Cell Assembly Model for Retinal Ganglion Cell population

Christopher Warner and Friedrich Sommer

January 25, 2019

Abstract

We use a Noisy-Or model, previously applied to disease modelling, to learn latent structure in observed spiking patterns from a population of rat retinal ganglion cells in response to white noise and natural movie stimulus. The probabalistic generative model, fit with the expectation maximization algorithm, captures approximately repeated groups of neurons often coactive together without requiring exact repeats. These cell assemblies capture the fine-temporal correlations in population spike trains that differ across stimulus conditions but are not necessarily stimulus locked. We discuss the spatial structure of cells in cell assemblies and what features of the stimulus is represented by a cell assembly above and beyond the union of the representations of its member cells.

1 Introduction

Retinal ganglion cells (RGCs) encode local visual stimulus features in their spike rates and rate-based models have captured retinal spike trains in response to white noise stimulus. However, when presented with more structured stimulus, independent rate-based models fail to capture neural responses. Models that take into account pairwise interactions between cells fair better with more realistic, structured data [1] but still fall short. We explore a Noisy-Or model [2] to capture more-than-pairwise interactions in the activity of latent variables. These "cell assemblies" represent groups of neurons that fire together on a fine temporal scale more frequently than would be expected if they acted like inhomogeneous poisson processes given rates. Here we learn a probabalistic generative cell assembly model from population spiking data in rat RGCs presented white noise and natural movie stimuli.

2 Probabalistic Generative Cell Assembly Model

Given data of binary "spike words" observed in a population of N neurons, $\vec{y} \in \{0,1\}^N$, we seek to represent the data by model with latent variables $\vec{z} \in \{0,1\}^M$. Each latent variable describes the activity of a cell assembly, that is, a group of cells that frequently activate synchronously within a small time window. The model has three types of variables that are learned: The vector $P_i = p(y_i = 0|z = 0) \in [0,1]^N$ represents the probabilities that individual cells are *silent*, given no cell assemblies are active. The matrix $P_{ia} = p(y_i = 0|z_a = 1) \in [0,1]^{NxM}$ contains the probabilities that cell i is silent when assembly a is active. The value $Q = p(z_a = 1) \in [0,1]^1$ is a scalar value indicating the probability that a cell assembly is active, assuming that the activities of different assemblies are independent. We compute the MAP estimate for the probability of individual cell's activity with a given $\vec{z}$ as:

$$p(y_i, \vec{z}) = \binom{M}{|\vec{z}|} \cdot Q^{|\vec{z}|} \cdot (1-Q)^{\left(M-|\vec{z}|\right)} \left[P_i^{\left(1-\frac{|\vec{z}|}{M}\right)} \prod_{a-1}^{M} (P_{ia})^{z_a}\right]^{(1-y_i)} \left[1 - P_i^{\left(1-\frac{|\vec{z}|}{M}\right)} \prod_{a-1}^{M} (P_{ia})^{z_a}\right]^{y_i} \tag{1}$$

Model parameters are learned by alternating Expectation Maximization (EM) where, for an observed $\vec{y}$, we first infer $\vec{z}$ with fixed model parameters and then adjust model parameters using gradient descent on the error between inferred and observed $\vec{y}$s.

3 Retinal data and model application

After verifying the model's behavior on synthetic data, we apply it to spike trains recorded from a population of 55 Off Brisk Transient *in vitro* rat retinal ganglion cells recorded from retina using multielectrode array. Spike trains are recorded with high (sub-millisecond) temporal precision from retinal tissue is presented with 200 trials of 5 second clips of white noise and natural movie stimuli. (reference Litke and Field 2004? or other paper?) We extract 500,000 "spike words" from population spike trains by recording groups of cells that fire within 3ms of eachother across all time points and all trials. We then randomly take 50,000 samples from this data distribution spike words that include at least 6 cells (that is, $|\vec{y}| \geq 6$), assuming that observed spike words caused by cell assemblies will contain multiple cells. The model is learned on these 50,000 samples.

4 Results, Evaluation and Discussion

From relatively few data observations, the model learns approximately 20 "non-trivial" cell assemblies (that is, containing more than one cell) from multi-cell spiking patterns in response to natural movies. It does not discover similar structured multi-cell spiking patterns in response to white noise stimulus (not shown). The learned model fit to RGC responses to natural movies is shown in figure 1. Column-sums of the thresholded P_{ia} matrix ($\Theta = 0.5$) give a measure of how many cells participate in each cell assembly (blue triangles in middle left panel). While many cells participate in only one cell assembly, but some particpate in multiple - presumably spiking in different groups under different contexts (blue triangles in top center). In the bottom panel, we raster the activity of the 10 most inferred non-trivial cell assemblies on top of individual cell PSTHs. The legend indicates the corresponding column of the P_{ia} matrix. In addition to these analyses, we will present further results on cell assembly receptive fields, that is the spatial relationship of cells in a given assembly and how the stimulus that drives a cell assembly differs from the stimulus that drives its cells individually. We will also determine the significance of found cell assemblies by comparing them to null models based on rate coding with neurons acting as inhomogeneous poisson processes.

