A Two-Dimensional Seperable Random Field Model of Within and Cross-Trial Neural Spiking Dynamics

Yingzhuo Zhang, Noa M Shinitski, Stephen Allsop, Kay Tye and Demba Ba

Summary

A fundamental problem in neuroscience is to characterize the dynamics of spiking from the neurons in a circuit that is involved in learning about a stimulus or a contingency. A key limitation of current methods to analyze neural spiking data is the need to collapse neural activity over time or trials, which may cause the loss of information pertinent to understanding the function of a neuron or circuit. We introduce a new method that can determine not only the trial when a neuron learns a contingency, but also the latency of this learning with respect to the onset of a cue within that trial. The backbone of the method is a separable two-dimensional (2D) random field (RF) model of neural spike rasters, in which the joint conditional intensity function of a neuron over time and trials depends on two latent Markovian state sequences that evolve separately but in parallel. We develop efficient statistical and computational tools to estimate the parameters of the separable 2D RF model, and apply these to data collected from neurons in the anterior cingulate cortex (ACC) in an experiment designed to characterize the neural underpinnings of the observational learning of fear in mice. We find that the trial at which ACC neurons exhibit a conditioned response to the auditory cue (learning trial) is robust across cells, occurring 3 to 4 trials into the conditioning period. Following the learning trial, we find that the time with respect to cue onset when we observe a significant change in neural spiking compared to baseline activity (learning time) varies significantly from cell to cell, occurring between 20 to 600 ms after cue onset. Overall, the separable 2D RF model provides a detailed characterization of the dynamics of neural spiking of ACC neurons during observational learning of fear.

Additional Details

The dynamics of neural spiking within and across trials are modeled as a separable 2D RF

$$\begin{cases} x_k = x_{k-1} + \epsilon_k, \epsilon_k \sim \mathcal{N}(0, \sigma_\epsilon^2) \\ z_r = z_{r-1} + \delta_r, \delta_r \sim \mathcal{N}(0, \sigma_\delta^2) \\ \log \frac{\lambda_{k,r}\Delta}{1-\lambda_{k,r}\Delta} = x_k + z_r \\ \Delta N_{k,r} | x_k, x_r \sim \text{Bernoulli}(\lambda_{k,r}\Delta) \end{cases}$$

where $r = 1, \cdots, R$ is the trial index and $k = 1, \cdots, K$ the time index within a trial. The 2D separable RF model expresses the trial-dependent conditional intensity function (CIF) of the neuron in discrete-time as $\lambda_{k,r}\Delta \approx e^{z_r} \cdot e^{x_k}\Delta$, that is, as the product of a within-trial component e^{x_k} in units of Hz (spikes/s) and a unitless quantity e^{z_r}. For a given trial r, e^{z_r} represents the excess spiking rate at that trial above what can be expected from the within-trial component. We call e^{z_r} the cross-trial component of the CIF. The within and cross-trial components are functions of two independent state sequences, $(x_k)_{k=1}^K$ and $(z_r)_{r=1}^R$, that evolve smoothly according to a first-order random walk. The parameters σ_ϵ^2 and σ_δ^2, which govern the smoothness of $(x_k)_{k=1}^K$ and $(z_r)_{r=1}^R$, must be estimated.

Parameter estimation for the 2D separable random field: The evolution of the latent states in the separable 2D RF cannot be written as a simple random walk in one dimension (e.g. time). The separable 2D RF is therefore not amenable to classic methods for estimating a state-space model from point process data [1]. We derive a Monte Carlo Expectation-Maximization algorithm to maximize the marginal likelihood of the data from the separable 2D RF with respect to the parameters σ_ϵ^2 and σ_δ^2. In the E-step, we leverage the Polya-Gamma latent representation of Bernoulli random variables [2] to generate Gibbs samples from each of the state sequences. In the M-step, we use the Monte-Carlo samples from the E-step to update the parameters σ_ϵ^2 and σ_δ^2.

Advantage over non-parametric methods: In addition to the need to collapse spiking activity over time or trials, two common pitfalls of non-parametric methods (e.g. Wilcoxon rank sum test) are their

reliance on large sample assumptions to justify comparing neural spiking rates, and the need to correct for multiple comparisons. The key advantages of our model-based approach stem from the fact that it yields a characterization of the joint posterior (over all trials and times within a trial) distribution of the instantaneous rate of spiking of as a function of both time and trials given the data. This not only obviates the need to correct for multiple comparisons, but also enables us to compare the instantaneous rate of any two trial time pairs at the millisecond resolution, where non-parametric methods break down because the sample size is 1.

