

MIT Open Access Articles

*Estimating a Separably Markov Random
Field from Binary Observations*

The MIT Faculty has made this article openly available. **Please share**
how this access benefits you. Your story matters.

Citation: Zhang, Yingzhuo et al. "Estimating a Separably Markov Random Field from Binary Observations." *Neural Computation* 30, 4 (April 2018): 1046–1079 © 2018 Massachusetts Institute of Technology

As Published: http://dx.doi.org/10.1162/neco_a_01059

Publisher: MIT Press

Persistent URL: <http://hdl.handle.net/1721.1/115400>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



Estimating a Separably Markov Random Field from Binary Observations

Yingzhuo Zhang

yingzhuozhang@g.harvard.edu

*John A. Paulson School of Engineering and Applied Sciences,
Harvard University, Cambridge, MA 02138, U.S.A.*

Noa Malem-Shinitski

noa@bccn-berlin.de

Technische Universität Berlin, Berlin 10623, Germany

Stephen A. Allsop

sa3@mit.edu

Kay M. Tye

kaytye@mit.edu

*Department of Brain and Cognitive Sciences and Picower Institute
for Learning and Memory, Cambridge, MA 02139, U.S.A.*

Demba Ba

demba@seas.harvard.edu

*John A. Paulson School of Engineering and Applied Sciences,
Harvard University, Cambridge, MA 02138, U.S.A.*

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-to-trial dynamics that accompany the learning of a contingency by a neuron, but also the latency of this learning with respect to the onset of a conditioned stimulus. 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. Classical tools to estimate state-space models cannot be applied readily to our 2D separable RF model. We develop efficient statistical and computational tools to estimate the parameters of the separable 2D RF model. We apply these to data collected from neurons in the prefrontal cortex in an experiment

designed to characterize the neural underpinnings of the associative learning of fear in mice. Overall, the separable 2D RF model provides a detailed, interpretable characterization of the dynamics of neural spiking that accompany the learning of a contingency.

1 Introduction

A fundamental problem in the analysis of electrophysiological data from neuroscience experiments is to determine the trial, and time within that trial, when a neuron or circuit first exhibits a conditioned response to a stimulus. This is a challenging problem because neural spike rasters resulting from such experiments can exhibit variability both within a given trial and across trials (Czanner et al., 2008). Fear conditioning experiments (Allsop, Vander Weele, Wichmann, & Tye, 2014) are a prime example of a scenario when this situation arises: a neutral stimulus, present across all trials of an experiment, gives rise to stereotypical within-trial spiking dynamics, while the associated aversive stimulus leads to changes in spiking dynamics across a subset of the trials.

State-of-the-art methods for analyzing neural spike rasters fall primarily within two classes. The more pervasive class of such methods neglects the inherent two-dimensional nature of neural spike rasters by aggregating the raster data across either time or trials and subsequently applying techniques applicable to one-dimensional signals (Smith & Brown, 2003; Zammit-Mangion, Sanguinetti, & Kadirkamanathan, 2012; Yuan, Girolami, & Niranjan, 2012; Scott & Pillow, 2012). In contrast to these one-dimensional methods, two-dimensional methods model both the within- and cross-trial dynamics of neural spiking (Czanner et al., 2008; Rad & Paninski, 2010). Within the class of one-dimensional methods, the past decade has seen a growing interest in approaches based on state-space models of neural spiking activity. These approaches treat neural spiking data as realizations of a stochastic point process whose conditional intensity function obeys a stochastic smoothness constraint in the form of a Markov process followed by a nonlinearity. The main challenge is to estimate the parameters of the model, and various solutions have been proposed toward this end (Smith & Brown, 2003; Zammit-Mangion et al., 2012; Yuan et al., 2012; Scott & Pillow, 2012). The main drawback of one-dimensional approaches applied to the analysis of neural spike rasters is the need, preceding analysis, for aggregation across one of the dimensions. Among one-dimensional methods, nonparametric methods based on rank tests (e.g., Wilcoxon rank sum test) have been the most popular, primarily due to their ease of application. In addition to the need to collapse neural activity of time or trials, two common pitfalls of nonparametric methods are their reliance on large sample assumptions to justify comparing neural spiking rates and the need to correct for multiple comparisons. For instance, tests that rely on estimates of

the neural spiking rate based on empirical averages are hard to justify when it is of interest to characterize the dynamics of neural spiking at the millisecond timescale. Consider a neural spike raster for which it is of interest to assess differences in instantaneous spiking rates between distinct time/trial pair. At the millisecond timescale, there would only be one observation per time/trial pair, violating the large sample assumptions that such nonparametric methods rely on. To the best of our knowledge, the work of Czanner et al. (2008) remains the most successful attempt to characterize simultaneously the within- and cross-trial dynamics of neural spiking. This approach uses a state-space model of the cross-trial dynamics, in conjunction with a parametric model of the within-trial dynamics. The use of a parametric model for the within-trial dynamics is convenient because it enables the estimation of the model parameters by expectation-maximization (EM), using a combination of point-process filtering and smoothing in the E-step (to fill in the missing cross-trial effect) and an M-step for the within-trial parameters that resembles a GLM (Truccolo, Eden, Fellows, Donoghue, & Brown, 2005). The main drawbacks of this approach are, on the one hand, the high-dimensionality of the state-space model that captures the cross-trial dynamics, and on the other hand, the lack of a simple interpretation, as in the one-dimensional models (Smith & Brown, 2003; Zammit-Mangion et al., 2012; Yuan et al., 2012), for the state sequence. Finally, a two-dimensional approach based on gaussian processes was proposed in Rad and Paninski (2010). One advantage of this approach, which is based on a gaussian process prior of the neural spiking rate surface, is its ability to model the interaction between the two dimensions through the use of a two-dimensional kernel. As is common with kernel methods, it does not scale well to multiple dimensions.

We propose a two-dimensional (2D) random field (RF) model of neural spike rasters, termed separably Markov random field (SMuRF), 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. Conventional methods for estimating state-space models from binary observations (Smith & Brown, 2003; Zammit-Mangion et al., 2012; Yuan et al., 2012) are not applicable to SMuRF. We derive a Monte Carlo EM algorithm to maximize the marginal likelihood of observed data under the SMuRF model. In the E-step, we leverage the Polya-Gamma (Polson, Scott, & Windle, 2013) representation of Bernoulli random variables to generate samples from the joint posterior distribution of the state sequences by Gibbs sampling. A similar strategy was adopted in Scott and Pillow (2012) for a one-dimensional state-space model. The sampler uses a highly efficient forward-filtering, backward-sampling algorithm for which the forward step can be implemented exactly and elegantly as Kalman filter, while the backward step uses Bayes' rule to correct the filter samples. The SMuRF model obviates the need for aggregation across either time or trials and yields a low-dimensional 2D characterization of neural spike rasters

that is interpretable in the sense that the posterior of the two state sequences captures the variability within and across trials, respectively. Moreover, being model based, the SMuRF model, unlike nonparametric methods, yields a characterization of the joint posterior (over all trials and time within a trial) distribution of the instantaneous rate of spiking, thus allowing us to precisely determine the dynamics of neural spiking that accompany the learning of a contingency. To demonstrate this, we apply the model to data collected from neurons in the prefrontal cortex (PFC) in an experiment designed to characterize the neural underpinnings of the associative learning of fear in mice. We find that the trial at which the cortical neurons begin to exhibit a conditioned response to the auditory conditioned stimulus is robust across cells, occurring three to four trials into the conditioning period. We also find that the time with respect to conditioned stimulus onset when we observe a significant change in neural spiking compared to baseline activity varies significantly from cell to cell, occurring 20 to 600 ms after conditioned stimulus onset. These findings are likely reflective of the variability in synaptic strength and connectivity that accompany learning, as well as the location of the neurons in the population.

The rest of our treatment begins in section 2 where we provide our motivation for the SMuRF model, define it, and introduce our notation. In section 3, we present the Monte Carlo EM algorithm for parameter estimation in the SMuRF model, as well as our process for inferring the dynamics of neural spiking that accompany the learning of a contingency by a neuron. Derivations relevant to this section are in the appendix. We present an application to the cortical data in section 4 and conclude in section 5.

2 Notation and the SMuRF Model

We begin this section with a continuous-time point-process formalism of a neural spike raster, characterized by a trial-dependent conditional intensity function (CIF). Then we introduce the SMuRF model, for the discrete-time version of the CIF.

2.1 Continuous-Time Point-Process Observation Model. We consider an experiment that consists of R successive trials. During each trial, we record the activity of a neuronal spiking unit. We assume, without loss of generality, that the duration of the observation interval during each trial is $(0, T]$. For trial r , $r = 1, \dots, R$, let the sequence $0 < t_{r,1} < \dots < t_{r,S_r} < T$ correspond to the times of occurrence of events from the neuronal unit—that is, the times when the membrane potential of the neuron crosses a given threshold. We assume that $\{t_{r,s}\}_{s=1}^{S_r}$ is the realization in $(0, T]$ of a stochastic point process with counting process $N_r(t) = \int_0^t dN_r(u)$, where $dN_r(t)$ is the indicator function in $(0, T]$ of $\{t_{r,s}\}_{s=1}^{S_r}$. A point process is fully characterized by its CIF. Let $\lambda_r(t|H_t)$ denote the trial-dependent CIF of $dN_r(t)$ defined as

$$\lambda_r(t|H_t) = \lim_{\Delta \rightarrow 0} \frac{P[N_r(t + \Delta) - N_r(t) = 1|H_t]}{\Delta}, \quad (2.1)$$

where H_t is the history of the point process up to time t .

We denote by $\{\Delta N_{k,r}\}_{k=1,r=1}^{K,R}$ the discrete-time process obtained by sampling $dN_r(t)$ at a resolution of Δ , $K = \lfloor \frac{T}{\Delta} \rfloor$. Let $\{\lambda_{k,r}\}_{k=1,r=1}^{K,R}$ denote the discrete-time, trial-dependent CIF of the neuron.

2.2 SMuRF Model of Within- and Cross-Trial Neural Spiking Dynamics

Definition 1. Let $\{y_{k,r}\}_{k=1,r=1}^{K,R} \in \mathbb{R}^{K \times R}$ be a collection of random variables. We say that this collection is a separable random field if $\exists \mathbf{x} \in \mathbb{R}^K, \mathbf{z} \in \mathbb{R}^R$ s.t. $\forall k, r \exists$ unique $(x_k, z_r) \in \mathbf{x} \times \mathbf{z}$ s.t. $y_{k,r}(\mathbf{x}, \mathbf{z}) \sim f(x_k, z_r)$. If, in addition, \mathbf{x} and \mathbf{z} are Markov processes, we say that $\{y_{k,r}\}_{k=1,r=1}^{K,R}$ is a separably Markov random field (SMuRF).

A 2D random field $\{y_{k,r}\}_{k=1,r=1}^{K,R} \in \mathbb{R}^{K \times R}$ is a collection of random variables indexed over a subset of $\mathbb{N}^+ \times \mathbb{N}^+$. We call this collection a separable field if there exist latent random vectors \mathbf{x} and \mathbf{z} (each indexed over a subset of \mathbb{N}^+) such that $\{y_{k,r}\}_{k=1,r=1}^{K,R}$ are independent conditioned on \mathbf{x} and \mathbf{z} and only a function of the outer product between \mathbf{x} and \mathbf{z} . If, in addition, \mathbf{x} and \mathbf{z} are Markov, we say that the field is a SMuRF. Intuitively, a separable random field is a random field that admits a stochastic rank one decomposition.

