the 7 features showed a moderate separation between the two classes consistent with our classifier performance (Figure 3B).

Soft DTW to extract changes in spatial coding over time: While Soft DTW nicely aligned many drifty bursting cells, it appeared to overfit in many cases (Figure 3C). The procedure created structure in the firing rate maps of cells that had no apparent spatial coding. Our attempt at quantifying the drift via least squares distances to the bary-centered time series found spurious coherent drift in some of these cases as well. Due to these failure modes, including any of these features during classification did not improve performance and in some cases worsened performance. We believe this is a promising procedure for quantifying and "correcting" the drift we see in spatially selective firing, but further work is needed to mitigate overfitting. For example, we may need to add a smoothness penalty when fitting warps to incorporate this notion that drifts from trial to trial should be gradual.

Summary and Future Directions

Here, we attempted to separate MEC grid cells from other MEC cell types using recordings made during exploration of a VR environment. The primary motivation was to allow us to study these cell populations using techniques that require head-fixation, such as imaging and whole-cell recording. A previous report claimed to be able to separate grid from non-grid cells based on VR recordings[2]. We were unable to achieve good separation based on features of the firing rate map alone, even with data augmentation. However, adding features derived from cells' responses to gain manipulations improved performance. This could reflect the fact that different cell types derive their spatial responses from different inputs. Error analysis suggested that 80-85% accuracy is the limit of performance using these features. Finally, we identified by eye and hand-labeled a new cell class which we call "drifty bursting cells." This pattern of activity suggests a network architecture that creates spatial structure independently of visual landmarks, and grid cells are thought to be organized in a network architecture that could produce this phenomenon[8]. Interestingly, 53.7% of grid cells were labeled as drifty bursting cells, compared to 28.5% of the rest of the population. Using a small set of features (n = 7), our classifiers separated drifty bursting cells from other cells with ~75% accuracy. We implemented Soft DTW in an attempt to quantify these changes in spatial coding over time but found that without further adjustments this method found structure in our data even when it was clearly absent.

We intend to continue this project after the class. There are several issues that need to be resolved. First, we would like to know how much unbalanced class sizes will affect the performance of our classifier. Relatedly, we are typically more interested in minimizing false positives than false negatives. We will develop penalties to encourage the classifiers to share this preference. Second, we will refine the definition of drifty bursting cells and develop features that allow our classifiers to better match human performance, including the confidence that the cell is indeed a drifty burster. From experience, we suspect that whether or not the cell is a drifty burster will predict its response to gain manipulations; we will test whether this is indeed the case. Finally, we believe dynamic time-warping is a promising approach to extract information about changes in spatial tuning over time. However, the algorithm overfit a substantial portion of our data. We intend to troubleshoot this method using various regularization approaches, such as penalizing the roughness of the fit across trials. In sum, we believe these findings will be a valuable contribution to the field, having both clarified the extent to which grid cells are separable from other MEC cells in VR, and revealed a new cell type that could reflect a specialized mechanism for spatial coding in the brain.

Table 1: Baseline Classification Accuracy

Grid vs. Non-Grid

	L2-Logistic	Linear-SVM	RBF-SVM	GDA
Train	54.25	54.26	85.98	90.59
Test	50	50	58.85	55.21

Grid vs. Border

	L2-Logistic	Linear-SVM	RBF-SVM	GDA
Train	52.19	61.08	84.35	98.97
Test	52.33	58.03	58.03	54.40

Table 2: Classification Accuracy with Data Augmentation

Grid vs. Non-Grid

	L2-Logistic	Linear-SVM	RBF-SVM	GDA
Train	97.93	97.88	97.92	97.5
Test	50.26	50	50	50

Grid vs. Border

	L2-Logistic	Linear-SVM	RBF-SVM	GDA
Train	97.51	97.51	97.93	97.71
Test	50.26	50.26	50.26	49.74

Table 3: Grid vs. Border Classification with Gain Manipulations

Gain Increase Sessions

	L2-Logistic	Linear-SVM	RBF-SVM	GDA
Train	80.7	84.43	92.69	1
Test	77.14	81.43	81.43	52.86

Gain Decrease Sessions

	L2-Logistic	Linear-SVM	RBF-SVM	GDA
Train	77.58	78.11	90.03	1
Test	73.91	72.46	71.01	53.62

Table 4: Drifty-Burster vs. Non-Drifty-Burster Classification Accuracy

	L2-Logistic	Linear-SVM	RBF-SVM	GDA
Train	73.97	74.01	87.51	78.68
Test	73.71	73.79	74.68	78.38

