

Towards an ensemble-level view of place encoding in hippocampal neuronal populations

D. Tomàs Cuesta^{1,2}, P. Jercog²

¹ Universitat Pompeu Fabra, Barcelona

² Institut d'investigacions Biomediques Agust Pi i Sunyer (IDIBAPS), Barcelona

Abstract

The study of Place Cells, hippocampal neurons tuned to spatial locations in the environment, is central to elucidate how the brain encodes and retrieves spatial information. Advances in genetic and imaging technologies have allowed keeping track of the dynamics of large ensembles of Place Cells across multiple days in mice. As the brain processes information at the neuronal population level, novel recording techniques such as in-vivo calcium imaging have the potential to unveil the mechanisms underlying the dynamics of place coding. However, with new recording paradigms comes the need to standardize and optimize the processing and first analysis stages of the data. In this work, we present our efforts in building a pipeline to process, extract, filter, track and analyze Place Cells from mice calcium imaging recordings in a linear-track experiment. To validate the pipeline, we show accurate prediction of the animal actions from the processed neural recordings. Finally, building on the previous steps, we present some tentative results on Place Cell turnover and the relation between predictive-accuracy and noise correlations.

1. Introduction

Of all brain capabilities, the ability to cognitively situate ourselves and our goal locations in the context of a physical environment is a highly sophisticated one, yet perhaps one of the many we take largely for granted. It is however difficult to imagine how the brain performs navigation and integrates other spatial functions without having an internal abstraction of the space. Finding and understanding this representation has been an ongoing task in psychology and neurosciences for numerous decades.

A milestone in this journey was the discovery in 1971 of Place Cells (PC) by John O'Keefe and John Dostrovsky [1]. PC are hippocampal neurons which show firing preference for specific spatial locations in the environment, and are now believed to be the main actors in how the brain encodes and retrieves spatial information. Across the years, electrophysiological recordings have enabled a profound characterization of PC. For example, it is known that, upon entering a new environment, they are tuned at the order of minutes [2], and remain robustly maintained from weeks to months, invariant under nonspatial changes [3].

However, due the single-cell nature of the recording techniques used, fundamental properties related to the information encoding and function of PC in the context of a neuronal population remain largely unexplored. Overcoming this limitation has motivated the adoption of

novel techniques allowing population-wide recordings in in-vivo animals. Underlying the aim for these analyses is the idea that information processing in the brain is carried out at the ensemble level. Single cell recordings, while invaluable from a physiological standpoint, provide but a limited picture of the mechanisms and computations involved in the different brain processes [4]. Furthermore, shifting the framework to single-trial, multiple neurons opens the door to the fields of Information Theory and Machine Learning, which could ultimately lead to novel insights on how brain encodes and processes spatial information.

One of the recording methods that have emerged in recent years consists on imaging the transient calcium signals in neurons, which act as surrogates of electrical activity. Upon calcium binding, genetically engineered sensors are able to produce fluorescence traces, which can then be captured by different types of cameras or microscopes [5].

In this work, we will focus on in-vivo one-photon calcium imaging recordings of sensor GCaMP6 [6] transfected PC, a method which achieves neuron population yields of >100 cells, albeit with some caveats, such as low temporal resolution, limited SNR, imaging artifacts, and lowpassing effect of spike dynamics.

While imaging calcium in cells opens new possibilities in studying neuronal circuit dynamics at ensemble level, changing the recording paradigm, together with dealing with the mentioned limitations, has also radically altered the data processing methods and conventions. There is now an urgent need for standardizing and optimizing the processing and first analysis stages of calcium imaging data.

The bulk of this work concerned designing, building, and automating a processing pipeline, in order to transform the raw calcium images into usable neurological data. We then coupled this processed data with the corresponding behavioral recordings of mice performing a simple linear track task, and validated the constructed pipeline by showing accurate behavioral prediction of the animal movement. Finally, we attempted to highlight the potential of ensemble level Place Cell Calcium Imaging recordings by investigating aspects related to population turnover and information coding reliability in the ensemble.

2. Processing the data

We divide our calcium data pipeline into a series of steps. For each day, pre-processing (figure 1A-C), concerns applying several transformations to the data, in order to make it suitable for cell extraction algorithms, in the extraction stage (figure 1D). After cells are extracted, they must then be filtered to remove false positives (figure 1E). Afterwards, the resulting neurons can be put in correspondence with the recordings of other days in the aligning stage (figure 1F). Besides, the fluorescence traces of the extracted cells (figure 1G) may be processed to extract the underlying spike train (figure 1H), using a de-convolution approach.