We propose the following SMuRF model of the discrete-time, trial-dependent CIF $\{\lambda_{k,r}\}_{k=1,r=1}^{K,R}$ of a neuronal spiking unit

$$\begin{cases} x_k = \rho_x x_{k-1} + \alpha_x u_{x,k} + \epsilon_k, \epsilon_k \sim \mathcal{N}(0, \sigma_\epsilon^2), \\ z_r = \rho_z z_{r-1} + \alpha_z u_{z,r} + \delta_r, \delta_r \sim \mathcal{N}(0, \sigma_\delta^2), \\ \log \frac{\lambda_{k,r} \Delta}{1 - \lambda_{k,r} \Delta} = x_k + z_r, \\ \Delta N_{k,r} | x_k, z_r \sim \text{Bernoulli}(\lambda_{k,r} \Delta). \end{cases} \quad (2.2)$$

By construction, this is a SMuRF of the trial-dependent CIF of a neuron. $u_{x,k}$ and $u_{z,k}$ are indicator functions of the presence of a cue. To provide some intuition, if we assume $x_k + z_r$ is small, then the SMuRF model approximates the trial-dependent CIF as $\lambda_{k,r} \Delta \approx e^{z_r} \cdot e^{x_k}$, 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 above what can be expected from the within-trial component at that trial, which we call the cross-trial component of the CIF. The within- and cross-trial components from the SMuRF model 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 stochastic difference equation. The parameters $\rho_x, \alpha_x, \sigma_\epsilon^2, \rho_z, \alpha_z$ and σ_δ^2 , which




Figure 1: Representation of the SMuRF model as a Bayesian network. The SMuRF model approximates the trial-dependent CIF, the product of a within-trial component, in units of Hz (spikes/s) and a unitless quantity. 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, each according to a first-order stochastic difference equation. Observations from the raster are Bernoulli random variables whose probability of occurrence is a nonlinear function of the sum of the two sequences.

govern the smoothness of $(x_k)_{k=1}^K$ and $(z_r)_{r=1}^R$, must be estimated from the raster data.

Remark 1. We note that in its generality, our model does not assume that $\lambda_{k,r} \Delta = e^{z_r} \cdot e^{x_k}$. In our model, $\lambda_{k,r} \Delta = \frac{e^{x_k + z_r}}{1 + e^{x_k + z_r}}$. The approximation $\lambda_{k,r} \Delta \approx e^{z_r} \cdot e^{x_k}$ holds for a neuron with small neural spiking rate (Truccolo et al., 2005).

Figure 1 shows a graphical representation of the SMuRF model as a Bayesian network. It is not mathematically possible to rewrite the state equations from the SMuRF model in standard state-space form without increasing significantly the dimension of the state space. We give a sketch

of an argument as to why in the appendix. Therefore, in an EM algorithm for parameter estimation, one cannot simply apply classical (approximate) binary filtering and smoothing in the E-step (Smith & Brown, 2003). We derive a Monte Carlo EM algorithm to maximize the likelihood of observed data under the SMuRF model with respect to the parameter vector $\theta = (\rho_x, \alpha_x, \sigma_\epsilon^2, \rho_z, \alpha_z, \sigma_\delta^2)$.

3 Parameter Estimation in the SMuRF by Maximum Likelihood

3.1 Maximum Likelihood Estimation by Expectation-Maximization. Let $\mathbf{x} = (x_1, \dots, x_K)^T$, $\mathbf{z} = (z_1, \dots, z_R)^T$, and $\Delta\mathbf{N} = \{\Delta N_{k,r}\}_{k=1, r=1}^{K,R}$. The goal is to maximize, with respect to θ , the likelihood $L(\theta|\Delta\mathbf{N})$ of the SMuRF model:

$$L(\theta|\Delta\mathbf{N}) = \log p(\Delta\mathbf{N}; \theta) = \log \int_{\mathbf{x}, \mathbf{z}} p(\Delta\mathbf{N}, \mathbf{x}, \mathbf{z}; \theta) d\mathbf{x} d\mathbf{z}. \quad (3.1)$$

This is a challenging problem because of the high-dimensional integral that must be carried out in equation 3.1. We propose to maximize the likelihood by EM.

Remark 2. For the moment, we treat \mathbf{x} and \mathbf{z} as missing data. In the sequel, we will augment the model with additional missing data that will simplify the EM algorithm.

Given a candidate solution $\theta^{(\ell)}$, EM (Dempster, Laird, & Rubin, 1977) maximizes $L(\theta|\Delta\mathbf{N})$ by building a sequence of successive approximations $\mathcal{Q}(\theta|\theta^{(\ell)})$ of $L(\theta|\Delta\mathbf{N})$ (the so-called E-step) such that maximizing these approximations, which in general is simpler than directly maximizing $L(\theta|\Delta\mathbf{N})$, is guaranteed to not decrease $L(\theta|\Delta\mathbf{N})$. That is, each iteration of EM generates a new candidate solution $\theta^{(\ell+1)}$ such that $L(\theta^{(\ell+1)}|\Delta\mathbf{N}) \geq L(\theta^{(\ell)}|\Delta\mathbf{N})$. By iterating this process, EM generates a sequence of iterates $\{\theta^{(\ell)}\}_{\ell=1}^\infty$ that under regularity conditions converges to a local optimum of $L(\theta|\Delta\mathbf{N})$ (Dempster et al., 1977).

In the context of the SMuRF model, the key challenge of EM is to compute $\mathcal{Q}(\theta|\theta^{(\ell)})$, defined as

$$\mathcal{Q}(\theta|\theta^{(\ell)}) = \mathbb{E}_{\mathbf{x}, \mathbf{z}} \left[\log p(\Delta\mathbf{N}, \mathbf{x}, \mathbf{z}; \theta) | \Delta\mathbf{N}, \theta^{(\ell)} \right], \quad (3.2)$$

the expected value of the complete-data likelihood with respect to the joint posterior distribution of the missing data (\mathbf{x}, \mathbf{z}) conditioned on the observed data $\Delta\mathbf{N}$ and the candidate solution $\theta^{(\ell)}$. This expectation is not tractable; it cannot be computed in closed form. The intractability stems not only from the lack of conjugacy between the Bernoulli observation model and our gaussian priors—also an issue for one-dimensional models (Smith & Brown, 2003)—but also because, as mentioned previously, the SMuRF

model cannot be reduced to a standard state-space model. We propose to approximate the required expectations using Markov Chain monte-carlo (MCMC) samples from $p(\mathbf{x}, \mathbf{z} | \Delta \mathbf{N}; \theta^{(\ell)})$. In particular, we will use Gibbs sampling (Casella & George, 1992), a Monte Carlo technique, to generate samples from a distribution by sampling from its so-called full conditionals (conditional distribution of one variable given all others), thus generating a Markov chain that, under regularity conditions, can be shown to converge to a sample from the desired distribution. Gibbs sampling is attractive in cases where sampling from the full conditionals is simple. However, it is prone to the drawbacks of MCMC methods, such as poor mixing and slow convergence, particularly if one is not careful in selecting the full conditionals from which to generate samples from. Two observations are in order that will lead to the derivation of an elegant block Gibbs sampler with attractive properties:

- Conditioned on \mathbf{z} , the joint distribution, $p(\Delta \mathbf{N}, \mathbf{x} | \mathbf{z}; \theta)$, of \mathbf{x} and $\Delta \mathbf{N}$ is equivalent to the joint distribution from a one-dimensional state-space model with binary observations (Smith & Brown, 2003). By symmetry, this is also true for $p(\Delta \mathbf{N}, \mathbf{z} | \mathbf{x}; \theta)$. This readily motivates a block Gibbs sampler that alternates between sampling from $\mathbf{x} | \Delta \mathbf{N}; \theta$ and $\mathbf{z} | \Delta \mathbf{N}; \theta$. This leaves us with one challenge: how to obtain samples from the posterior distribution of the state in a one-dimensional state-space model with Bernoulli (more generally binomial) observations?
- We introduce a new collection of independent and identically distributed (i.i.d.), Pólya-Gamma distributed (Polson, Scott, & Windle, 2013) random variables $\mathbf{w} = \{w_{k,r}\}_{k=1, r=1}^{K,R}$, such that sampling from $\mathbf{x} | \Delta \mathbf{N}, \mathbf{w}; \theta$ is equivalent to sampling from the posterior of the state in a linear gaussian state-space model (we prove this in the appendix) using a forward-filtering, backward-sampling algorithm (Frühwirth-Schnatter, 1994). Moreover, it has been shown that the Gibbs sampler based on this Pólya-Gamma augmentation scheme (Choi & Hobert, 2013) is uniformly ergodic and possesses superior mixing properties to alternate a data-augmentation scheme for logit-based models (Polson et al., 2013; Choi & Hobert, 2013). The intuition behind the introduction of the Pólya-Gamma random variables is the following: they are missing data that, if we could observe, would make the Bernoulli observations gaussian. Stated otherwise, the Pólya-Gamma random variables are scale variables in a gaussian scale mixture (Andrews & Mallows, 1974) representation of Bernoulli random variables.

Remark 3. The random vector \mathbf{w} in the preceding bullet point is the vector additional missing data alluded to in remark 2.

Together, these two observations form the basis of an efficient block-Gibbs sampler we use for maximum-likelihood estimation of the

parameters from the SMuRF model by Monte Carlo EM, also referred to as empirical Bayes (Casella, 2001). We introduce the basic ideas behind PG augmentation and its utility in Bayesian estimation for logit-based models. In the appendix, we provide detailed derivations for the PG sampler adapted to the SMuRF model.

3.2 Pólya-Gamma Augmentation and Sampling in One Dimension.

Let $\Delta N \in \{0, 1\}$ and $X \in \mathbb{R}$ and suppose that, conditioned on $X = x$, ΔN is Bernoulli with mean $\frac{e^x}{1+e^x}$:

$$p(\Delta N|x) = \frac{(e^x)^{\Delta N}}{1 + e^x}. \quad (3.3)$$

We begin with a definition of Pólya-Gamma (PG) random variables, followed by a PG augmentation scheme for the Bernoulli/binomial likelihood. We will see that the augmentation scheme leads to an attractive form for the posterior of x given the observation ΔN and the augmented variable. Finally, we will see that the posterior of the augmented variable itself follows a PG distribution. Our treatment follows closely that of Choi and Hobert (2013).