Detailed changes in neural spiking dynamics: Figure 1 summarizes the result of the 2D separable RF model applied to 45 2-second-long trials of data from an ACC neuron. The cue occurs 500 ms into a trial. The left panel shows the raster from this neuron, along with the estimated within-trial (bottom) and cross-trial (left) components. The within and cross-trial components indicate significant changes respectively in response to the cue and to conditioning. In the right panel, we quantify using Monte-Carlo samples the significance of these changes by computing the probability that the spiking rate at a given time during one of the conditioning trials (trials 16 through 45) is bigger than the baseline spiking rate at that trial (region A, panel (c)) *and* the average spiking rate at the same time during the habituation period (region B, panel (c)). We determine the trial and time pair when the neuron first shows a conditioned response to the cue (learning trial and time) as the first trial time pair for which this probability remains consistently above a certain significance level (e.g. 0.9). The unit in this figure begins to exhibit a conditioned response at trial 20, 614 ms following the cue.

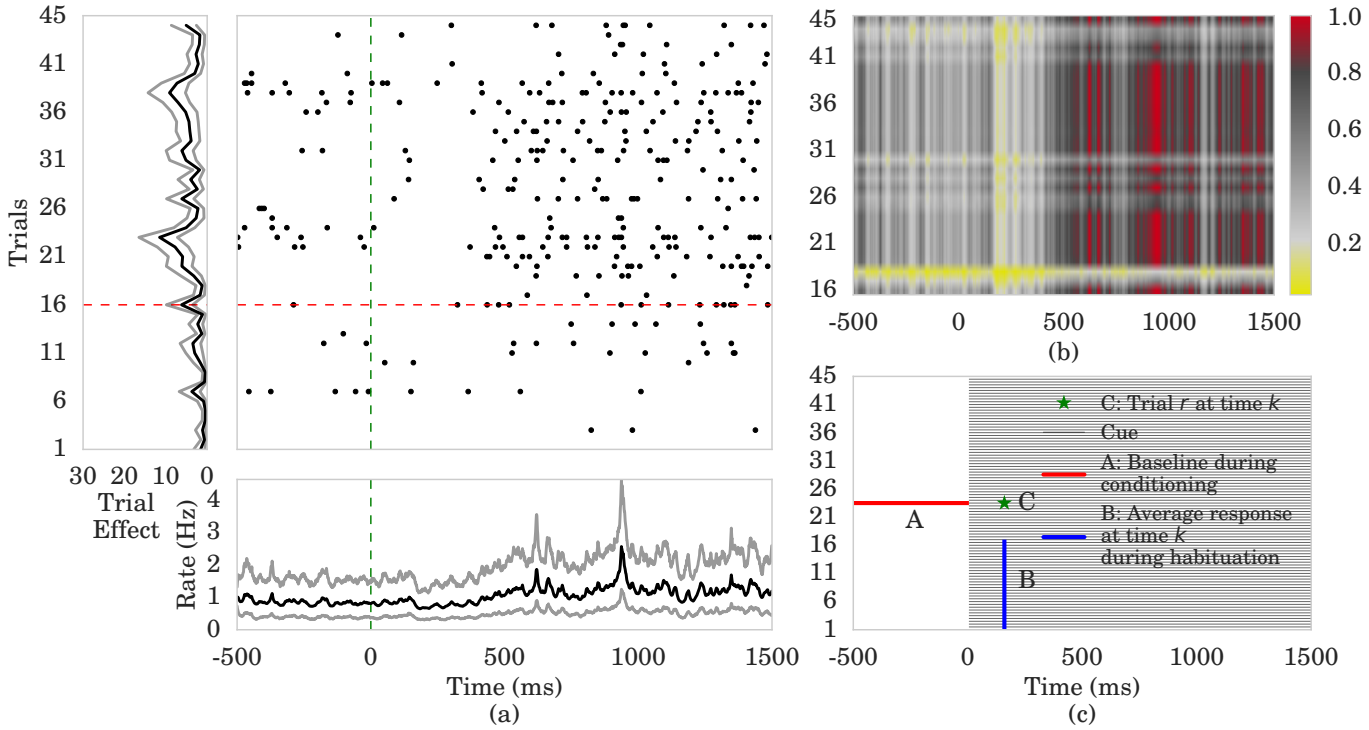

Figure 1: Characterization of the trial and time when an ACC neuron exhibits a conditioned response. (a) Raster and posterior estimates of within and across trial effects. The horizontal red line indicates the beginning of conditioning. The vertical green line indicates cue onset. (b) Comparison of spiking rate at a given trial and time to the baseline rate at the same trial and the average rate during habituation at the same time. (c) Regions compared in (b).

Support: Demba thanks the Alfred P. Sloan Foundation. Kay Tye thanks the McKnight Foundation and the NIH (Grant R01-MH102441-01).

[1] Anne C Smith and Emery N Brown. Estimating a state-space model from point process observations. *Neural Computation*, 15(5):965–991, 2003.

[2] Nicholas G Polson, James G Scott, and Jesse Windle. Bayesian inference for logistic models using pólya–gamma latent variables. *Journal of the American statistical Association*, 108(504):1339–1349, 2013.