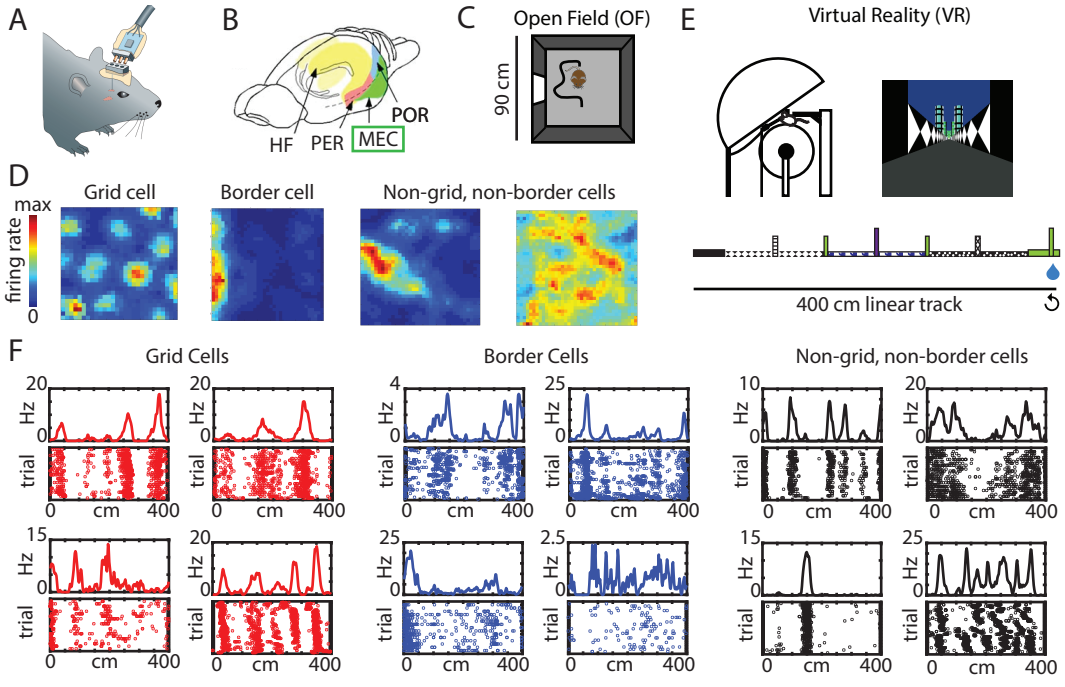

Figure 1

A) Tetrode recordings in mouse. Wires are implanted in the mouse's brain and connected to a detachable cable. Electrical signals from individual neurons are recorded while the animal explores its environment. **B)** Drawing of a mouse brain. Neurons were recorded in the medial entorhinal cortex (MEC). **C)** Cells were classified based on recordings in an "open field" (OF) environment (a square box). **D)** Examples of classical MEC cells in OF. Color coding shows the firing rate of the cell at each location in the environment. **E)** Diagram of the VR setup. **F)** Examples of MEC cells in VR. Four examples of each cell type are shown. Bottom panel: Locations of spikes on the VR track across trials. Top panel: Average firing rate of the cell over the track.

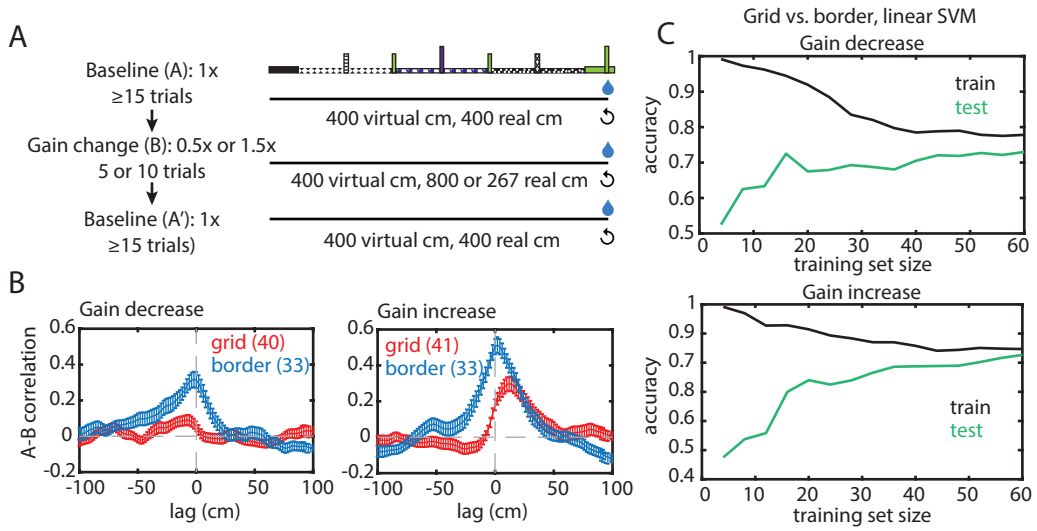

Figure 2

A) Schematic of gain manipulation experiments. **B)** Average cross correlations between firing rate maps computed in A and B periods for grid and border cells in gain decrease and gain increase experiments. **C)** Accuracy of the linear SVM classifier on train and test data as a function of training set size. Random downsampling was performed 10 times per training set size and averaged to generate these plots.