3.2.1 Definition of Pólya-Gamma Random Variables. Let $\{E_m\}_{m=1}^\infty$ be a sequence of an i.i.d. exponential random variable with parameter equal to 1. The random variable

$$W \stackrel{d}{=} \frac{2}{\pi^2} \sum_{m=1}^{\infty} \frac{E_m}{(2m-1)^2} \quad (3.4)$$

follows a PG(1,0) distribution, where $\stackrel{d}{=}$ denotes equality in distribution. The moment-generating function of W is

$$\mathbb{E}[e^{-tW}] = \cosh^{-1} \left(\frac{\sqrt{t}}{2} \right). \quad (3.5)$$

An expression for its density $p_W(w)$, expressed as an infinite sum, can be found in Choi and Hobert (2013) and Polson et al. (2013). The PG(1, c) random variable is obtained by exponential tiling of the density of a PG(1,0) random variable. Letting $p_W(w|c)$ denote the density of a PG(1, c) random variable,

$$p_W(w|c) = \cosh \left(\frac{c}{2} \right) e^{-\frac{c^2 w}{2}} p_W(w). \quad (3.6)$$

3.2.2 PG Augmentation Preserves the Bernoulli Likelihood. Following the treatment of Choi and Hobert (2013), conditioned on $X = x$, let W be a $\text{PG}(1, |x|)$ random variable. Furthermore, suppose that, conditioned on $X = x$, ΔN and W are independent. Then

$$p(\Delta N, w|x) = p(\Delta N|x)p_W(w|x). \quad (3.7)$$

Integrating out W , we see that the augmentation scheme does not alter $p(\Delta N|x)$. One may then ask, What is the utility of the augmentation scheme? The answer lies in the following identity, discussed in detail in Choi and Hobert (2013), and is the key idea behind PG augmentation,

$$p(\Delta N|x)p_W(w|x) = \frac{(e^x)^{\Delta N}}{1 + e^x} \cdot \text{PG}(1, |x|) \propto e^{-\frac{1}{2} \frac{(\tilde{y}-x)^2}{1/w}} \propto \mathcal{N}\left(\tilde{y}; x, \frac{1}{w}\right), \quad (3.8)$$

where $\tilde{y} = \frac{y - \frac{1}{2}}{w}$, and \propto indicates that we are dropping terms independent of x . Equation 3.8 states that given $X = x$ and a logit model, a Bernoulli random variable is, up to a constant independent of x , a scale mixture of gaussians (Andrews & Mallows, 1974), that is, a gaussian random variable with random variance $1/w$, where $W = w$ follows a PG distribution (Polson et al., 2013; Choi & Hobert, 2013). If we assume $X \sim p_X(x)$, then (Choi & Hobert, 2013)

$$p(\Delta N, x, w) = p(\Delta N|x)p_W(w|x)p_X(x) \propto \mathcal{N}\left(\tilde{y}; x, \frac{1}{w}\right) p_X(x). \quad (3.9)$$

3.2.3 Implications of Augmentation on $p(x|\Delta N, w)$ and $p(w|\Delta N, x)$.

$$\begin{aligned} p(x|\Delta N, w) &= \frac{p(\Delta N, x, w)}{p(\Delta N, w)} \propto p(\Delta N|x)p_W(w|x)p_X(x) \\ &\propto \mathcal{N}\left(\tilde{y}; x, \frac{1}{w}\right) p_X(x), \end{aligned} \quad (3.10)$$

where we make use of equation 3.9. If X is gaussian, then $p(x|\Delta N, w)$ is gaussian and available in closed form (see the appendix):

$$p_W(w|\Delta N, x) = \frac{p(\Delta N, x, w)}{p(\Delta N, x)} = \frac{p(\Delta N|x)p_W(w|x)p_X(x)}{\int_w p(\Delta N|x)p_W(w|x)p_X(x)} = p_W(w|x), \quad (3.11)$$

that is, $p(w|\Delta N, x) = p_W(w|x) = \text{PG}(1, |x|)$. Together, equations 3.10 and 3.11 form the basis of a uniformly ergodic (Choi & Hobert, 2013) Gibbs sampler to obtain a sample from $p(x, w|\Delta N)$.

3.3 Block Gibbs Sampler for a PG-Augmented SMuRF Model. Consider the following version of the SMuRF model with PG augmentation:

$$\begin{cases} x_k = \rho_x x_{k-1} + \alpha_x u_{x,k} + \epsilon_k, \epsilon_k \sim N(0, \sigma_\epsilon^2), \\ z_r = \rho_z z_{r-1} + \alpha_z u_{z,r} + \delta_r, \delta_r \sim N(0, \sigma_\delta^2), \\ \lambda_{k,r} \Delta = \frac{e^{x_k + z_r}}{1 + e^{x_k + z_r}}, \\ \Delta N_{k,r} | x_k, z_r \sim \text{Bernoulli}(\lambda_{k,r} \Delta), \\ w_{k,r} | x_k, z_r \sim \text{PG}(1, |x_k + z_r|), k = 1, \dots, K; r = 1, \dots, R. \end{cases} \quad (3.12)$$

We can apply the basic results from section 3.2 to derive the following result (the proof is in the appendix):

Theorem 1. Suppose \mathbf{w} , \mathbf{x} , \mathbf{z} , and $\Delta \mathbf{N}$ come from the PG-augmented SMuRF model (equation 3.12), then $p(\Delta \mathbf{N}, \mathbf{x} | \mathbf{w}, \mathbf{z}; \theta)$ is equivalent in distribution to the following linear-gaussian state-space model

$$\begin{cases} x_k = \rho_x x_{k-1} + \alpha_x u_{x,k} + \epsilon_k, \epsilon_k \sim \mathcal{N}(0, \sigma_\epsilon^2) \\ \Delta \tilde{N}_k = x_k + \tilde{v}_k, \tilde{v}_k \sim \mathcal{N}\left(0, 1 / \sum_{r=1}^R w_{k,r}\right), \text{ i.i.d., } r = 1, \dots, R \\ \Delta \tilde{N}_k = x_k - \frac{K}{2} - \sum_{r=1}^R z_r w_{k,r}. \end{cases} \quad (3.13)$$

Following the discussion from section 3.2, it is not hard to see that such a result would hold. The proof of this result is in the appendix, as well as the derivation of an elegant forward-filtering, backward-sampling algorithm (Frühwirth-Schnatter, 1994) for drawing samples from $p(\Delta \mathbf{N}, \mathbf{x} | \mathbf{w}, \mathbf{z}; \theta)$. By symmetry, it is not hard to see that a similar result holds for $p(\Delta \mathbf{N}, \mathbf{z} | \mathbf{w}, \mathbf{x}; \theta)$.

3.3.1 Block Gibbs Sampling from a PG-Augmented SMuRF Model. The E-step of the Monte Carlo EM algorithm consists in sampling from $p(\mathbf{x}, \mathbf{z} | \Delta \mathbf{N}; \theta^{(\ell)})$ by drawing from $p(\mathbf{x}, \mathbf{z}, \mathbf{w} | \Delta \mathbf{N}; \theta^{(\ell)})$ using a block Gibbs sampler that uses the following full conditionals:

- $p(\mathbf{x} | \Delta \mathbf{N}, \mathbf{w}, \mathbf{z}; \theta^{(\ell)})$, which, according to theorem 1, is equivalent to the posterior distribution of the state sequence in a linear-gaussian state-space model
- $p(\mathbf{z} | \Delta \mathbf{N}, \mathbf{w}, \mathbf{x}; \theta^{(\ell)})$, which obeys properties similar to the previous full-conditional (by symmetry)

- $p(w_{k,r}|\mathbf{x}, \mathbf{z}) = p(w_{k,r}|x_k, z_r) = \text{PG}(1, |x_k + z_r|), k = 1, \dots, K, r = 1, \dots, R$ (Polson et al., 2013).

In the appendix, we detail how we initialize the algorithm and monitor convergence.

In practice, we found that estimating ρ_x and ρ_z is difficult. We hypothesize that including those parameters yields an unwieldy likelihood function. In the results we report, we assume $\rho_x = \rho_z = 1, \alpha_x = \alpha_z = 0$ and focus on estimating a simple model with two parameters σ_ϵ^2 and σ_δ^2 . The assumption $\rho_x = \rho_z = 1$ gives the random walk priors more freedom, thus allowing us to capture the variability of the within- and cross-trial processes. We have run simulations, not reported here, that show that the joint estimation of $\sigma_\epsilon^2, \sigma_\delta^2, \alpha_x$, and α_z is stable and that our EM algorithm converges. This demonstrates the ability of the SMuRF model (see equation 3.12) to incorporate exogenous input stimuli.

Assuming $\rho_x = \rho_z = 1$, in the M-step, the update equations for the parameters $\sigma_\epsilon^2, \sigma_\delta^2, \alpha_x$, and α_z follow standard formulas (Smith & Brown, 2003),

$$\alpha_x^{(\ell+1)} = \frac{\sum_{k=1}^K (\mathbb{E}_{\mathbf{x}} [x_k - x_{k-1} | \Delta \mathbf{N}, \theta^{(\ell)}]) u_{x,k}}{\sum_{k=1}^K u_{x,k}^2}, \quad (3.14)$$

$$\alpha_z^{(\ell+1)} = \frac{\sum_{r=1}^R (\mathbb{E}_{\mathbf{z}} [z_r - z_{r-1} | \Delta \mathbf{N}, \theta^{(\ell)}]) u_{z,r}}{\sum_{r=1}^R u_{z,r}^2}, \quad (3.15)$$

$$\sigma_\epsilon^{2(\ell+1)} = \mathbb{E}_{\mathbf{x}} \left[\frac{1}{K} \sum_{k=1}^K (x_k - x_{k-1} - \alpha_x^{(\ell+1)} u_{x,k})^2 | \Delta \mathbf{N}, \theta^{(\ell)} \right], \quad (3.16)$$

$$\sigma_\delta^{2(\ell+1)} = \mathbb{E}_{\mathbf{z}} \left[\frac{1}{R} \sum_{r=1}^R (z_r - z_{r-1} - \alpha_z^{(\ell+1)} u_{z,k})^2 | \Delta \mathbf{N}, \theta^{(\ell)} \right], \quad (3.17)$$

where we set $x_0 = z_0 = 0$ and approximate the expectations with respect to $p(\mathbf{x} | \Delta \mathbf{N}, \theta^{(\ell)})$ and $p(\mathbf{z} | \Delta \mathbf{N}, \theta^{(\ell)})$ using Gibbs samples from the E-step.

3.4 Assessment of Within-Trial and Cross-Trial Spiking Dynamics.

Bayesian estimation of the SMuRF model (see equation 3.12) enables us to infer detailed changes in neural dynamics, in particular to extract the within-trial and cross-trial components of the neural spiking dynamics that accompany the learning of a contingency by a neuron. This is because, following estimation, inference in the SMuRF model yields the joint posterior distribution of the instantaneous spiking rate of a neuron as a function of trials, and time within a trial, conditioned on the observed data. We can use this posterior distribution, in turn, to assess instantaneous changes in neural spiking dynamics and without the need to correct for multiple comparisons as with nonparametric methods.

In what follows, we let $p(\mathbf{x}, \mathbf{z} | \Delta \mathbf{N}; \hat{\theta}_{ML})$ denote the posterior distribution of \mathbf{x} and \mathbf{z} , given the raster data $\Delta \mathbf{N}$ and the maximum likelihood estimate $\hat{\theta}_{ML}$ of θ . In what follows, it is understood that we use Gibbs samples $(\mathbf{x}_i, \mathbf{z}_i)_{i=1}^n$ from $p(\mathbf{x}, \mathbf{z} | \Delta \mathbf{N}; \hat{\theta}_{ML})$ to obtain an empirical estimate of the distribution.

