

Bayesian spike inference from calcium imaging data

Eftychios A. Pnevmatikakis, Josh Merel, Ari Pakman and Liam Paninski

Department of Statistics

Center for Theoretical Neuroscience

Grossman Center for the Statistics of Mind

Columbia University, New York, NY

Abstract—We present efficient Bayesian methods for extracting neuronal spiking information from calcium imaging data. The goal of our methods is to sample from the posterior distribution of spike trains and model parameters (baseline concentration, spike amplitude etc) given noisy calcium imaging data. We present discrete time algorithms where that the existence of a spike at each time bin using Gibbs methods, as well as continuous time algorithms that sample over the number of spikes and their locations at an arbitrary resolution using Metropolis-Hastings methods for point processes. We provide Rao-Blackwellized extensions that (i) marginalize over several model parameters and (ii) provide smooth estimates of the marginal spike posterior distribution in continuous time. Our methods serve as complements to standard point estimates and allow for quantification of uncertainty in estimating the underlying spike train and model parameters.

I. INTRODUCTION

Calcium imaging is an increasingly popular technique for large scale data acquisition in neuroscience [1]. The method detects underlying single neuron activity indirectly through observations of fluorescent indicators for calcium concentration. A key problem in the analysis of calcium imaging data is the inference of exact spike times from the noisy calcium signal, which has slower dynamics compared to neural spiking and is sampled at a relatively low acquisition rate. A variety of methods have been proposed to deal with this problem, including particle filtering [2], fast nonnegative deconvolution for approximate maximum-a-posteriori (MAP) inference [3], greedy template matching [4], and methods for estimating signals with finite rate of innovation [5]. In these methods, parameter estimation is typically performed offline or in an iterative manner, using for example the expectation maximization (EM) algorithm.

In this paper we propose Bayesian methods for sampling from the joint posterior distribution of the spike train and the model parameters given the fluorescence observations. We present two efficient approaches for sampling the spikes. The first is a discrete time binary sampler that samples whether a spike occurred at each timebin using Gibbs sampling. By exploiting the weak interaction between spikes at distant timebins, we show that a full sample can be obtained with just $O(T)$ complexity, where T is the number of timebins, and that parallelization is also possible. Our second sampler operates in continuous time and samples the number of spikes and the spike times at arbitrary resolution using Metropolis-Hastings (MH) techniques. We use a proposal distribution to move the spike

times around a local neighborhood that is based on the resulting signal residual. This proposal distribution enables fast mixing and tractable inference; each full Gibbs update is obtained with just $O(K)$ complexity, where K is the total number of spikes, rendering this algorithm particularly efficient for recordings that are sparse and/or are obtained at a fine resolution. Moreover, in high-SNR conditions, this method enables super-resolution spike inference (i.e. determining where each spike occurred within each timebin) and smooth estimates of the marginal spike posterior using a Rao-Blackwellized scheme. We also show that is possible to marginalize over several of the model parameters and derive collapsed Gibbs samplers that exhibit faster mixing.

II. MODEL DESCRIPTION AND BLOCK GIBBS SAMPLING

We assume that we observe a single neuron calcium trace for a duration of T timesteps. Let $\mathbf{s} \in \{0, 1\}^T$ the binary spiking vector of the neuron that indicates the existence of the spike at each timebin. The calcium activity that is generated by \mathbf{s} can be described by a simple first order autoregressive process¹ as

$$c(t) = \gamma c(t-1) + As(t),$$

where γ is a discrete time constant with $0 < \gamma < 1$, A indicates the amplitude of each spike, and $c(1) = c_1 + As(1)$, with c_1 an initial condition for the calcium concentration. Our fluorescence observation vector \mathbf{y} , can be written as

$$\mathbf{y}(t) = c(t) + b + \varepsilon_t,$$

where b is the baseline concentration and $\varepsilon_t \sim \mathcal{N}(0, \sigma^2)$ is some random Gaussian noise. Our goal is to estimate the spiking vector \mathbf{s} given the observation vector \mathbf{y} , which we assume is normalized in the interval $[0, 1]$. Approximate MAP methods [3] have been shown to perform well under high SNR assumptions. However in the low SNR regime their performance degrades, and the parameter estimation becomes more challenging. To overcome these limitations we introduce a block-Gibbs sampler that produces samples from the joint posterior distribution of the spikes and model parameters given \mathbf{y} . Let $G \in \mathbb{R}^{T \times T}$ and $\mathbf{v} \in \mathbb{R}^T$ be

$$G = \begin{bmatrix} 1 & 0 & \dots & 0 \\ -\gamma & 1 & \dots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \dots & -\gamma & 1 \end{bmatrix}, \quad \mathbf{v} = \begin{bmatrix} 1 \\ \gamma \\ \vdots \\ \gamma^{T-1} \end{bmatrix}.$$

¹Generalization to higher order AR processes is straightforward.