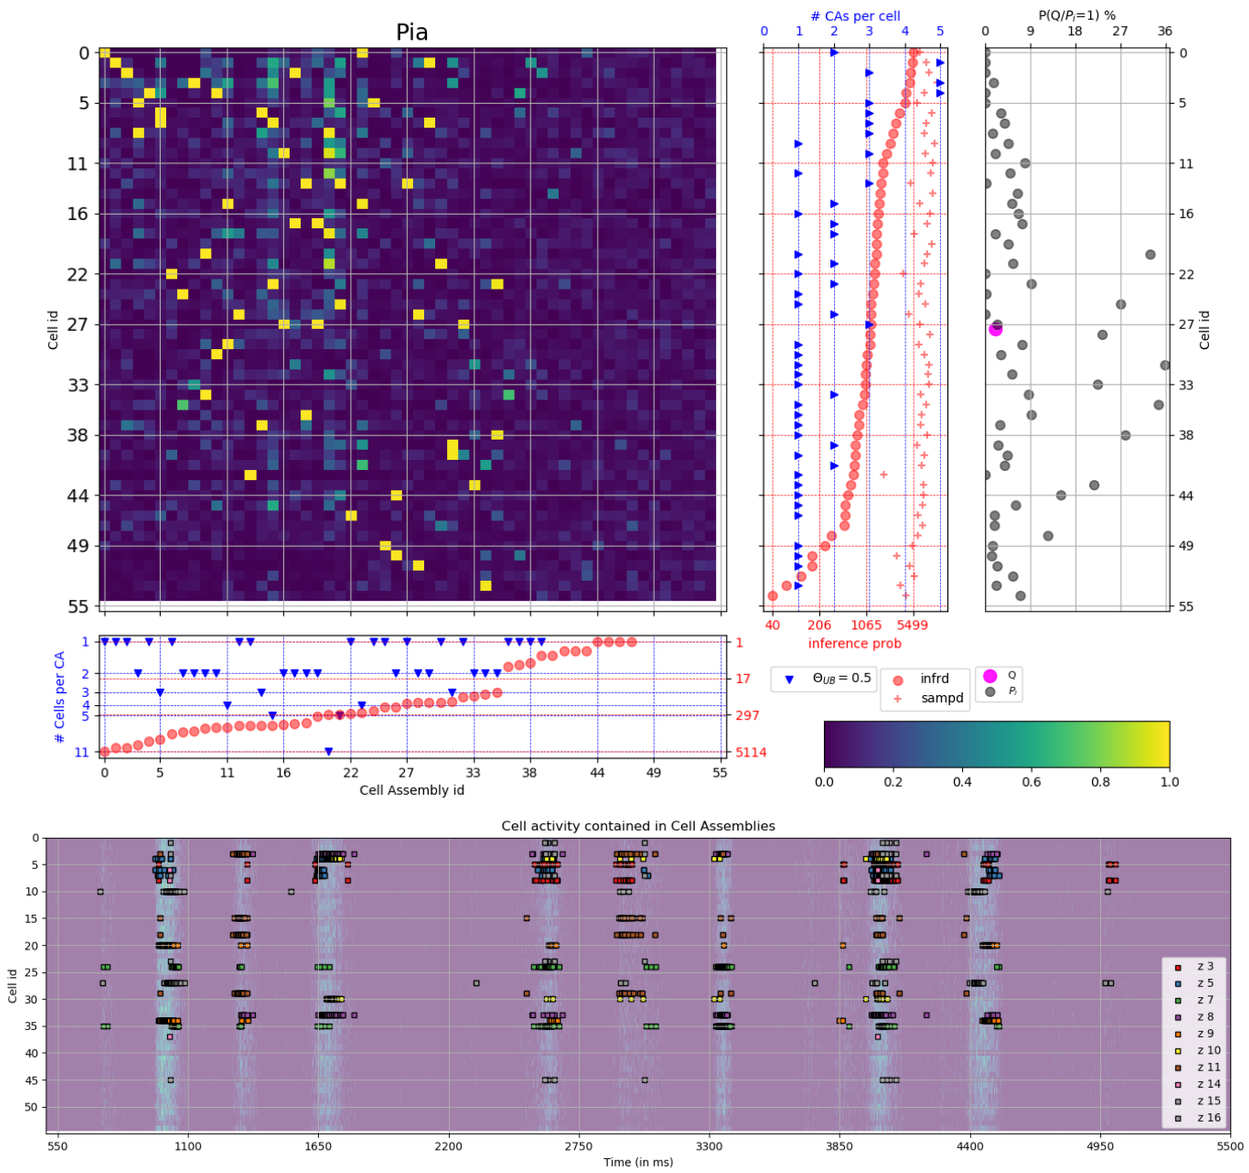

Figure 1: Probabalistic generative model of cell assemblies learned on retinal ganglion cell (RGC) spike trains presented with natural movies: *Top Left:* Columns in the P_{ia} matrix indicate cell assemblies and large values within indicate cells belonging to a given assembly. *Two flanking subpanels*: Blue triangles show the number of cells in each assembly and the number of assemblies that each cell participates in. Red circles indicate the number of times an assembly or a cell is inferred. *Top right:* P_i values (black circles) quantify the innate "noisiness" of neurons and Q (pink circle) indicates the probability a cell assembly is active. *Bottom:* Color coded cell assembly activity overlayed on individual cell PSTHs indicates observed activity explanable by latent structure.

References

[1] Jonathan W Pillow, Jonathon Shlens, Liam Paninski, Alexander Sher, Alan M Litke, EJ Chichilnisky, and Eero P Simoncelli. Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature*, 454(7207):995, 2008.

[2] David Heckerman. A tractable inference algorithm for diagnosing multiple diseases. In *Machine Intelligence and Pattern Recognition*, volume 10, pages 163–171. Elsevier, 1990.