3.4.1 Posterior Distribution of the Joint CIF over Time and Trials. We can use these posterior samples to approximate the posterior distribution, at $\hat{\theta}_{ML}$, of any quantities of interest. Indeed, it is well known from basic probability that if $(\mathbf{x}_i, \mathbf{z}_i)$ is a sample from $p(\mathbf{x}, \mathbf{z} | \Delta \mathbf{N}; \hat{\theta}_{ML})$, then $f(\mathbf{x}_i, \mathbf{z}_i)$ is a sample from $p(f(\mathbf{x}, \mathbf{z}) | \Delta \mathbf{N}; \hat{\theta}_{ML})$. In particular, if the instantaneous spiking rate of a neuron at time k and trial r $\lambda_{k,r} \Delta = \frac{e^{x_k + z_r}}{1 + e^{x_k + z_r}}$, we can use the Gibbs samples to approximate the joint posterior distribution of $\{\lambda_{k,r} \Delta\}_{k=1, r=1}^{K,R}$ given $\Delta \mathbf{N}$ and $\hat{\theta}_{ML}$.

Let $\{\lambda_{k,r}^p \Delta\}_{k=1, r=1}^{K,R}$ be the random variable that represents the a posteriori instantaneous spiking rate of the neuron at time trial r and time k within that trial. The superscript p highlights the conditioning on the data $\Delta \mathbf{N}$ and $\hat{\theta}_{ML}$ and the fact that this quantity is a function of (\mathbf{x}, \mathbf{z}) distributed according to $p(\mathbf{x}, \mathbf{z} | \Delta \mathbf{N}; \hat{\theta}_{ML})$.

3.4.2 Within-Trial Effect. We define the within-trial effect as the a posteriori instantaneous spiking rate at time k , averaged over all trials:

$$e_k^{WT} = \frac{1}{R} \sum_{r=1}^R \lambda_{k,r}^p(x_k, z_r), k = 1, \dots, K. \quad (3.18)$$

It is important to note that the averaging is performed after characterization of the joint CIF as a function of time and trials, which is not the same as first aggregating the data across trials and applying one of the one-dimensional methods for analyzing neural data (Smith & Brown, 2003; Zammit-Mangion et al., 2012; Yuan et al., 2012). In practice, every Gibbs sample pair $(\mathbf{x}_i, \mathbf{z}_i)$, $i = 1, \dots, n$ leads to a scalar quantity:

$$\hat{e}_{i,k}^{WT} = \frac{1}{R} \sum_{r=1}^R \lambda_{k,r}(\mathbf{x}_{i,k}, \mathbf{z}_{i,r}), k = 1, \dots, K. \quad (3.19)$$

Performing this computation over all Gibbs samples and times $k = 1, \dots, K$ leads to a joint empirical distribution for the within-trial effect $\{e_k^{WT}\}_{k=1}^K$.

3.4.3 Cross-Trial Effect. We define the cross-trial effect as the a posteriori excess instantaneous spiking at trial r and time k (above the within-trial effect e_k^{WT}) averaged across all times k :

$$e_r^{\text{CT}} = \frac{1}{K} \sum_{k=1}^K \frac{\lambda_{k,r}^p(x_k, z_r)}{e_k^{\text{WT}}}, r = 1, \dots, R. \quad (3.20)$$

In practice, every Gibbs sample pair $(\mathbf{x}_i, \mathbf{z}_i)$, $i = 1, \dots, n$ leads to a scalar quantity:

$$\hat{e}_{i,r}^{\text{CT}} = \frac{1}{K} \sum_{k=1}^K \frac{\lambda_{k,r}(\mathbf{x}_{i,k}, \mathbf{z}_{i,r})}{\hat{e}_{i,k}^{\text{WT}}}, r = 1, \dots, R. \quad (3.21)$$

Performing this computation over all Gibbs samples and R_c trials of interest, $r = R - R_c + 1, \dots, R$, leads to a joint empirical distribution for the cross-trial effect $\{e_r^{\text{CT}}\}_{r=R-R_c}^R$.

Remark 4. The next paragraph explains the meaning of R_c in the context of an associative learning experiment.

3.5 Assessment of Neural Spiking Dynamics across Time and Trials.

Consider an associative learning (conditioning) experiment characterized by the pairing of a conditioned stimulus (e.g., auditory) to an aversive stimulus (e.g., a shock). Let R_c be the number of conditioning trials and K_h the length of the habituation period. Gibbs samples from the SMuRF model (see equation 3.12) parameterized by θ_{ML} let us approximate the a posteriori probability that the spiking rate at a given point (see point C in Figure 2) during one of the conditioning trials (trials 16 through 45 in this example) is bigger than the baseline spiking rate at that trial (see region A in Figure 2) and the average spiking rate at the same time during the habituation period (see region B in Figure 2). This yields a probabilistic description of the intricate dynamics of neural spiking that accompany the learning of the contingency by a neuron. Let

$$\text{Event U} = \left\{ \lambda_{k,r}^p(x_k, z_r) > \overbrace{\frac{1}{R_c} \sum_{m=1}^{R_c} \lambda_{k,m}^p(x_k, z_m)}^{\text{Average rate in Region B}} \right\}, \quad (3.22)$$

$$\text{Event V} = \left\{ \lambda_{k,r}^p(x_k, z_r) > \overbrace{\frac{1}{K_h} \sum_{s=1}^{K_h} \lambda_{s,r}^p(x_s, z_r)}^{\text{Average rate in Region A}} \right\}. \quad (3.23)$$




Figure 2: Regions defined to quantify changes in neural spiking dynamics in an associative learning experiment. The SMuRF model lets us approximate the a posteriori probability that the spiking rate at a given point C during one of the conditioning trials (trials 16–45 in this example) is bigger than the baseline spiking rate at that trial (region A) and the average spiking rate at the same time during the habituation period (region B). This yields a probabilistic description of the intricate dynamics of neural spiking that accompany the learning of the contingency in the experiment by a neuron.

For a given pair (k, r) s.t. $k \geq K_h, r \geq R_c$, this probability is

$$\mathbb{P}[\text{Event U} \cap \text{Event V}] \tag{3.24}$$

$$\approx \frac{1}{n} \sum_{i=1}^n \mathbb{I}_{\{\lambda_{k,r}^p(\mathbf{x}_{i,k}, \mathbf{z}_{i,r}) > \frac{1}{R_c} \sum_{m=1}^{R_c} \lambda_{k,m}^p(\mathbf{x}_{i,k}, \mathbf{z}_{i,m}) \cap \lambda_{k,r}^p(\mathbf{x}_{i,k}, \mathbf{z}_{i,r}) > \frac{1}{K_h} \sum_{s=1}^{K_h} \lambda_{s,r}^p(\mathbf{x}_{i,s}, \mathbf{z}_{i,r})\}}, \tag{3.25}$$

where the second line approximates the probability of the event of interest using its frequency of occurrence in the n posterior samples. As we demonstrate in section 4, we thus obtain a detailed characterization of the dynamics of neural spiking that accompany learning.

In section 4, we use simulated and real data examples to demonstrate the utility of the SMuRF model (see equation 3.12) for the characterization of detailed neural spiking dynamics.




Figure 3: Setup used to simulate neural spike raster data from a neuron that exhibits a conditioned response to the conditioned stimulus in an associative learning experiment. The experiment consists of 45 trials, each of which lasts 2 s. The conditioned stimulus becomes active 1 s into a trial, and the aversive stimulus becomes active from the 16th trial onward. An idealized neuron that exhibits a conditioned response would exhibit two distinct regions of activity. Region B consists of all trials preceding trial 16, along with the period from all trials before the conditioned stimulus is presented. We assume that the rate of spiking of the neuron is $\lambda_B = 20$ Hz. Region A consists of the period from trials following trial 15 after the conditioned stimulus is presented. The rate of spiking of the neuron in this region is $\lambda_A = 60$ Hz.

4 Applications

4.1 Simulation Studies. We simulated neural spike raster data from a neuron that exhibits a conditioned response to the conditioned stimulus (see Figure 4) in an associative learning experiment. The experiment consists of 45 trials, each of which lasts 2 s. The conditioned stimulus becomes active 1 s into a trial, while the aversive stimulus becomes active after trial 15. We obtain the simulated data by dividing the raster into two predefined regions as shown in Figure 3. Region B consists of all trials before trial 16, along with the period from all trials before the conditioned stimulus is presented. We assume that the rate of the spiking of the neuron is $\lambda_B = 20$ Hz. Region A consists of the period from trials following trial 15 after the conditioned stimulus is presented. The rate of spiking of the neuron in this region is $\lambda_A = 60$ Hz.




Figure 4: (a) Simulated neural spike raster, along with estimated within- and cross-trial effects from the SMuRF model. The horizontal red line indicates the beginning of conditioning. The vertical green line indicates the onset of the conditioned stimulus. The left panel of the figure shows the cross-trial effect, which, following conditioning, increases above its average initial value of ≈ 1 to ≈ 2 . (b) Empirical probability that the spiking rate at a given trial and time is bigger than the average rate at the same trial and the average rate during habituation at the same time. With probability close to 1, the spiking rate at a given time/trial pair—following the conditioned stimulus and during conditioning (Figure 2C)—is larger than the average rate at the same trial (Figure 2A) and the average rate at the same time (Figure 2B)

We applied the SMuRF model (see equation 3.12) to the analysis of this simulated neural spike raster. Figure 4a shows that during conditioning, learning is accompanied by a doubling of the spiking rate above the within-trial spiking rate of the neuron. Indeed, the left panel of the figure shows the cross-trial effect, which, following conditioning, increases above its average initial value of ≈ 1 to ≈ 2 . Figure 4b provides a more detailed characterization of the neural spiking dynamics. With probability close to 1, the spiking rate at a given time/trial pair, following the conditioned stimulus and during conditioning (see Figure 2C), is bigger than the average rate at the same trial (see Figure 2A) and the average rate at the same time (see Figure 2B). We conclude that with high probability, the simulated neuron exhibits a conditioned response to the conditioned stimulus.

4.2 Neural Dynamics during Associative Learning of Fear

4.2.1 Basic Experimental Paradigm. The ability to learn through observation is a powerful means of learning about aversive stimuli without direct experience of the stimuli.

We use the SMuRF model (see equation 3.12) to analyze data from a fear-conditioning paradigm designed to elucidate the nature of the circuits that facilitate the associative learning of fear. The experimental paradigm is described in detail in Allsop et al. (2017). Briefly, an observer mouse observes a demonstrator receive conditioned stimulus-shock pairings through a perforated transparent divider. The experiment consists of 45 to 50 trials, divided into two phases. During the first 15 trials of the experiment, termed the habituation period, both the observer and the demonstrator simply hear an auditory conditioned stimulus. From the 16th trial onward, the auditory conditioned stimulus is followed by the delivery of a shock to the demonstrator. The data are recorded from the prefrontal cortex (PFC) of the observer mouse.

4.2.2 Results. Figure 5a shows the within- and cross-trial effects estimated using the SMuRF model applied to a cortical neuron from the experiment described. The estimated within-trial (bottom) and cross-trial (left) components indicate significant changes respectively in response to the conditioned stimulus and to conditioning. By definition (see equation 3.20), the cross-trial effects takes into account the increase in spiking rate due to the presentation of the conditioned stimulus. The bottom panel suggests that this neuron exhibits a delayed response to the conditioned stimulus, beginning at ≈ 400 ms following conditioned stimulus presentation. Accounting for this increase in the within-trial spiking rate due to the conditioned stimulus, the left panel shows a multiplicative increase in spiking rate due to conditioning from an average initial value of less than ≈ 1 (indicative of suppression, as can be seen through the sparseness of the raster during trials 1 through 5) to a peak average value of ≈ 4 at trial 23. This increase, however, does not persist as in the case of the simulated data (see Figure 4), suggesting that conditioning is accompanied by intricate dynamics in neural modulation.