By denoting $\theta = [A, b, c_1]^T$, the likelihood can be written as

$$p(\mathbf{y}|\mathbf{s}, \theta, \sigma^2) \propto \exp\left(-\frac{A^2}{2\sigma^2}\mathbf{s}^T G^{-T} G^{-1} \mathbf{s} + \frac{A}{\sigma^2}\mathbf{s}^T G^{-T} \tilde{\mathbf{y}}\right),$$

with $\tilde{\mathbf{y}} = \mathbf{y} - b\mathbf{1}_T - c_1\mathbf{v}$ ($\mathbf{1}_T$ denotes a vector of ones of length T). We here place an i.i.d. Bernoulli process prior on the spike trains so the probability of a spike in any timebin is π . Under this uniform spiking assumption the discrete time constant γ can be estimated robustly from the autocovariance function of \mathbf{y} . For the prior probability π we adopt a Monte Carlo EM scheme where at each step we set π equal to its maximum likelihood estimate (MLE) $\pi_{\text{MLE}} = \mathbf{1}_T^T \mathbf{s} / T$. An alternative approach is to set a hyper-prior $\pi \sim \text{Beta}(\alpha, \beta)$, where α, β reflect any prior knowledge. In practice we found these approaches to be largely equivalent. For the rest of parameters θ we assume a joint half-normal (nonnegative) distribution, $\theta \sim \mathcal{N}(\boldsymbol{\mu}, \Sigma) \mathbf{1}_{\{\theta \in \mathbb{R}_+^3\}}$. The parameters $\boldsymbol{\mu}$ and Σ can also be learned from the data, but we selected general values that assume little prior knowledge and model a wide prior. Other prior choices, e.g. exponential, yield practically the same results. Finally, for the noise variance σ^2 we set an inverse Gamma prior $\sigma^2 \sim \text{InvGamma}(1, 0.1)^2$, which is a weak and relatively flat prior for both high and low-SNR regions for traces \mathbf{y} normalized to the $[0, 1]$ interval. Under these assumptions, the block Gibbs sampler proceeds as follows to draw samples from the joint posterior:

$$\begin{aligned} \log p(\mathbf{s}|\pi, \theta, \sigma^2, \mathbf{y}) &\propto -\frac{1}{2}\mathbf{s}^T W \mathbf{s} + \mathbf{s}^T \left(\frac{A}{\sigma^2} G^{-T} \tilde{\mathbf{y}} + l\mathbf{1}_T\right) \\ \pi &= \mathbf{1}_T^T \mathbf{s} / T \\ \theta|\mathbf{s}, \sigma^2, \mathbf{y} &\sim \mathcal{N}(\Lambda(\sigma^{-2} S^T \mathbf{y} + \Sigma^{-1} \boldsymbol{\mu}), \Lambda) \mathbf{1}_{\{\theta \in \mathbb{R}_+^3\}} \\ \sigma^2|\mathbf{s}, \theta, \mathbf{y} &\sim \text{InvGamma}(1 + T/2, 0.1 + \|\mathbf{y} - S\theta\|^2/2), \end{aligned}$$

with $W = A^2(G^{-T}G^{-1})/\sigma^2$, $l = \log(\pi/(1-\pi))$, S is a $T \times 3$ matrix that depends on the current state of \mathbf{s} , given by

$$S = [G^{-1} \mathbf{s}, \mathbf{1}_T, \mathbf{v}],$$

and $\Lambda = (\Sigma^{-1} + \sigma^{-2} S^T S)^{-1}$. We now turn to the main problem of sampling from the posterior of the spike vector \mathbf{s} .

III. DISCRETE BINARY SAMPLER

To sample \mathbf{s} we take a Gibbs-sampling approach where at each iteration we sample $s(i)$ given the current state of the rest of the entries. We first use a standard Metropolized Gibbs (MG) sampler with the flip of each binary entry being the possible move [6], and defer until section V for a more efficient collapsed Gibbs sampler. If $s_i^{(k)}$ is the state of $s(i)$ at the k -th sample, and define $\mathbf{s}_{\text{cur}} = [s_{1:i-1}^{(k)}, s_{i:T}^{(k-1)}]^T$ the current state of \mathbf{s} . The sampling algorithm flips $s_i^{(k)}$ with probability

$$\mathbb{P}_{\text{flip}}(i) = \min\left(1, \frac{\mathbb{P}(\mathbf{s}_{\text{flipped}})}{\mathbb{P}(\mathbf{s}_{\text{cur}})}\right) \quad (1)$$

²The variance σ^2 can also be estimated from the autocovariance function via the Yule-Walker equations but we favor a Bayesian approach here.

