

Efficient methods for sampling spike trains in networks of coupled neurons

Yuriy Mishchenko and Liam Paninski
Columbia University
www.stat.columbia.edu/~liam

October 31, 2010

Abstract

Monte Carlo approaches have recently been proposed to quantify connectivity in neuronal networks. The key problem is to sample from the conditional distribution of a single neuronal spike train, given the activity of the other neurons in the network. Dependencies between neurons are usually relatively weak; however, temporal dependencies within the spike train of a single neuron are typically strong. In this paper we develop several specialized Metropolis-Hastings samplers which take advantage of this dependency structure. These samplers are based on two ideas: 1) an adaptation of fast forward-backward algorithms from the theory of hidden Markov models to take advantage of the local dependencies inherent in spike trains, and 2) a first-order expansion of the conditional likelihood which allows for efficient exact sampling in the limit of weak coupling between neurons. We also demonstrate that these samplers can effectively incorporate side information, in particular noisy fluorescence observations in the context of calcium-sensitive imaging experiments. We quantify the efficiency of these samplers in a variety of simulated experiments in which the network parameters are closely matched to data measured in real cortical networks, and also demonstrate the sampler applied to real calcium imaging data.

1 Introduction

One of the central goals of neuroscience is to understand how the structure of neural circuits underlies the processing of information in the brain, and in recent years a considerable effort has been focused on measuring neural connectivity empirically (Shepherd et al., 2003; Bureau et al., 2004; Briggman and Denk, 2006; Hagmann et al., 2007; Sato et al., 2007; Smith, 2007; Hagmann et al., 2008; Luo et al., 2008; Bohland, J. et al, 2009; Helmstaedter et al., 2009). “Functional” approaches to this neural connectivity problem rely on statistical analysis of neural activity observed with experimental techniques such as multielectrode extracellular recording (Hatsopoulos et al., 1998; Harris et al., 2003; Stein et al., 2004; Paninski, 2004; Truccolo et al., 2005; Santhanam et al., 2006; Luczak et al., 2007; Pillow et al., 2008) or calcium imaging (Tsien, 1989; Cossart et al., 2003; Yuste et al., 2006; Ohki et al., 2005). Although functional approaches (Brillinger, 1988; Nykamp, 2003; Nykamp, 2005a; Okatan et al., 2005; Srinivasan et al., 2006; Rigat et al., 2006; Nykamp, 2007; Yu et al., 2006; Kulkarni and Paninski, 2007) have not yet been demonstrated to directly yield the true physical structure of a neural circuit, the estimates obtained via this type of analysis play

an increasingly important role in attempts to understand information processing in neural circuits (Bureau et al., 2004; Sato et al., 2007; Pillow et al., 2008; Mishchenko et al., 2010).

Perhaps the biggest challenge for inferring neural connectivity from functional data — and indeed in network analysis more generally — is the presence of hidden nodes which are not observed directly (Nykamp, 2005b; Nykamp, 2007; Kulkarni and Paninski, 2007; Vidne et al., 2009; Pillow and Latham, 2007; Vakorin et al., 2009). Despite swift progress in simultaneously recording activity in massive populations of neurons, it is still beyond the reach of current technology to monitor a complete set of neurons that provide the presynaptic inputs even for a single neuron (though see e.g. (Petreanu et al., 2009) for some recent progress in this direction). Since estimation of functional connectivity relies on the analysis of the inputs to target neurons in relation to their observed spiking activity, the inability to monitor all inputs can result in persistent errors in the connectivity estimation due to model misspecification. Developing a principled and robust approach for incorporating such unobserved neurons, whose spike trains are unknown and constitute hidden or latent variables, is an area of active research in connectivity analysis (Nykamp, 2007; Pillow and Latham, 2007; Vidne et al., 2009; Vakorin et al., 2009).

Incorporating these latent variables into connectivity estimation is a challenging task. The standard approach for estimating the connectivity requires us to compute the likelihood of the observed neural activity given an estimate of the underlying connectivity. Computing this likelihood, however, requires us to integrate out the probability distribution over the activity of all hidden neurons. This latent activity variable will typically have very high dimensionality, making any direct integration methods infeasible.

Thus it is natural to turn to Markov chain Monte Carlo (MCMC) approaches here (Rigat et al., 2006). The best design of such an MCMC sampler is not at all obvious in this setting, since it may be necessary to develop sophisticated proposal densities to capture the dependence of the hidden spike trains on the observed spiking data in order to guarantee a reasonable proposal acceptance rate (Pillow and Latham, 2007). For example, the simplest Gibbs sampling approaches may not perform well given the strong dependencies between adjacent spiking timebins that are inherent to neural activity.

In this paper we develop Metropolis-Hastings (MH) algorithms for efficiently sampling from the probability distribution over the activity of hidden neurons. We derive a proposal density that is asymptotically correct in the limit of weak interneuronal couplings, and demonstrate the utility of this proposal density on simulated networks of neurons with neurophysiologically feasible parameters (Sayer et al., 1990; Abeles, 1991; Braitenberg and Schuz, 1998; Gomez-Urquijo et al., 2000; Song et al., 2005; Lefort et al., 2009). We also consider a hybrid MH strategy based on fast hidden Markov model (HMM) sampling approaches that can be exploited to extend the range of applicability of the basic algorithm to more strongly coupled networks of neurons. In each case, the resulting MCMC chain mixes quickly and each step of the chain may be computed quickly.