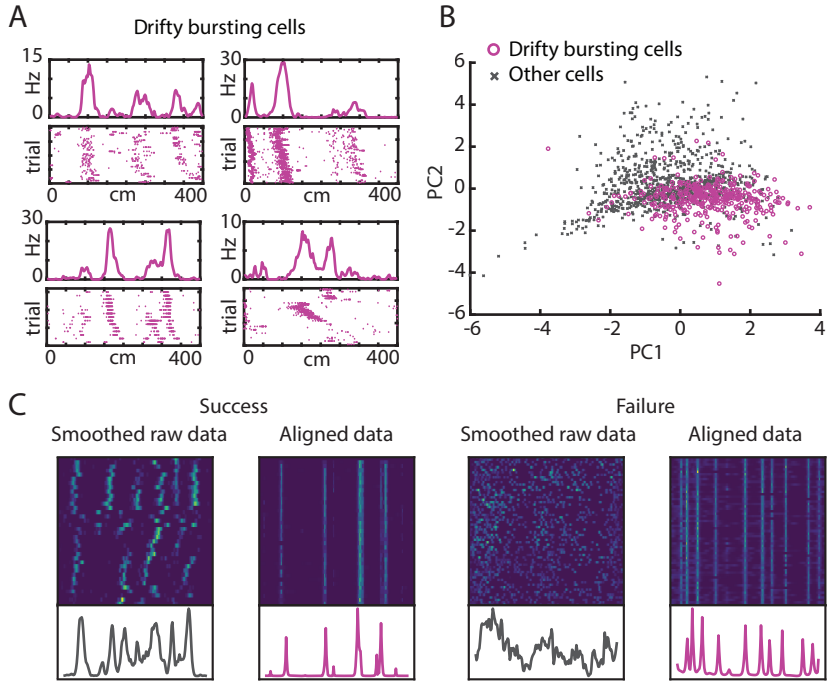

Figure 3

A) Example drifty bursting cells in VR. These cells were identified and labeled by hand. Note how spikes are clustered into bursts which drift over time and occasionally disappear or abruptly jump to a new location (examples of this behavior can also be seen in Figure 1). **B)** PCA of 7 hand-crafted features revealed moderate separation of drifty bursting cells from other cells. **C)** We attempted to use Soft DTW to align drifting cells and extract features that describe the way they warp over time. Left: Success of the algorithm. The drifting spikes are aligned in a manner that matches what we expect by eye. Right: Failure of the algorithm. No spatial structure is apparent in the raw data, but the algorithm creates spatial structure in the aligned data. Because of failures such as this one, we could not separate drifty bursting cells from other cells using Soft DTW features.

Contributions

Malcolm Campbell collected the data used in this project prior to the start of the class as part of his research in Lisa Giocomo's laboratory (Neurobiology). Mark Plitt wrote the code for training the classifiers and running cross-validations, and implemented soft dynamic time-warping. Malcolm Campbell implemented data augmentation, generated gain manipulation features, and identified and labeled "drifty bursting cells." Both authors contributed to making the poster and writing the final report.

References

1 Hartley, T., Lever, C., Burgess, N. & O'Keefe, J. Space in the brain: how the hippocampal formation supports spatial cognition. *Philosophical Transactions of the Royal Society B: Biological Sciences* **369**, 20120510, doi:10.1098/rstb.2012.0510 (2014).

2 Domnisoru, C., Kinkhabwala, A. A. & Tank, D. W. Membrane potential dynamics of grid cells. *Nature* **495**, 199-204 (2013).

3 Hafting, T., Fyhn, M., Molden, S., Moser, M. B. & Moser, E. I. Microstructure of a spatial map in the entorhinal cortex. *Nature* **436**, 801-806 (2005).

4 Solstad, T., Boccara, C. N., Kropff, E., Moser, M. B. & Moser, E. I. Representation of geometric borders in the entorhinal cortex. *Science* **322**, 1865-1868 (2008).

5 Cuturi, M. & Blondel, M. Soft-DTW: A Differentiable Loss Function for Time-Series. *arXiv* **1703** (2017).

6 Sakoe, H. & Chiba, S. Dynamic Programming Algorithm Optimization for Spoken Word Recognition. *IEEE Transactions on Acoustics, Speech, and Signal Processing* **ASSP-26**, 43-49 (1978).

7 Petitjean, F., Ketterlin, A. & Gançarski, P. A global averaging method for dynamic time warping, with applications to clustering. *Pattern Recognition* **44**, 678-693 (2011).

8 McNaughton, B. L., Battaglia, F. P., Jensen, O., Moser, E. I. & Moser, M. B. Path integration and the neural basis of the 'cognitive map'. *Nat Rev Neurosci* **7**, 663-678 (2006).