The ratio inside (1) can be more efficiently computed in the log-domain. By denoting $\alpha = \frac{A}{\sigma^2} G^{-T} \tilde{\mathbf{y}} + l\mathbf{1}_T$ the log-ratio becomes

$$(1 - 2s_i^{(k-1)})((2s_i^{(k-1)} - 1)[W]_{ii}/2 - \mathbf{s}_{\text{cur}}^T W \mathbf{e}_i + \alpha(i)), \quad (2)$$

with \mathbf{e}_i the i -th standard basis vector. It is important to note that the log-ratio of (2) can be computed efficiently in just $O(1)$ time. The matrix G^{-1} is lower triangular and Toeplitz with $[G^{-1}]_{ij} = \gamma^{i-j} \mathbf{1}_{i \geq j}$. Therefore it can be approximated by a banded Toeplitz matrix with bandwidth that depends on γ . Consequently W can again be approximated by a banded matrix and is also asymptotically Toeplitz for $T \rightarrow \infty$. It follows that the products $\mathbf{s}_{\text{cur}}^T W \mathbf{e}_i$ can be computed in just $O(1)$ time (technically in $O(m/\log(1/\gamma))$ for m bits of accuracy). This gives a total $O(T)$ complexity per full sweep to obtain a new sample from $p(\mathbf{s}|\mathbf{y}, \pi, \theta, \sigma^2)$. For large T , the algorithm can also be parallelized – notice that the entries of \mathbf{s} which are at least $L = O(1/\log(1/\gamma))$ timesteps apart, are approximately conditionally independent. \mathbf{s} can be partitioned into chunks of length L , that can be sampled in parallel.

IV. CONTINUOUS TIME SAMPLER

The spike samplers presented in section III operate on a discrete domain, with the length of each timebin set by the frame rate of the calcium imaging. While simple and effective, these methods can have several disadvantages. First, when the calcium signal is acquired by raster scanning, the frame rate is typically in the range of 10-30Hz. The length of each timebin is too large to assume that the underlying neuron can fire at most 1 spike per bin. In certain applications, higher resolution of the spike time is needed. A typical example is in “connectomics” where the order with which the various neurons fire spikes is crucial for determining their connectivity [7], [8]. Moreover, the complexity of the discrete samplers scales with the number of observed timebins (i.e. with the temporal resolution). In experiments where this resolution is very fine, binary samplers can become computationally expensive. Compounding this, neural activity may be very sparse and thus sampling at intervals where no spikes occur is uninformative. Instead, a method that scales with the number of spikes (instead of the number of bins) is more desirable.

To address these issues we propose a sampler that samples directly the spike times in continuous time, given the calcium observations. The state space corresponds to the set of spike times, and transdimensional moves change the dimensionality of this state space. At every iteration we sample the new time of the i -th spike t_i given the current location of the rest of the spikes and the hyperparameters. Since the number of spikes K is in general unknown, we also allow for spike insertions and deletions at each iteration. Similar approaches have been proposed before, see e.g. [9] for the context of signals with finite rate of innovation, where however the number of “spikes” K was considered known. In the continuous time setup, the calcium evolution is described by the differential equation

$$\frac{d}{dt} c(t) = -\frac{1}{\tau} c(t) + A s(t),$$

where τ is the (continuous) time constant of the calcium indicator. For a timebin of size Δ with $\Delta \ll \tau$ the discrete and continuous time constants are related through $\gamma = 1 - \Delta/\tau$. If t_1, \dots, t_k are the spike times, denoted by \mathfrak{T} , the calcium signal is given by

$$c_{\mathfrak{T}}(t) = c_1 e^{-\frac{t}{\tau}} + A \sum_{k=1}^T e^{-\frac{t-t_k}{\tau}} \mathbf{1}_{\{t \geq t_k\}}, \quad (3)$$

where the subscript \mathfrak{T} indicates the dependence of the calcium signal on the set of spike times \mathfrak{T} , and the observations are interpreted as $y(n) = c(n\Delta) + b + \varepsilon_n$, giving a likelihood of the form (we ignore c_1 for simplicity)

$$\log p(\mathbf{y}|\mathfrak{T}, \boldsymbol{\theta}, \sigma^2) \propto -\frac{1}{2\sigma^2} \|\mathbf{y} - b\mathbf{1}_T - c_{\mathfrak{T}}\|^2,$$

where $c_{\mathfrak{T}} = [c_{\mathfrak{T}}(\Delta), \dots, c_{\mathfrak{T}}(T\Delta)]^T$. Finally, since the algorithm now operates in the continuous domain, we replace the i.i.d. Bernoulli prior for the spikes with a homogeneous continuous Poisson process prior with parameter λ . If we set a prior $\lambda \sim \text{Gamma}(\alpha, \beta)$ (with β denoting the rate), then

$$\lambda|\mathfrak{T} \sim \text{Gamma}(\alpha + |\mathfrak{T}|, \beta + T\Delta).$$