Figure 5b provides a more detailed characterization of the neural spiking dynamics of this neuron. The figure shows the evolution, as a function of time and trials, of the probability that the spiking rate at a given time/trial pair (see Figure 2C) is bigger than the average rate at the same trial (see Figure 2A) and the average rate at the same time (see Figure 2B). The figure indicates that this neuron exhibits a delayed conditioning to the conditioned stimulus (beginning ≈ 400 ms following the conditioned stimulus presentation) and that the extent of the condition is highest first between trials 18 and 24 and then between trials 31 and 41.

Figure 6 shows an application of the SMuRF model (see equation 3.12) to a cortical neuron that does not exhibit a conditioned response to the

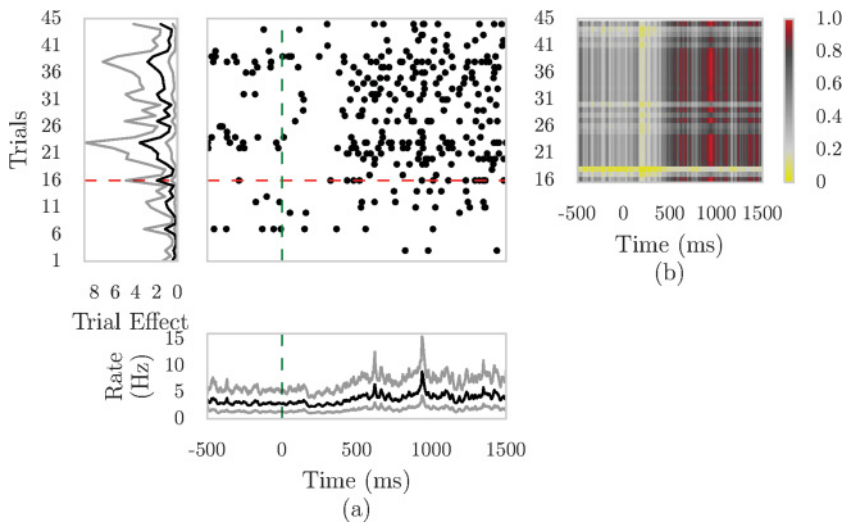


Figure 5: (a) Raster, along with SMuRF within and cross-trial components, from a cortical neuron that exhibits a conditioned response to the conditioned stimulus. The horizontal red line indicates the beginning of conditioning. The vertical green line indicates conditioned stimulus onset. The bottom panel suggests that this neuron exhibits a delayed response to the conditioned stimulus, beginning at ≈ 400 ms following conditioned stimulus presentation. Accounting for this increase in the within-trial spiking rate due to the conditioned stimulus, the left panel shows a multiplicative increase in spiking rate due to conditioning from an average initial value of less than ≈ 1 (indicative of suppression, as can be seen through the sparseness of the raster during trials 1 through 5) to a peak average value of ≈ 4 at trial 23. This increase, however, does not persist as in the case of the simulated data (see Figure 4). (b) Empirical probability that the spiking rate at a given trial and time is bigger than the average rate at the same trial and the average rate during habituation at the same time (refer to Figure 2). This panel indicates that this neuron exhibits a delayed conditioning to the conditioned stimulus (beginning ≈ 400 ms following conditioned stimulus presentation) and that the extent of the condition is highest first between trials 18 and 24 and then between trials 31 and 41. Panels a, and, in particular, b suggest that conditioning is accompanied by intricate dynamics in neural modulation.

conditioned stimulus. The bottom of Figure 6a indicates no significant increase in the within-trial spiking rate in response to the conditioned stimulus, while the left panel shows that the cross-trial effect remains constant throughout the experiment at an average value of ≈ 1 . This indicates that conditioning does not result in a significant increase in spiking rate. Figure 6b corroborates these findings: for all points C following the conditioned stimulus and during conditioning, there is a small probability that the

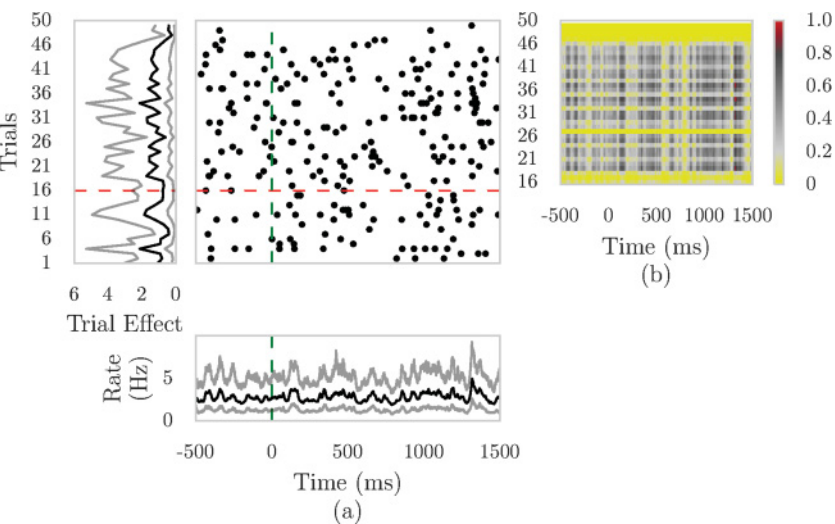


Figure 6: (a) Raster, along with SMuRF within- and cross-trial components, from a cortical neuron that does not exhibit a conditioned response to the conditioned stimulus. The horizontal red line indicates the beginning of conditioning. The vertical green line indicates the onset of the conditioned stimulus. This neuron does not exhibit a conditioned response to the conditioned stimulus. The bottom of the panel indicates no significant increase in the within-trial spiking rate in response to the conditioned stimulus, while the left panel shows that the cross-trial effect remains constant throughout the experiment at an average value of ≈ 1 . This indicates that conditioning does not result in a significant increase in spiking rate. (b) Empirical probability that the spiking rate at a given trial and time is bigger than the average rate at the same trial and the average rate during habituation at the same time This panel corroborates the observations from panel a: for all points C following the conditioned stimulus and during conditioning, there is a small probability that the instantaneous spiking rate is significantly different from the average spiking rates in regions A and B.

instantaneous spiking rate is significantly different from the average spiking rates in regions A and B. Figures 10 and 11 in the appendix show results for two additional cortical neurons that exhibit a transient conditioned response to the conditioned stimulus.

4.2.3 Using SMuRF Inference to Determine a Neuron’s Learning Time and Trial. The power of the Bayesian approach, and the SMuRF model (see equation 3.12) in particular, lies in the fact that it lets us approximate the a posteriori probability that the spiking rate at a given point (point C in Figure 2) during one of the conditioning trials (trials 16 through 45 in this example) is bigger than the baseline spiking rate at that trial (region A in

Table 1: Learning Trial and Time Computed for the Cortical Neurons Analyzed.

	Neuron in Figure 5	Neuron in Figure 6	Neuron in Figure 10	Neuron in Figure 11
Learning time (ms)	617	1316	202	20
Learning trial	16	34	18	18

Notes: The learning times and trials reported are consistent with the detailed inference provided by the respective figures for these neurons. The cortical unit from Figure 5, for instance, shows a delayed response, significant 617 ms after conditioned stimulus onset and at trial 16. The cortical unit from Figure 6 exhibits a significant change only in neural spiking 1316 ms following the conditioned stimulus and at trial 34. This is consistent with our previous observation from Figure 6 that this neuron does not exhibit a conditioned response to the stimulus.

Figure 2) and the average spiking rate at the same time during the habituation period (region B in Figure 2) (see equation 3.25). This yields an instantaneous probabilistic quantification of the extent of learning for any given time and trial pair.

Figures 5b, 6b, 10b, and 11b provide detailed characterizations of the dynamics of learning and its extent for all times following the onset of the conditioned stimulus and all conditioning trials.

Here, we provide some guidance for practitioners to summarize the results of our inference to a single learning time/trial pair. We stress, however, that the power of our methods lies in the detail provided by Figures 5b, 6b, 10b and 11b. Since the SMuRF model enables us to compute an empirical probability that the spiking rate at a given time/trial pair (see Figure 2, point C) is bigger than the average rate at the same trial (see Figure 2, region A) and the average rate at the same time (see Figure 2, region B), we can identify learning time and learning trial for each neuron by finding the first time after cue and after conditioning that this probability exceeds a certain threshold. Table 1 reports the learning time and trial computed using a threshold of 95%. Note that the learning time is computed with respect to the onset of the conditioned stimulus (time = 0 ms).

The learning times and trials reported in Table 1 are consistent with the detailed inference provided by the respective figures for these neurons. Indeed, the cortical unit from Figure 5 shows a delayed response, significant at 617 ms after conditioned stimulus onset and at trial 16. The cortical unit from Figure 6 exhibits a significant change in neural spiking only 1316 ms following the conditioned stimulus and at trial 34. This is consistent with our previous observation from Figure 6 that this neuron does not exhibit a conditioned response to the stimulus.

In the appendix, we perform a simulation that demonstrates the ability of SMuRF inference to identify learning time and trial when learning of a contingency is accompanied by sustained changes in neural spiking following

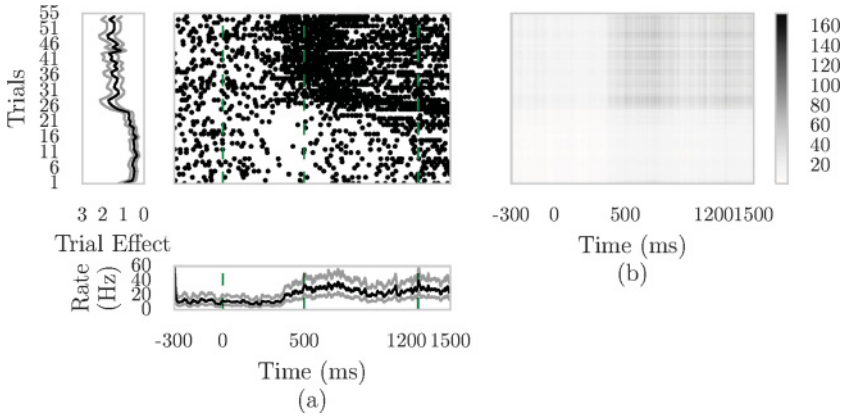


Figure 7: Application of the SMuRF model to neural spiking activity from a hippocampal neuron recorded during an experiment designed for a location-scene association learning task. The green lines represent the end of the fixation, scene presentation, and delay periods, respectively. (a) Neural spike raster from the hippocampal neuron. The bottom panel shows the estimate of the within-trial effect from the SMuRF model, while the left panel shows the cross-trial effect. (b) Estimate of the a posteriori instantaneous spiking rate $\{\hat{\lambda}_{k,r}^p, \Delta\}_{k=1, r=1}^{K,R}$ (in Hz) of the hippocampal neuron at time trial r and time k within that trial. This figure shows that while the SMuRF model is able to characterize the detailed changes in spiking dynamics of this neuron during the task, it fails to capture the nonseparable nature of the raster data.

conditioned stimulus onset and during conditioning. We also demonstrate through simulation that the SMuRF model is robust to the presence of error trials.