Figure 1. Overall processing pipeline, from the raw camera footage to the extracted cells, traces and spike trains. To build the pipeline, code from multiple sources; ImageJ and batch processing [7] for B,C,F, CELLMAX [8] and ICA [9] for D,G, OASIS [10] for H, was integrated.

2.1. Preprocessing

Using the miniscope software (Inscopix decompressor from Inscopix, inc.) we generated 3D image stacks from the miniscope output, corresponding to 20Hz movies of $\sim 560 \times 560 \mu\text{m}$ of tissue patch. The resulting dataset (see figure 1A, for a snapshot), is however not suitable for cell extraction, due to large file size ($>10\text{gb}$ for $\sim 20\text{min}$ recordings), external interferences, motion jitter and lens artifacts. To target those first issues we followed the Schnitzer lab pipeline [7] and used a set of MATLAB functions from the miniscope analysis toolbox written by Biafra Ahanonu. Using bi-linear interpolation, we spatially downsampled the files by a factor of 4, yielding $250 \times 250 \times \text{frames}$ matrices and effectively reducing the file size by a factor of 16. This allows holding the data in RAM memory for the most memory-demanding operations, like filtering. Next, motion jitter, which comes from the relative motion between the miniscope and the tissue, had to be removed. This task can be achieved using registration methods, which align the different

frames of the stack to a reference image by using landmarks known to be spatially fixed. We applied the Turboreg algorithm [11] from the ImageJ software (imagej.nih.gov), which uses a pyramid approach (starting the registration with a low-resolution image, progressively increasing the resolution to avoid getting stuck in a local minima) based on image intensity to align the different frames to the reference. Before the registration, the image was lowpassed to reduce high frequency noise, which further reduces local minima impact. To lowpass the movie, each frame was convolved with a Gaussian filter ($\sigma = 3, 80 \text{ pixels}^2$).

Besides motion, raw calcium imaging may also suffer from photo-bleaching (loss of fluorophore effectivity), external light fluctuations, tissue pulsation and other unwanted external effects. These issues can be minimized by using $\Delta F/F_0$, a transformation on the data that removes the baseline of the movie across frames [5]. In our definition, it corresponds to:

$$F'(t)_{ij} = \frac{F(t)_{ij}}{\frac{1}{N} \sum_{t=1}^N F(t)_{ij}} - 1$$

Where $F(t)_{ij}$ is the fluorescence at frame t and pixels i, j , and N is the number of frames in the movie. Then, if $F(t)_{ij}$ has the same value as its time average, its new value $F'(t)_{ij}$ maps to 0. As PC firing is sparse [2], removing the baseline will highlight the temporal calcium differences. As such, $\Delta F/F_0$ is a measure of the activity of the neurons (figure 1C).

2.2. Cell Extraction

Cell extraction is the process in which neurons, the sources of the fluorescence traces are inferred from the pre-processed movie. While it is technically possible to manually assign a ROI (region of interest) to the perceived fluorescence spikes, the length of the movies and multiple recording days, together with the variability of intensities and shapes of the fluorescence flashes, make an automated, systematic method necessary in practice. We explored two approaches to cell extraction, Independent Components Analysis (ICA) and Expectation Maximization (EM). Both are probabilistic methods rooted in statistical analysis, and generate a set of filters describing the most likely locations and shapes of the different cells, but they use distinct strategies to achieve this.

ICA is a method to find an underlying set of independent basis generating the data [12]. In the context of calcium imaging, the aim is to find the neurons that generated each frame by searching independent basis that maximize spatial and temporal sparseness of the resulting cells and traces [13]. We used the ICA implementation FastICA [9], and prior to extraction, reduced the dimensionality of the data using Principal Component Analysis (PCA), a method for finding a set of orthonormal basis minimizing reconstruction error [14].

Expectation minimization is a general iterative technique to solve maximum loglikelihood estimate problems in the presence of unknown latent variables [12]. The algorithm

we used for the extraction is CELLMax [8], and works in the following manner: First, a set of potential underlying neurons are initialized as a uniform grid of 2D Gaussians. These sources generate the observed frames via a quantum physical Bose-Einstein model that predicts the fluorescence photon counts. Comparing with the actual movie, the source image is modified to maximize the expected likelihood of the observed data.