For the hyper-parameters α, β we can adopt an empirical Bayes approach. For a spike sample \mathfrak{T} the evidence function [10] can be written as:

$$p(\mathfrak{T}|\alpha, \beta) = \int_0^\infty p(\mathfrak{T}|\lambda) p(\lambda|\alpha, \beta) d\lambda = \frac{\beta^\alpha \Gamma(\alpha + |\mathfrak{T}|)}{(\beta + T\Delta)^{\alpha + |\mathfrak{T}|} \Gamma(\alpha)}.$$

$p(\mathfrak{T}|\alpha, \beta)$ is maximized for $\alpha, \beta \rightarrow \infty$ with $\beta = \alpha\Delta T/|\mathfrak{T}|$. By choosing these values the posterior of λ puts all the mass in the maximum likelihood estimate $\lambda_{\text{MLE}} = |\mathfrak{T}|/\Delta T$, and therefore this approach coincides with a Monte-Carlo EM algorithm. Alternatively, we can use the empirical Bayes Gibbs sampling method proposed in [11], which at the m -th sample chooses

$$(\alpha^{(m)}, \beta^{(m)}) = \arg \max_{\alpha, \beta} \prod_{i=1}^m p(\mathfrak{T}^{(i)}|\alpha, \beta),$$

with solution $\alpha^{(m)}, \beta^{(m)} \rightarrow \infty$ with $\beta^{(m)} = \alpha^{(m)} m \Delta T / (\sum_{i=1}^m |\mathfrak{T}^{(i)}|)$. We next describe the sampling of the spike times in more detail.

A. Move existing spikes

We can update individual spike times either using random walk Metropolis-Hastings (MH) moves or a Gibbs algorithm. For the MH algorithm, we propose a move $t_i \rightarrow t'_i$, using a Gaussian density, centered at t_i and with standard deviation equal to e.g. 10Δ , as a proposal distribution. Since this proposal is symmetric, the proposed move is accepted with probability

$$\mathbb{P}(t_i \rightarrow t'_i) = \min \left(1, \frac{p(\mathfrak{T}_{\text{new}}|\mathbf{y}, \boldsymbol{\theta}, \sigma^2)}{p(\mathfrak{T}_{\text{cur}}|\mathbf{y}, \boldsymbol{\theta}, \sigma^2)} \right). \quad (4)$$

Similarly to the discrete case, each spike contributes an exponentially decaying calcium trace (3), that has a temporally localized effect to the whole vector $c_{\mathfrak{T}}$. By exploiting this, we can implement the spike shifting operation in just $O(1)$ time, provided that we keep in memory the residual vector $\mathbf{w} \in \mathbb{R}^T$

with $\mathbf{w} = \mathbf{y} - b\mathbf{1}_T - c_{\mathfrak{T}}$. After any operation $\mathfrak{T}_{\text{cur}} \rightarrow \mathfrak{T}_{\text{new}}$, the vectors \mathbf{c} and \mathbf{w} can be locally updated, and the ratio in (4) can be computed in just $O(1)$ time.

An alternative random walk MH method is to use full Gibbs sampling. We can sample each new location t'_i from the likelihood $p(\mathbf{y}|\mathfrak{T}_{\setminus i} \cup \{t'_i\}, \boldsymbol{\theta}, \sigma^2)$. While this approach mixes with less samples than using MH (intuitively the spike can move to anywhere, instead of only locally), the computational cost for moving each spike is $O(T)$ which is undesirable. In practice we use a localized version of Gibbs; for each spike time, we construct a distribution from the residual between the data and the current calcium signal, restricted to small time interval \mathcal{I} centered on the current spike time, i.e., $\mathcal{I} = [t_i - L, t_i + L]$ with e.g., $L = 10\Delta$. The proposal distribution can then be expressed as

$$\log A(t_i \rightarrow t'_i) \propto -\frac{1}{2\sigma^2} \sum_{n:n\Delta \in \mathcal{I}} (y(n) - b - c_{\mathfrak{T}_{\setminus i} \cup \{t'_i\}}(n\Delta))^2,$$

where $\mathfrak{T}_{\setminus i}$ is the set of spike times excluding t_i . Note that in this case, the proposal distribution is no longer symmetric (since the local intervals are different), and therefore the Hastings ratio is also included in the acceptance probability. However, this can also be computed in $O(1)$ time³ and thus this scheme remains efficient. Finally we note that this proposal distribution can be used to derive a Rao-Blackwellized scheme for updating the continuous time posterior spike distribution, instead of using the actual spike samples. This method works quite well empirically.