4.3 Application of SMuRF Model to a Nonseparable Example. We demonstrate the limitations of the separability assumption in the SMuRF model (see equation 3.12) by applying it to the neural spike raster data from Czanner et al. (2008). We briefly describe the experiment here and refer readers to Wirth et al. (2003) for a more detailed description. Figure 7a shows neural spiking activity from a hippocampal neuron recorded during an experiment designed for a location-scene association learning task. The same scene was shown to a macaque monkey across 55 trials, each lasting 1700 ms. The first 300 ms of every trial is a fixation period, and the scene is presented to the monkey from 300 to 800 ms. A delay period takes place from 800 to 1500 ms, followed by a response period from 1500 to 1700 ms.

The data from the experiment are shown in the center of Figure 7a. The raster suggests that the time- and trial-dependent CIF of this neuron is not separable. Intuitively, this can be seen from the fact that the region in which

there are significant changes in neural spiking does not follow the rectangular form from Figure 3. Nevertheless, the CIF could be well approximated by a separable model. We apply the SMuRF model to these data to uncover some of its limitations in nonseparable settings. The bottom panel of Figure 7a shows the estimate of the within-trial effect from the SMuRF model, while the left panel shows the cross-trial effect. These two panels indicate that the SMuRF model is able to capture within- and cross-trial dynamic changes in the spiking activity of the neuron. Figure 7b shows the estimate of the a posteriori mean instantaneous spiking rate $\{\hat{\lambda}_{k,r}^P \Delta\}_{k=1,r=1}^{K,R}$ (in Hz) of the neuron at time trial r and time k within that trial. This figure shows that while the SMuRF model is able to characterize the detailed changes in spiking dynamics, it does not fully capture the nonseparable nature of the raster data.

Remark 5. Unlike for the cortical neurons, this experiment does not have a conditioning period. Hence, it does not make sense to generate plots such as Figure 5b.

5 Conclusion

We proposed a 2D separably Markov random field (SMuRF) for the analysis of neural spike rasters that obviates the need to aggregate data across time or trials, as in classical one-dimensional methods (Smith & Brown, 2003; Zammit-Mangion et al., 2012; Yuan et al., 2012), while retaining their interpretability. The SMuRF model approximates the trial-dependent conditional intensity function (CIF) of a neuron as the product of a within-trial component, in units of Hz (spikes/s), and a unitless quantity, which we call the cross-trial effect, that represents the excess spiking rate above what can be expected from the within-trial component at that trial. One key advantage of our 2D model-based approach over nonparametric methods stems 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 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 nonparametric methods break down because the sample size is 1.

We applied the SMuRF model to data collected from neurons in the pre-frontal cortex (PFC) in an experiment designed to characterize the neural underpinnings of the associative learning of fear in mice. We found that as a group, the recorded cortical neurons exhibit a conditioned response to the auditory conditioned stimulus, occurring three to four trials into conditioning. We also found intricate and varied dynamics of the extent to which the cortical neurons exhibit a conditioned response (e.g., delays, short-term

conditioning). This is likely reflective of the variability in synaptic strength, connectivity, and location of the neurons in the population.

In future work, we plan to investigate nonseparable random field models of neural spike rasters, such as Markov random fields (MRFs; Besag, 1974). Compared to the SMuRF model, MRFs are 2D models for which the dimensionality of the putative state-space is as large as the dimensionality of the raster, suggesting that MRFs may provide a more detailed characterization of neural spike rasters. Indeed, the SMuRF model makes the strong assumption that the neural spiking dynamics are decomposable into two timescales, with the additional simplifying assumption that there is only one component per timescale. This simplifying assumption is motivated by one-dimensional state-space models of neural data (Smith & Brown, 2003) in which a neuron's time-dependent CIF is only a function of one hidden state sequence. We will investigate the inclusion of additional components in future work. We also plan to investigate analogs of the SMuRF model for population-level data. MRFs, multicomponent and population-level SMuRF models, naturally lead to model selection problems, and to the investigation of tools, based on sequential Monte Carlo methods (Chopin, Jacob, & Papaspiliopoulos, 2013) (also known as particle filters), to compare state-space models of neural spike rasters, such as one-dimensional models (Smith & Brown, 2003), the SMuRF model, and MRFs. The development of such tools for model comparisons is, in our opinion, the ultimate measure of the ability of different models to capture the intricate dynamics present in neural spike rasters. Finally, the SMuRF model can be interpreted as a two-dimensional gaussian process prior on the neural spiking rate surface (Rad & Paninski, 2010), with a separable kernel that is the Kronecker product of kernels from Gauss-Markov processes (one process for each dimension). The choice of kernels in the SMuRF model leads to the very efficient algorithms for estimation and inference derived in this article. Moreover, these algorithms scale well to more than two dimensions, unlike classical kernel methods. We plan to explore this connection to gaussian process inference in future work.

Appendix: Supplementary Derivations and Results for SMuRF

A.1 The SMuRF Model Cannot Be Converted Easily to a Standard State-Space Model. We focus on the simple case when $\rho_x = \rho_z = 1$, and $\alpha_x = \alpha_z = 0$:

$$\begin{cases} x_k = x_{k-1} + \epsilon_k, \epsilon_k \sim \mathcal{N}(0, \sigma_\epsilon^2), k = 1, \dots, K \\ z_r = z_{r-1} + \delta_r, \delta_r \sim \mathcal{N}(0, \sigma_\delta^2), r = 1, \dots, R \end{cases}$$

Let $t = (r - 1) \times K + k, r = 1, \dots, R, k = 1, \dots, K$. The index t is obtained by unstacking the raster trials and serializing them.

The question we ask is whether the state equations from the SMuRF model can be turned into ones of the form

$$\mathbf{s}_t = \mathbf{A}\mathbf{s}_{t-1} + \mathbf{v}_t, \quad (\text{A.1})$$

where $\mathbf{s}_t \in \mathbb{R}^2$. Let $\mathbf{s}_{t,1}$ and $\mathbf{s}_{t,2}$ denote the first and second components of \mathbf{s}_t , respectively. We ask that $\mathbf{s}_t \in \mathbb{R}^2$ because of the two dimensions present in the SMuRF model. Allowing the dimensionality of \mathbf{s}_t to increase up to K would allow a representation of the form of equation A.1. However, this would become a very high-dimensional, unwieldy state-space model.

Intuitively, this cannot be done for the following reason: the dimensionality of the latent states in the SMuRF model is $K + R$, while the dimensionality of the state sequence in equation A.1 is $2 \times (K \times R)$. For there to be an equivalence, the sequence \mathbf{s}_t must necessarily be redundant (i.e., some of the states must be copies of previous states). Storing these copies would necessarily mean having to increase the dimensionality of the state-space.

Let $\mathbf{s}_t = \begin{bmatrix} x_{t - (\lceil \frac{t}{K} \rceil - 1) \times K} \\ z_{\lceil \frac{t}{K} \rceil} \end{bmatrix} \in \mathbb{R}^2$. The quantity $\lceil \frac{t}{K} \rceil$ gives the trial index r corresponding time index t . The within-trial index corresponding to index t is then obtained by subtracting $(r - 1) \times K$ from t .

Note, for instance, that $\mathbf{s}_1 = \begin{bmatrix} x_1 \\ r_1 \end{bmatrix}$ and $\mathbf{s}_{K+1} = \begin{bmatrix} x_1 \\ r_2 \end{bmatrix}$. In general, $\mathbf{s}_{t,1} = \mathbf{s}_{t',1} = x_{k_0}$ for some $1 \leq k_0 \leq K$ if and only if $t > t'$ s.t. $t - t' = p \times K$ for some integer p , where we assume without loss of generality that $t > t'$. That is, two different indices t and t' share the same within-trial component if and only if they are apart by an integer multiple of K . Stated otherwise, the first component of \mathbf{s}_t exhibits circular symmetry. Therefore, for equation A.1 to hold, $\mathbf{s}_{t-1,1}$ must equal $\mathbf{s}_{t,1}$, which is not possible because t and $t - 1$ are not apart by an integer multiple of p !

The argument shows that in order to write the SMuRF state equations in the form of equation A.1, one would need to augment the state \mathbf{s}_t to dimension $K + 1$, which would lead to a very high-dimensional standard state-space model, thus increasing the complexity of performing inference.

A.2 Derivation of Gibbs Sampler for a PG-Augmented SMuRF Model. We first derive theorem 1, which leads to the forward-filter, backward-sampling algorithm from the full conditionals for \mathbf{x} and \mathbf{z} in the Gibbs sampler.

Since $w_{k,r}|x_k, z_r$ is drawn from a PG distribution, we can write the log pdf of $w_{k,r}|x_k, z_r$ as

$$\begin{aligned} \log p(w_{k,r}|x_k, z_r) &= \log \left(\cosh \left(\frac{x_k + z_r}{2} \right) \right) \\ &+ \log \left(\sum_{i=1}^{\infty} (-1)^i \frac{(2i+1)}{\sqrt{2\pi} w_{k,r}^2} e^{-\frac{(2i+1)^2}{8w_{k,r}} - \frac{x_k^2 w_{k,r}}{2}} \right). \end{aligned} \quad (\text{A.2})$$

The complete data likelihood of the SMuRF model is

$$p(\Delta \mathbf{N}, \mathbf{x}, \mathbf{w}; \theta) = p(\Delta \mathbf{N} | \mathbf{x}) p(\mathbf{w} | \mathbf{x}) p(\mathbf{x}) \quad (\text{A.3})$$

$$= \prod_{k=1}^K \prod_{r=1}^R \{p(\Delta N_{k,r} | x_k, z_r) p(w_{k,r} | x_k, z_r)\} \prod_{k=1}^K p(x_k | x_{k-1}; \sigma_\epsilon^2) \\ \prod_{r=1}^R p(z_r | z_{r-1}; \sigma_\delta^2). \quad (\text{A.4})$$

The log of the complete data likelihood is therefore

$$\log p(\Delta \mathbf{N}, \mathbf{x}, \mathbf{w}; \theta) \\ = \sum_{k=1}^K \sum_{r=1}^R \{ \log p(\Delta N_{k,r} | x_k, z_r) + \log p(w_{k,r} | x_k, z_r) \} + \pi(\mathbf{x}, \mathbf{z}) \quad (\text{A.5})$$

$$= \sum_{k=1}^K \sum_{r=1}^R \left[\Delta N_{k,r} \log \left(\frac{e^{x_k + z_r}}{1 + e^{x_k + z_r}} \right) + (1 - \Delta N_{k,r}) \log \left(\frac{1}{1 + e^{x_k + z_r}} \right) \right. \\ \left. + \log \left(\cosh \left(\frac{x_k + z_r}{2} \right) \right) \right. \\ \left. + \log \left(\sum_{i=1}^{\infty} (-1)^i \frac{(2i+1)}{\sqrt{2\pi} w_{k,r}^2} e^{-\frac{(2i+1)^2}{8w_{k,r}} - \frac{(x_k + z_r)^2 w_{k,r}}{2}} \right) \right] + \pi(\mathbf{x}, \mathbf{z}) \quad (\text{A.6})$$

$$= \sum_{k=1}^K \sum_{r=1}^R \left[\Delta N_{k,r} (x_k + z_r) - \log(1 + e^{x_k + z_r}) + \log \left(\frac{1 + e^{x_k + z_r}}{2e^{\frac{x_k + z_r}{2}}} \right) \right. \\ \left. + \log \left(e^{-\frac{(x_k + z_r)^2 w_{k,r}}{2}} \sum_{i=1}^{\infty} (-1)^i \frac{(2i+1)}{\sqrt{2\pi} w_{k,r}^2} e^{-\frac{(2i+1)^2}{8w_{k,r}}} \right) \right] + \pi(\mathbf{x}, \mathbf{z}) \quad (\text{A.7})$$