As most datasets do not completely fulfill the sparseness assumptions required by ICA, we expected to EM have an edge in versatility and extraction accuracy. To test this, we computed several salient features for the extracted neurons of 19 recorded sessions. The extractions were performed with ICA and EM at different initial conditions and minimal iteration criteria, to check for convergence. We compared several features, relating the extracted cell shapes with the actual movie spike transients (correlation and overlap), comparing the SNR and shape of the transients of the different traces around the spikes, and computing properties of the extracted cells shape (area, solidity, eccentricity). The results revealed rapid convergence of EM with the number of iterations, and similar overlap, solidity, eccentricity and trace shape features between ICA and EM. Having found no conclusive discerning results, EM was used for the rest of the work on convenience grounds. An example of extraction results can be seen on figure 1D.

2.3. Cell Filtering

Despite the sophistication of the extraction algorithms, manual inspection of the resulting potential cells revealed a significant portion of them capturing noise, fragments of other cells, comprising multiple sources or being otherwise unconvincing as proper neurons. Filtering has the goal of getting rid of these false cells while conserving the good ones, in a consistent and objective manner. We explored three approaches to achieve this:

Automatic method: We manually filtered 17,578 cells and used the gathered features described in the previous step, to train several models with the aim to predict whether a cell was valid or not given the array of features. An accuracy around 90% was achieved for the best models (Decision Tree ensembles and Support Vector Machine), suggesting the viability of a fully automated filtering approach.

Semi-automatic method: We set conservative bounds on the features aiming to remove the maximum number of false cells while losing the least possible valid cells, in order to constrain manual selection latter on. Using this approach, we were able to remove a 41% of cells a priori, with minimal valid cell loss (~1%), reducing the burden of manual sorting considerably.

Manual, event-based method: An issue with our feature-based methods is their lack of robustness to overlapping neurons or false spikes, which mislead the extracted statistics. To solve this, we developed a software to, with human supervision; evaluate each event of a spike separately to assess the overall validity of the cell. To set an event as valid, we asked for a minimum

overlap with the extracted cell shape, penalizing for missing and overflowing areas.

2.4. Cell Alignment and spike train extraction

Aligning concerns finding the correspondence between the filtered cells of two or more days. As the miniscope was removed and reintroduced, we expect to find a 2D transformation (rotation dilation and translation) that relates the cells of neighboring days. In this sense, aligning the cells is similar to frame registration in the preprocessing stage, though performed at the cell-map level, and allowing for larger (coherent) displacements, and we used the Turboreg algorithm [11] once again.

Finally, to recover the underlying spike train from the cell fluorescence traces we used a deconvolution approach by Vogelstein [10], which assumes the spikes come from an exponential distribution and produce a decreasing exponential transient in the Ca^{2+} . In this and posterior steps, we used the original sampled (20Hz) movies.

3. Pipeline validation

The experimental paradigm in which the data was gathered consists on the task of a mouse switching between two water ports on the extremes of a linear platform. The ports were triggered by the animal licking in an alternate manner, which meant that no water was awarded for sticking to one end of the track. Mice suffered no other punishment from performing badly, and were trained in advance. Two mice populations were used in the experiment, one performing the task every 2 days, and the other every 5 days.

We extracted the animal trajectory from behavioral videos by discarding all pixels brighter than the mouse, subtracting background to isolate moving parts and applying a morphological opening and erosion, which removes small unconnected regions. The position of the mouse in a given frame was taken as the centroid of the largest remaining region, an approach we observed to be reliable enough in practice. Outliers were removed and replaced using a piecewise cubic interpolation of the values at neighboring points, and the traces were smoothed using a zero phase digital filter with uniform weights and a 10 point window. Movement in the perpendicular direction to the track was discarded for being negligibly small.

To validate the pipeline, we binned the track in 20 segments (~5cm bins) and trained a Bayesian decoder to predict, for each camera frame, the bin the mouse was occupying, using the whole population of fluorescence traces. For each recording session, the decoder was trained with 90% of the data and tested with the remaining 10%. This was repeated 10 times with different segments. The error in the decoding was taken as the mean root squared difference between the true and predicted trajectories, averaged over the 10 bootstraps. Taking only the frames in which there was movement across the track (velocity > 4cm/s), the mean error was around 0.9bins (figure 2, green), indicating we were successfully extracting positional information from the data.




Figure 2. Predictions on the trajectory of the mouse along the track, with the slow frames removed.

4. Other results

Expanding on the decoding applications, we took a preliminary look on some aspects of PC ensemble coding.

4.1. Decoder accuracy and correlations

We wanted to see whether the accuracy of our decoder was limited by correlations between neurons. We shuffled among trials (a trial being one frame) the frames where the mouse occupied the same spatial bin, for each cell of the population separately, in order to remove interneuron correlations across trials. This resulted in decreased decoding error (figure 2, red).