B. Adding and removing spikes

The sampler over spike trains also requires transdimensional moves in order to sample over the number of spikes. To add or remove spikes, we follow a standard birth-death MH algorithm [12]. We choose a fixed probability z for proposing new spikes and uniform proposal densities for adding and removing spikes. Suppose we want to add a new spike $t_{k+1} = \xi$ to the existing set of spikes \mathfrak{T} . This is accepted with probability

$$\mathbb{P}(\mathfrak{T} \rightarrow \mathfrak{T} \cup \{\xi\}) = \min(1, r_{\text{birth}}),$$

$$\text{with } r_{\text{birth}} = \frac{p(\mathfrak{T} \cup \{\xi\}|\mathbf{y}, \boldsymbol{\theta}, \lambda, \sigma^2)(1-z)q_d(\mathfrak{T} \cup \{\xi\}, \{\xi\})}{p(\mathfrak{T}|\mathbf{y}, \boldsymbol{\theta}, \lambda, \sigma^2)zq_b(\mathfrak{T}, \{\xi\})},$$

where $q_d(\mathfrak{T} \cup \{\xi\}, \{\xi\})$ is the proposal probability for removing $\{\xi\}$ from $\mathfrak{T} \cup \{\xi\}$, and $q_b(\mathfrak{T}, \{\xi\})$ is the proposal density for adding $\{\xi\}$ to \mathfrak{T} . $p(\mathfrak{T} \cup \{\xi\}|\mathbf{y}, \boldsymbol{\theta}, \lambda, \sigma^2)$ and $p(\mathfrak{T}|\mathbf{y}, \boldsymbol{\theta}, \lambda, \sigma^2)$ are posterior probabilities of the spike train given the data and the spiking prior respectively after and before the proposed moves. Under uniform proposals and $|\mathfrak{T}| = K$, we have $q_d(\mathfrak{T} \cup \{\xi\}, \{\xi\}) = 1/(K+1)$, $q_b(\mathfrak{T}, \{\xi\}) = (\Delta T)^{-1}$, giving

$$r_{\text{birth}} = \frac{p(\mathbf{y}|\mathfrak{T} \cup \{\xi\}, \boldsymbol{\theta}, \sigma^2)(1-z)\lambda\Delta T}{p(\mathbf{y}|\mathfrak{T}, \boldsymbol{\theta}, \sigma^2)z(K+1)}.$$

³Technically, this computation scales with the resolution at which we want to discretize the proposal distribution. In practice one can use a coarse resolution to choose a bin from the proposal distribution and then sample inside this bin uniformly.

Similarly the acceptance ratio for removing spike $t_i = \eta$ is

$$\begin{aligned} r_{\text{death}} &= \frac{p(\mathcal{T} \setminus \{\eta\} | \mathbf{y}, \boldsymbol{\theta}, \sigma^2) z q_b(\mathcal{T} \setminus \{\eta\}, \{\eta\})}{p(\mathcal{T} | \mathbf{y}, \boldsymbol{\theta}, \sigma^2) (1-z) q_d(\mathcal{T}, \{\eta\})} \\ &= \frac{p(\mathbf{y} | \mathcal{T} \setminus \{\eta\}, \boldsymbol{\theta}, \sigma^2) z K}{p(\mathbf{y} | \mathcal{T}, \boldsymbol{\theta}, \sigma^2) (1-z) \lambda \Delta T}. \end{aligned}$$

Following [13], a typical choice of the prior proposal probability is $z = 1/2$, and we repeat the birth-death sampling process 10 times per iteration. Each iteration of the algorithm is schematically presented in Alg. 1.

Algorithm 1 Schematic representation of each iteration

for $i = 1$ to current number of spikes K **do**
 Sample $t_i | \mathcal{T} \setminus i, \lambda, \boldsymbol{\theta}, \sigma^2$ using MH.
for $j = 1$ to 10 **do**
 Propose addition of spikes. Update K .
 Propose removal of spikes. Update K .
 Sample parameters $\lambda, \boldsymbol{\theta}, \sigma^2$.