$$= \sum_{k=1}^K \sum_{r=1}^R \left[\Delta N_{k,r} (x_k + z_r) - \log(2) - \frac{x_k + z_r}{2} - \frac{(x_k + z_r)^2 w_{k,r}}{2} \right. \\ \left. + \log \left(\sum_{i=1}^{\infty} (-1)^i \frac{(2i+1)}{\sqrt{2\pi} w_{k,r}^2} e^{-\frac{(2i+1)^2}{8w_{k,r}}} \right) \right] \\ + \sum_{k=1}^K \left[\frac{1}{2} \log(2\pi \sigma_\epsilon^2) - \frac{(x_k - x_{k-1})^2}{2\sigma_\epsilon^2} \right] + \pi(\mathbf{x}, \mathbf{z}) \quad (\text{A.8})$$

$$\begin{aligned}
&= K \log(2) + \pi(\mathbf{x}, \mathbf{z}) + \sum_{k=1}^K \sum_{r=1}^R \left\{ \Delta N_{k,r}(x_k + z_r) - \frac{x_k + z_r}{2} - \frac{(x_k + z_r)^2 w_{k,r}}{2} \right\} \\
&\quad + \sum_{k=1}^K \sum_{r=1}^R \left\{ \log \left(\sum_{i=1}^{\infty} (-1)^i \frac{(2i+1)}{\sqrt{2\pi} w_{k,r}^2} e^{-\frac{(2i+1)^2}{8w_{k,r}}} \right) \right\}, \tag{A.9}
\end{aligned}$$

where

$$\begin{aligned}
\pi(\mathbf{x}, \mathbf{z}) &= \sum_{k=1}^K \left[\frac{1}{2} \log(2\pi\sigma_\epsilon^2) - \frac{(x_k - \rho_x x_{k-1} - \alpha_x u_{x,k})^2}{2\sigma_\epsilon^2} \right] \\
&\quad + \sum_{r=1}^R \left[\frac{1}{2} \log(2\pi\sigma_\delta^2) - \frac{(z_r - \rho_z z_{r-1} - \alpha_z u_{z,k})^2}{2\sigma_\delta^2} \right]. \tag{A.10}
\end{aligned}$$

From the complete data log likelihood, we see that

$$\log p(\mathbf{x}|\mathbf{\Delta N}, \mathbf{z}, \mathbf{w}; \theta)$$

$$\propto \sum_{k=1}^K \sum_{r=1}^R \left\{ \Delta N_{k,r}(x_k + z_r) - \frac{x_k + z_r}{2} - \frac{(x_k + z_r)^2 w_{k,r}}{2} \right\} + \pi(\mathbf{x}, \mathbf{z}) \tag{A.11}$$

$$\propto - \sum_{k=1}^K \sum_{r=1}^R \frac{1}{2} \frac{(\Delta \tilde{N}_{k,r} - (x_k + z_r))^2}{1/w_{k,r}} + \pi(\mathbf{x}, \mathbf{z}), \tag{A.12}$$

where $\Delta \tilde{N}_{k,r} = \frac{\Delta N_{k,r} - \frac{1}{2}}{w_{k,r}}$. Therefore, we can rewrite the augmented model as a linear gaussian state-space model as stated in theorem 1:

$$\begin{cases} x_k = \rho_x x_{k-1} + \alpha_x u_{x,k} + \epsilon_k, \epsilon_k \sim \mathcal{N}(0, \sigma_\epsilon^2) \\ \tilde{y}_k = x_k + \tilde{v}_k, \tilde{v}_k \sim \mathcal{N}\left(0, \left(\sum_{r=1}^R w_{k,r}\right)^{-1}\right) \\ \tilde{y}_k = \Delta \tilde{N}_k = x_k - \frac{K}{2} - \sum_{r=1}^R z_r w_{k,r} \end{cases}$$

Let $H_{k,r} = \Delta \tilde{N}_{1,r}, \dots, \Delta \tilde{N}_{k-1,r}$ denote the history of the observed process up to and including $k-1$. We can now write

$$p(\mathbf{x}|\mathbf{\Delta N}, \mathbf{z}, \mathbf{w}; \theta) \propto \prod_{k=1}^K p(\Delta N|x_k) p(x_k|H_k). \tag{A.13}$$

The forward filtering equations for this linear gaussian state-space model are

$$x_{k|k-1} = \rho_x x_{k-1|k-1} + \alpha_x u_{x,k}, \quad (\text{A.14})$$

$$\sigma_{k|k-1}^2 = \rho_x^2 \sigma_{k-1|k-1}^2 + \sigma_\epsilon^2, \quad (\text{A.15})$$

$$x_{k|k} = \rho_x x_{k|k-1} + \frac{\left(\sum_{r=1}^R w_{k,r} \right) \sigma_{k|k-1}^2}{1 + \left(\sum_{r=1}^R w_{k,r} \right) \sigma_{k|k-1}^2} \times \left(\frac{\alpha_x u_{x,k}}{\left(\sum_{r=1}^R w_{k,r} \right) \sigma_{k|k-1}^2} + \frac{\sum_{r=1}^R \Delta N_{k,r} - \frac{R}{2} - \sum_{r=1}^R z_r w_{k,r}}{\left(\sum_{r=1}^R w_{k,r} \right)} - \rho_x x_{k|k-1} \right), \quad (\text{A.17})$$

$$\sigma_{k|k}^2 = \frac{\sigma_{k|k-1}^2}{1 + \left(\sum_{r=1}^R w_{k,r} \right) \sigma_{k|k-1}^2}. \quad (\text{A.18})$$

After running the forward-filtering algorithm, we obtain $x_{K|K}$ and $\sigma_{K|K}^2$ from the final iteration of the filter. We can then draw $x_K \sim N(x_{K|K}, \sigma_{K|K}^2)$. Now we can treat x_K as the new observations and use the Kalman filter again to draw samples for x_{K-1} , and repeat this process iteratively for x_{K-1}, \dots, x_1 . The new observation equation reads

$$\begin{cases} x_k = \rho_x x_{k-1} + \epsilon_k, \epsilon_k \sim N(0, \sigma_\epsilon^2) \\ x_{k+1} = x_k + \epsilon_k \end{cases}.$$

From Bayes rule we have

$$p(x_k | x_{k+1}, H_k) = \frac{p(x_{k+1} | x_k) p(x_k | H_k)}{p(x_{k+1} | H_k)}. \quad (\text{A.19})$$

Denote the densities of $x_k | x_{k+1}, H_k$ as

$$x_k | x_{k+1}, H_k \sim N(x_{k|k}^*, \sigma_{k|k}^{2*}).$$




Figure 8: Plot of average (a) learning time and (b) learning trial identified by SMuRF inference as a function of ratio of neural spiking rate following and preceding learning. The learning time and trial are defined as the first time and trial pair when the empirical probability (see equation 3.25) that the spiking rate at a given time/trial pair (see Figure 2, point C) is bigger than the average rate at the same trial (see Figure 2, region A) and the average rate at the same time (see Figure 2, region B) is larger than 95%. The average is taken over 10 independently simulated rasters for each of the values for the ratio $\frac{\lambda_A}{\lambda_B}$. The figure demonstrates the ability of inference performed using the SMuRF model to reliably identify learning.

Then the update equations are

$$\log p(x_k|H_k) \propto \log p(x_{k+1}|x_k) + \log p(x_k|H_{k-1}), \quad (\text{A.20})$$

$$x_{k|k}^* = \rho_x x_{k|k-1} + \frac{\sigma_{k|k}^2}{\sigma_\epsilon^2} (x_{k+1} - \rho_x x_{k|k-1}), \quad (\text{A.21})$$

$$\sigma_{k|k}^{2*} = \frac{\sigma_\epsilon^2 \sigma_{k|k-1}^2}{\sigma_\epsilon^2 + \sigma_{k|k-1}^2}. \quad (\text{A.22})$$

With this backward-sampling algorithm, we can draw $x_k \sim N(x_{k|k}^*, \sigma_{k|k}^{2*})$, where $i = K - 1, \dots, 1$. The forward-filtering and backward-sampling algorithm are symmetric for x_k and z_r .

A.3 Initialization of the EM Algorithm and the Gibbs Sampler. We initialize the Monte Carlo EM algorithm with values for σ_ϵ^2 and σ_δ^2 obtained by applying the one-dimensional state-space model from Smith and Brown (2003) to the raster data aggregated across either trials or time. We initialize the Gibbs sampler using trajectories drawn from posterior distribution of the state in the one-dimensional state-space models (Smith & Brown, 2003)

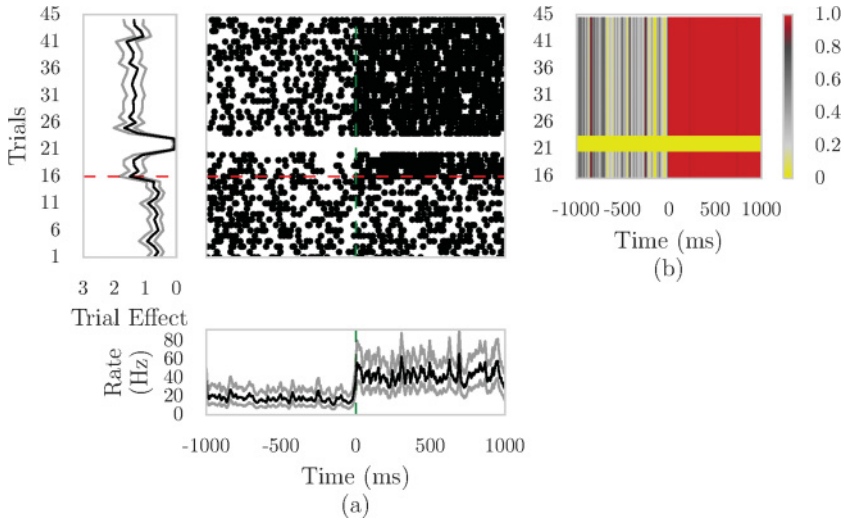


Figure 9: (a) Simulated neural spike raster with three consecutive error trials, along with estimated within- and cross-trial effects from the SMuRF model. The horizontal red line indicates the beginning of conditioning. The vertical green line indicates the onset of the conditioned stimulus. Despite the presence of the error trials, the left panel shows the cross-trial effect, which, following conditioning, increases above its average initial value of ≈ 1 to ≈ 2 . (b) Empirical probability that the spiking rate at a given trial and time is bigger than the average rate at the same trial and the average rate during habituation at the same time. Despite the presence of the error trials, we can see that with probability close to 1, the spiking rate at a given time/trial pair, following the conditioned stimulus and during conditioning (see Figure 2C), is bigger than the average rate at the same trial (see Figure 2A) and the average rate at the same time (see Figure 2B).

used to initialize σ_{ϵ}^2 and σ_{δ}^2 . The Gibbs sampler draws 5000 samples for \mathbf{x} , \mathbf{z} , and \mathbf{w} at every iteration. The algorithm reaches convergence when the absolute change in σ_{ϵ}^2 and σ_{δ}^2 is less than a certain threshold (10^{-5}).