To explore this finding, which could link to aspects of information coding reliability in PC, we computed the Place Fields (PF), the desired zones of firing of each neuron. For this, we extracted spike trains to find firing rates for the 20 spatial bins, normalizing by mean firing rate and animal's bin occupancy. We asked for PF significance by demanding a significant decrease in the mutual information between the firing rates and the mouse position when time-shifting the spike trains. With the PF computed, we then measured pair-wise correlations between cells as a function of the distance between their PF, and found that pairs with nearby PFs were positively correlated and that correlations faded away for pairs with distant PFs. This spatial pattern could deleterious for population decoding.

4.2. Turnover

It has previously been reported [7] that PC coding is highly dynamic, with different subsets of PC activated for a particular environment in each day. We investigated whether this turnover depended on the frequency that the animal had been exposed to the task. We computed decoding accuracy across days for two mice; Mouse A, which performed the task every 2 days and Mouse B, which performed the task every 5 days. We tracked the cells in the first recording day across the sessions, and measured how the decoding accuracy degraded when using a decoder trained in the first day. We also checked accuracy when training in each day to account for changes in the overall decoding performance due to photo-bleaching or other phenomena. While the results (figure 3) suffer from outliers (2nd session of B), they hint at a higher turnover for Mouse A than B, indicating

turnover increases when a spaced learning frequency is higher. Future work in this direction will be required to test this and related hypothesis.




Figure 3. Decoding accuracy across days reveals turnover.

5. References

- [1] O'Keefe J, Nadel L. The Hippocampus as a Cognitive Map. Oxford University Press 1978 (ISBN: 978-0198572060)
- [2] Wilson MA, McNaughton BL. Dynamics of the hippocampal ensemble code for space. *Science*, vol 261, sup 5155, 1993, pp 1055–1058 (ISSN: 1095-9203)
- [3] Lenck-Santini PP, Rivard B, Muller RU, Poucet B. Study of cal place cell activity and exploratory behavior following spatial and nonspatial changes in the environment. *Hippocampus*, vol 15, sup 3, 2004, pp 356–369. (ISSN: 1098-1063)
- [4] Quiroga R.Q, Panzeri S. Extracting information from neuronal populations: information theory and decoding approaches, *Nature Neuroscience*, vol 10, sup 3, 2009, pp 173–185. (ISSN 1546-1726)
- [5] Grienberger C, Konnerth A. “Imaging calcium in neurons,” *Neuron*, vol 73, sup 5, 2012, pp 862–885. (ISSN: 0896-6273)
- [6] Chen TW, Wardill TJ, Svoboda K, Kim DS. Ultrasensitive fluorescent proteins for imaging neuronal activity, *Nature*, vol 499, sup 7458, 2013, pp 295–300. (ISSN: 0028-0836)
- [7] Ziv Y, Burns LD, Schnitzer MJ. Long-term dynamics of cal hippocampal place codes, *Nature Neuroscience*, vol 16, sup 3, 201, pp 264–266, (ISSN 1546-1726)
- [8] Kitch LJ. Machine learning meets mammalian learning: Statistical tools for large-scale calcium imaging and the study of changing neural codes. *Stanford Doctoral Thesis*, 2015.
- [9] Hyvarinen A, Oja E. Independent component analysis: algorithms and applications, *Neural Networks*, vol 13, sup 4–5, 2000, pp 411–430. (ISSN: 0893-6080)
- [10] Vogelstein JT, Packer AM, Paninski L. Fast nonnegative deconvolution for spike train inference from population calcium imaging, *Neurophysiology*, vol 104, sup 6, 2010, pp 3691–3704. (ISSN: 1573-9007)
- [11] Thevenaz P. A pyramid approach to subpixel registration based on intensity, *IEEE transactions on image processing*, vol 7, sup 1, 1998, pp 27–41. (ISSN: 1057-7149)
- [12] Murphy KP. Machine Learning, A Probabilistic Perspective. MIT press 2012 (ISBN: 978-0262018029)
- [13] Mukamel E.A., Nimmerjahn A., Schnitzer M.J., Automated analysis of cellular signals from large-scale calcium imaging data, *Neuron*, vol 63, sup 6, 2009 pp 747–760. (ISSN: 0896-6273)
- [14] Bishop CM. Pattern Recognition and Machine Learning. Springer 2006 (ISBN: 978-0387310732)