Similarly to the case of shifting spikes, the addition and removal can also be implemented in just $O(1)$ time. Using a similar argument it is also easy to show that $\boldsymbol{\theta}$ can also be sampled in $O(1)$ time using the method described in section II. As a result, the complexity of each iteration scales linearly with the number of spikes and not the number of timebins, rendering this algorithm particularly attractive for recordings that are very sparse and/or have a fine temporal resolution.

V. COLLAPSED GIBBS SAMPLER

It is possible to marginalize over the baseline b and initial value c_1 and enhance the mixing rates of our algorithm [14]. We present this approach here for the discrete case and note that the continuous case can be treated similarly. For this part we assume σ^2 to be fixed. By denoting the marginal prior for $[b; c_1] \sim \mathcal{N}(\boldsymbol{\mu}_b, \Sigma_b)$, the marginal likelihood is computed as⁴

$$\begin{aligned} p(\mathbf{y} | \mathbf{s}, A, \pi) &= \iint_{b, c_1 \in \mathbb{R}^2} p(\mathbf{y} | \mathbf{s}, p, \boldsymbol{\theta}) p(b, c_1) db dc_1 \\ &\propto \exp \left(-\frac{A^2}{2\sigma^2} \mathbf{s}^T G^{-T} V G^{-1} \mathbf{s} + \frac{A}{\sigma^2} \mathbf{s}^T G^{-T} V \mathbf{y}_b \right), \end{aligned}$$

with $\mathbf{y}_b = \mathbf{y} - [\mathbf{1}_T, \mathbf{v}] \boldsymbol{\mu}_b$

$$V = I + [\mathbf{1}_T, \mathbf{v}] (\sigma^2 \Sigma_b^{-1} + [\mathbf{1}_T, \mathbf{v}]^T [\mathbf{1}_T, \mathbf{v}])^{-1} [\mathbf{1}_T, \mathbf{v}]^T$$

and we used the fact that for $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \Phi)$ and a symmetric matrix C such that $\Phi^{-1} - C$ is positive definite, we have

$$\mathbb{E}_{\mathbf{x}} \left[e^{\frac{1}{2} \mathbf{x}^T C \mathbf{x} + \mathbf{b}^T \mathbf{x}} \right] = |I - \Phi C|^{-1/2} e^{\frac{1}{2} \mathbf{b}^T (\Phi^{-1} - C)^{-1} \mathbf{b}}.$$

The marginalized likelihood has the same functional form and therefore \mathbf{s} and A can again be sampled with the same methods. Moreover, by exploiting the structure of V it is easy to see that each multiplication of the form $V G^{-1} \mathbf{s}$ can be

⁴Technically this approach is approximate since $[b, c_1]$ have truncated normal priors but here are integrated over the whole \mathbb{R}^2 . However, in practice the posterior of $[b, c_1]$ puts very little (if not negligible) mass on negative values and the approximation is very tight.

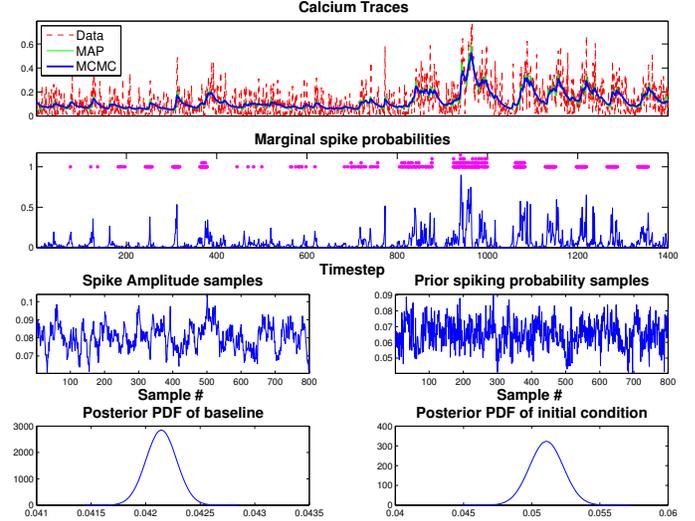


Fig. 1. Application of the collapsed, discrete Metropolisized-Gibbs sampler on real spinal cord data (in-vitro). Top row: Real data (red dashed), MAP estimate (green) and mean from 800 Gibbs samples (blue). Second row: Marginal spike probabilities per timestep. The true spike times are also shown (purple circles). Third row: 800 samples from the spike amplitude (left) and the prior spiking probability (right). Bottom row: Rao-Blackwellized estimates of the posterior distributions for the baseline (left) and initial value (right). The collapsed MG sampler mixes fast and provides low variance estimates for the model parameters.