A.4 Ability of the SMuRF Model to Identify Learning Time and Trial in Simulated Data. The results of our analysis of the cortical data in section 4 demonstrate that learning of a contingency by a neuron is a dynamic process that cannot be easily quantified in terms of a static time and trial of learning. We also demonstrated (see Table 1) how to use inferences from the SMuRF model to identify a learning time and a trial.

Here, we use simulated data to determine the ability of the SMuRF model to identify learning time and trial when learning is accompanied by sustained changes in neural spiking following conditioned stimulus onset and

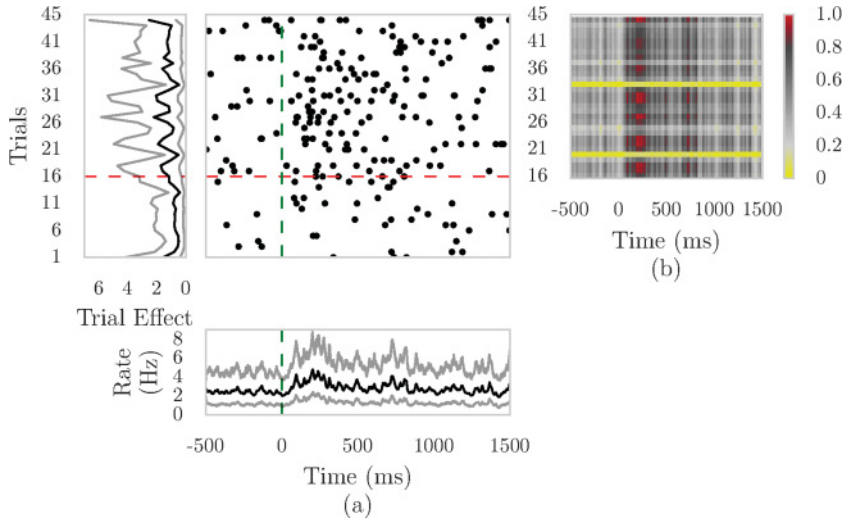


Figure 10: (a) Raster, along with SMuRF within and cross-trial components, from a cortical neuron that exhibits a transient conditioned response to the conditioned stimulus. The horizontal red line indicates the beginning of conditioning. The vertical green line indicates the onset of the conditioned stimulus. The bottom panel suggests that this neuron exhibits a delayed response to the conditioned stimulus, beginning at less than ≈ 100 ms following conditioned stimulus presentation. The response then decreases and is followed by a slight increase at ≈ 700 ms. Accounting for this increase in the within-trial spiking rate due to the conditioned stimulus, the left panel shows a multiplicative increase in spiking rate due to conditioning from an average initial value of ≈ 1 to a peak average value of ≈ 2 between trial ≈ 18 at trial 30. (b) Empirical probability that the spiking rate at a given trial and time is bigger than the average rate at the same trial and the average rate during habituation at the same time (refer to Figure 2). This panel provides a more detailed account of the intricate dynamics in neural modulation that accompanies conditioning to the stimulus for this neuron.

during conditioning. In particular, we assess the sensitivity of our method to the extent of the change in neural spiking rate following conditioned stimulus onset and during conditioning.

We simulated neural spike raster data in the same manner as described in section 4.1. In that section, the raster is divided into two regions (see Figure 3). We assume that the rate of spiking of the neuron in region B is fixed and equals $\lambda_B = 20$ Hz. We vary the rate of spiking λ_A of the neuron in region A from 20 to 45 Hz in 5 Hz increments. For each value of λ_A , we simulated 10 independent rasters and determined the learning time and trial as in Table 1. We use the average over the 10 rasters as the learning time/trial pair. When our method detects no change, we declare the learning time

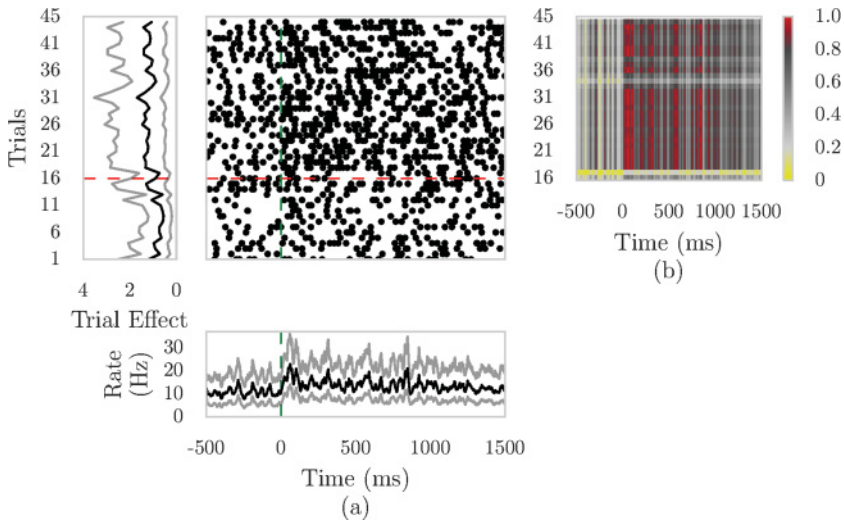


Figure 11: (a) Raster, along with SMuRF within and cross-trial components, from a cortical neuron that exhibits a transient conditioned response to the conditioned stimulus. The horizontal red line indicates the beginning of conditioning. The vertical green line indicates the onset of the conditioned stimulus. The bottom panel suggests that this neuron exhibits a delayed response to the conditioned stimulus, beginning at ≈ 20 ms following the conditioned stimulus presentation. The response is sustained until ≈ 800 ms and then decreases. Accounting for this increase in the within-trial spiking rate due to the conditioned stimulus, the left panel shows a multiplicative increase in spiking rate due to conditioning from an average initial value of ≈ 1 to a peak average value of less than ≈ 2 between trial ≈ 18 at trial 30. (b) Empirical probability that the spiking rate at a given trial and time is bigger than the average rate at the same trial and the average rate during habituation at the same time (refer to Figure 2). This panel provides a more detailed account of the intricate dynamics in neural modulation that accompanies conditioning to the stimulus for this neuron.

and trial as the last time and trial pair in the simulated data (i.e., 1000 ms and trial). Figures 8a and 8b show the averages of the identified learning times and trials as a function of the ratio $\frac{\lambda_A}{\lambda_B}$. The true learning time is at 0 ms with respect to conditioned stimulus onset and the true learning trial is trial 16. The figures demonstrate that the inference performed from the SMuRF model is able to detect the true learning time and trial when the rate in region A is 1.8 and 2 times larger than that in region B. Moreover, the lower the ratio $\frac{\lambda_A}{\lambda_B}$, the larger the delay is. The intuitive reason that it is easier to determine the learning trial is that for a given trial, there are many more observations compared to the number of trials for a give time instant.

A.5 Robustness of SMuRF Model to the Presence of Error Trials. We simulated neural spike raster data in the same manner as described in section 4.1. We picked three consecutive trials, starting from trial 21, to be error trials in which all of the observations were 0. Note that the data were simulated in the same manner as in Figure 4, except for the presence of the error trials. Figure 9 shows the result of applying the SMuRF model to these simulated raster data. (See also Figures 10 and 11.) The presence of the error trials does not affect our remarks for Figure 4.

Acknowledgments

We thank Anne C. Smith for her generous feedback on the manuscript for this letter and extensive discussions regarding the SMuRF model. D.B. thanks the Alfred P. Sloan Foundation. K.T. is a New York Stem Cell Foundation–Robertson Investigator and McKnight Scholar, and this work was supported by funding from the JPB Foundation, the PIIF and PIIF Engineering Award, PNDRE, JFDP, Alfred P. Sloan Foundation, New York Stem Cell Foundation, McKnight Foundation, R01-MH102441-01 (NIMH), RF1-AG047661-01 (NIA), R01-AA023305-01 (NIAAA), and NIH Directors New Innovator Award DP2-DK-102256-01 (NIDDK).

References

- Allsop, S. A., Vander Weele, C. M., Wichmann, R., & Tye, K. M. (2014). Optogenetic insights on the relationship between anxiety-related behaviors and social deficits. *Frontiers in Behavioral Neuroscience*, 8, 241.
- Allsop, S. A., Wichmann, R., Mills, F., Felix-Ortiz, A. C., Viennel, A., Beyeler, A., . . . Tye, K. M. (2017). *Observational learning requires cortico-amygdala transfer of socially-derived information*. Manuscript submitted for publication.
- Andrews, D. F., & Mallows, C. L. (1974). Scale mixtures of normal distributions. *Journal of the Royal Statistical Society, Series B (Methodological)*, 36, 99–102.
- Besag, J. (1974). Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society, Series B (Methodological)*, 36, 192–236.
- Casella, G. (2001). Empirical Bayes Gibbs sampling. *Biostatistics*, 2(4), 485–500.
- Casella, G., & George, E. I. (1992). Explaining the Gibbs sampler. *American Statistician*, 46(3), 167–174.
- Choi, H. M., & Hobert, J. P. (2013). The poly-gamma Gibbs sampler for Bayesian logistic regression is uniformly ergodic. *Electronic Journal of Statistics*, 7, 2054–2064.
- Chopin, N., Jacob, P. E., & Papaspiliopoulos, O. (2013). SMC2: An efficient algorithm for sequential analysis of state space models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 75(3), 397–426.
- Czanner, G., Eden, U. T., Wirth, S., Yanike, M., Suzuki, W. A., & Brown, E. N. (2008). An analysis of between-trial and within-trial neural spiking dynamics. *J. Neurophysiol.*, 99(5), 2672–2693.

- Dempster, A. P., Laird, N. M., & Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B (Methodological)*, 39, 1–38.
- Frühwirth-Schnatter, S. (1994). Data augmentation and dynamic linear models. *Journal of Time Series Analysis*, 15(2), 183–202.
- Polson, N. G., Scott, J. G., & Windle, J. (2013). Bayesian inference for logistic models using Pólya–Gamma latent variables. *Journal of the American Statistical Association*, 108(504), 1339–1349.
- Rad, K. R., & Paninski, L. (2010). Efficient, adaptive estimation of two-dimensional firing rate surfaces via gaussian process methods. *Network: Computation in Neural Systems*, 21(3–4), 142–168.
- Scott, J., & Pillow, J. W. (2012). Fully Bayesian inference for neural models with negative-binomial spiking. In F. Pereira, C. J. C. Burges, L. Bottou, & K. Q. Weinberger (Eds.), *Advances in neural information processing systems*, 25 (pp. 1898–1906). Red Hook, NY: Curran.
- Smith, A. C., & Brown, E. N. (2003). Estimating a state-space model from point process observations. *Neural Computation*, 15(5), 965–991.
- Truccolo, W., Eden, U. T., Fellows, M. R., Donoghue, J. P., & Brown, E. N. (2005). A point-process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *J. Neurophysiol.*, 93, 1074–1089.
- Wirth, S., Yanike, M., Frank, L. M., Smith, A. C., Brown, E. N., & Suzuki, W. A. (2003). Single neurons in the monkey hippocampus and learning of new associations. *Science*, 300(5625), 1578–1581.
- Yuan, K., Girolami, M., & Niranjana, M. (2012). Markov chain Monte Carlo methods for state-space models with point process observations. *Neural Computation*, 24(6), 1462–1486.
- Zammit-Mangion, A., Sanguinetti, G., & Kadirkamanathan, V. (2012). Variational estimation in spatiotemporal systems from continuous and point-process observations. *IEEE Transactions on Signal Processing*, 7(60), 3449–3459.