performed in $O(1/\log(1/\gamma))$ time and therefore the complexity of each full Gibbs sweep still scales as $O(T)$ and the algorithm remains efficient. After the initial burn-in period, the (Rao-Blackwellized) posterior of the marginalized parameters b, c_1 can be approximated as

$$p(b, c_1 | \mathbf{s}, A) \approx \frac{1}{M} \sum_{i=1}^M p(b, c_1 | \mathbf{s}^{(i)}, A^{(i)}),$$

where $\mathbf{s}^{(i)}, A^{(i)}$ are the sampled values and each $p(b, c_1 | \mathbf{s}^{(i)}, A^{(i)})$ can be computed by conditioning $p(\boldsymbol{\theta} | \mathbf{s}^{(i)})$. Note that σ^2 is treated as a known parameter here. If we assume an inverse Gamma prior then the posterior is no longer inverse Gamma. However σ^2 can still be sampled with standard rejection sampling methods. Finally, marginalization over A is also possible; however the marginalized posterior of \mathbf{s} has no longer the nice quadratic form, and the computational complexity of the algorithm increases.

VI. RESULTS

We apply our methods to calcium imaging data from spinal cord neurons in-vitro, collected using the calcium indicator GCaMP6s with a temporal resolution 15Hz [15]. The neurons were stimulated antidromically (so we know the spike times) and fired small bursts of spikes. In several timebins multiple spikes were fired. Fig. 1 shows an application of the discrete algorithm with b, c_1 marginalized out. 1000 samples were collected with the first 200 discarded (burn-in period). This particular trace is of low SNR, however the algorithm predicts

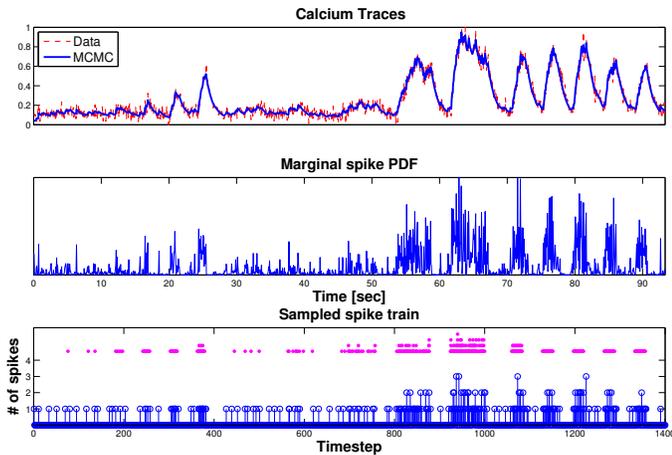


Fig. 2. Application of the collapsed, discrete Metropolized-Gibbs sampler on real spinal cord data (in-vitro). Top row: Real data (red dashed) and mean from 700 Gibbs samples (blue). Second row: Estimated marginal spike PDF in continuous time. Third row: A typical samples of the spike train in continuous time, binned in the resolution defined by the imaging rate. The true spike times are also shown (purple circles). The continuous time sampler can assign multiple spikes at each timebin and provide better approximation of the spiking behavior in high-SNR conditions.

most of the spikes and provides low variance estimates of the model parameters.

A limitation of the discrete algorithm is that it can only produce one spike per timebin. To resolve this issue we run the continuous time sampler on the same dataset. All the traces from imaged pixels that correspond to the same neuron were averaged to produce a high-SNR trace. The algorithm produced 500 samples after a burn in period of 200 samples. The results are shown in Fig. 2. Again the algorithm predicts well the produced calcium trace and provides estimates of the spike posterior in continuous time. By binning the produced sample in the original bin size (bottom row) we see that the algorithm can assign multiple spikes per timebin and better approximate the true spikes.

VII. CONCLUSIONS - FUTURE WORK

We presented two classes of Bayesian methods for spike train inference from calcium imaging data. Our methods provide a principled approach for estimating the posterior distribution of the spike trains and provide robust estimates of the model parameters, especially in low-SNR conditions. We also derived collapsed Gibbs samplers that exhibit faster mixing with no significant computational cost per sample.

In future work, we plan to explore the use of Hamiltonian Monte Carlo [16] and particle Markov chain Monte Carlo [17] methods for more efficient spike sampling, and extend our methods to allow for a slowly time-varying baseline, a phenomenon that is often observed in vivo experimental conditions. We also plan to scale up to a spatial setup where each measurement corresponds to a pixel that is part of a neuron (or is shared across a small number of neurons) [18], and in the case of compressive calcium imaging [19], where

each measurement is formed by projecting the calcium activity of the whole imaged spatial onto a random vector.

ACKNOWLEDGMENT

We thank S. Chandramouli and D. Carlson for useful discussions, and T. Machado for sharing his spinal cord data. LP is supported from an NSF CAREER award. This work is also supported by ARO MURI W911NF-12-1-0594.

REFERENCES

- [1] M. B. Ahrens, M. B. Orger, D. N. Robson, J. M. Li, and P. J. Keller, "Whole-brain functional imaging at cellular resolution using light-sheet microscopy," *Nature methods*, vol. 10, no. 5, pp. 413–420, 2013.
- [2] J. Vogelstein, B. Watson, A. Packer, R. Yuste, B. Jedynek, and L. Paninski, "Spike inference from calcium imaging using sequential monte carlo methods," *Biophysical journal*, vol. 97, no. 2, pp. 636–655, 2009.
- [3] J. Vogelstein, A. Packer, T. Machado, T. Sippy, B. Babadi, R. Yuste, and L. Paninski, "Fast non-negative deconvolution for spike train inference from population calcium imaging," *Journal of Neurophysiology*, vol. 104, no. 6, pp. 3691–3704, 2010.
- [4] B. F. Grewe, D. Langer, H. Kasper, B. M. Kampa, and F. Helmchen, "High-speed in vivo calcium imaging reveals neuronal network activity with near-millisecond precision," *Nature methods*, vol. 7, no. 5, pp. 399–405, 2010.
- [5] J. Oñativia, S. R. Schultz, and P. L. Dragotti, "A finite rate of innovation algorithm for fast and accurate spike detection from two-photon calcium imaging," *Journal of neural engineering*, vol. 10, no. 4, p. 046017, 2013.
- [6] J. S. Liu, "Metropolized Gibbs sampler: an improvement," Technical report, Dept. Statistics, Stanford Univ, Tech. Rep., 1996.
- [7] Y. Mishchenko, J. Vogelstein, and L. Paninski, "A Bayesian approach for inferring neuronal connectivity from calcium fluorescent imaging data," *The Annals of Applied Statistics*, vol. 5, no. 2B, pp. 1229–1261, 2011.
- [8] S. Keshri, E. Pnevmatikakis, A. Pakman, B. Shababo, and L. Paninski, "A shotgun sampling approach for solving the common input problem in neural connectivity inference." 2013, arXiv preprint arXiv:1309.372.
- [9] V. Y. F. Tan and V. K. Goyal, "Estimating signals with finite rate of innovation from noisy samples: A stochastic algorithm," *Signal Processing, IEEE Transactions on*, vol. 56, no. 10, pp. 5135–5146, 2008.
- [10] C. Bishop, *Pattern Recognition and Machine Learning*. Springer, 2006.
- [11] G. Casella, "Empirical Bayes Gibbs sampling," *Biostatistics*, vol. 2, no. 4, pp. 485–500, 2001.
- [12] J. Møller and R. P. Waagepetersen, *Statistical inference and simulation for spatial point processes*. CRC Press, 2003.
- [13] R. P. Adams, I. Murray, and D. J. MacKay, "Tractable nonparametric bayesian inference in poisson processes with gaussian process intensities," in *Proceedings of the 26th Annual International Conference on Machine Learning*. ACM, 2009, pp. 9–16.
- [14] J. S. Liu, "The collapsed Gibbs sampler in Bayesian computations with applications to a gene regulation problem," *Journal of the American Statistical Association*, vol. 89, no. 427, pp. 958–966, 1994.
- [15] T. Machado, L. Paninski, and T. Jessell, "Functional organization of motor neurons during fictive locomotor behavior revealed by large-scale optical imaging," in *SFN Neuroscience*, 2013.
- [16] A. Pakman and L. Paninski, "Auxiliary-variable exact Hamiltonian Monte Carlo samplers for binary distributions," in *Advances in Neural Information Processing Systems 26*, 2013.
- [17] C. Andrieu, A. Doucet, and R. Holenstein, "Particle Markov chain Monte Carlo methods," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 72, no. 3, pp. 269–342, 2010.
- [18] E. Pnevmatikakis, T. Machado, L. Grosenick, B. Poole, J. Vogelstein, and L. Paninski, "Rank-penalized nonnegative spatiotemporal deconvolution and demixing of calcium imaging data," in *Computational and Systems Neuroscience Meeting COSYNE*, 2013.
- [19] E. Pnevmatikakis and L. Paninski, "Sparse nonnegative deconvolution for compressive calcium imaging: algorithms and phase transitions," in *Advances in Neural Information Processing Systems 26*, 2013.