Of special interest is the problem of sampling from the probability distribution over unknown spike trains when a neuron is indirectly observed using calcium imaging (instead of electrical recordings) (Tsien, 1989; Yuste et al., 2006; Cossart et al., 2003; Ohki et al., 2005; Vogelstein et al., 2009). This problem can be naturally related to the sampling problem described above. We develop an approach to incorporate calcium fluorescence imaging observations directly into our proposal density. This allows us to efficiently sample from the probability distribution over the activity of multiple hidden neurons given calcium fluorescence observations, improving on the methods introduced in (Mishchenko et al., 2010).

2 Methods

2.1 Model definition

We model the activity of individual neurons with a discrete-time generalized linear model (GLM) (Brillinger, 1988; Chornoboy et al., 1988; Brillinger, 1992; Plesser and Gerstner, 2000; Paninski et al., 2004; Paninski, 2004; Rigat et al., 2006; Truccolo et al., 2005; Nykamp, 2007; Kulkarni and Paninski, 2007; Pillow et al., 2008; Vidne et al., 2009; Stevenson et al., 2009):

$$\begin{aligned} n_i(t) &\sim \text{Bernoulli}[f(J_i(t))\Delta], \\ J_i(t) &= b_i(t) + \sum_{j=1}^N \sum_{t' < t} w_{ij}(t-t')n_j(t'), \end{aligned} \quad (1)$$

where the spike indicator function for neuron i , $n_i(t) \in \{0, 1\}$, is defined on time bins t of size Δ . Connectivity between neurons is described via the “connectivity matrix,” $w_{ij}(t)$; the self-terms $w_{ii}(t-t')$ describe refractory effects (there is a minimal “refractory” interspike interval of a millisecond or two in most neurons, for example), and $w_{ij}(t-t')$ represents the statistical effect of a spike in neuron j at time t' upon the spiking rate of neuron i at time t . N is the number of neurons in the neural population, including hidden and observed neurons, and $b_i(t)$ denotes a baseline driving term (which might depend on some observed covariates such as a sensory stimulus). We assume that the maximal firing frequency Δ^{-1} is much larger than the typical spiking rate, $\Delta^{-1} \gg E[n_i(t)]$, in which case the Bernoulli spiking model in (1) is closely related to models in which $n_i(t)$ is drawn from a Poisson distribution at every time step (see the references above for a number of examples).

In the presence of fluorescence observations from a calcium-sensitive indicator dye (the calcium imaging setting), the model (1) is supplemented by two variables per neuron i : the intracellular calcium concentration $C_i(t)$ (which is not observed directly) and the fluorescence observations $F_i(t)$. We model these terms according to a hidden Markov model governed by a simple driven autoregressive process (Vogelstein et al., 2009; Mishchenko et al., 2010),

$$\begin{aligned} C_i(t) &= C_i(t-\Delta) - \frac{\Delta}{\tau_i^c}(C_i(t-\Delta) - C_i^b) + A_i n_i(t), \\ F_i(t) &\sim \mathcal{N}[S(C_i(t)), V(C_i(t))]. \end{aligned} \quad (2)$$

Thus, under nonspiking conditions, $C_i(t)$ is set to the baseline level of C_i^b . Whenever the neuron fires a spike, $n_i(t) = 1$, the calcium variable $C_i(t)$ jumps by a fixed amount A_i , and subsequently decays with time constant τ_i^c ; τ_i^c is on the order of hundreds of milliseconds in the cases of most interest here. The fluorescence signal $F_i(t)$ corresponds to the count of photons collected at the detector per neuron per imaging frame. This photon count may be modeled with approximately Gaussian statistics (Poisson photon count models are also tractable in this context (Vogelstein et al., 2009), though we will not pursue this detail here), with the mean given by a saturating Hill-type function $S(C) = C/(C + K_d)$ (Yasuda et al., 2004) and the variance $V(C)$ scaling with the mean. See (Vogelstein et al., 2009) for full details and further discussion. Note that observations of calcium-indicator fluorescence are typically performed at a low “frame-rate”, FR , measured in frames per second. Thus we will restrict our attention to the case $FR < \Delta^{-1}$, i.e., $F_i(t)$ observations are available only every few timesteps Δ . In addition, it will be useful to define an effective SNR, following (Mishchenko et al., 2010), as

$$eSNR = \frac{E[F_i(t) - F_i(t-\Delta) \mid n_i(t) = 1]}{E[(F_i(t) - F_i(t-\Delta))^2/2 \mid n_i(t) = 0]^{1/2}}, \quad (3)$$

i.e., the size of a spike-driven fluorescence jump divided by a rough measure of the standard deviation of the baseline fluorescence.

In (Mishchenko et al., 2010) we introduced an expectation-maximization approach for estimating the model parameters given calcium fluorescence data; (Pillow and Latham, 2007) discuss a related approach for fitting the model given spiking data recorded from extracellular electrodes. The maximization step in this context is often quite tractable, involving a separable convex optimization problem (Smith and Brown, 2003; Paninski, 2004; Kulikarni and Paninski, 2007; Escola and Paninski, 2008). The expectation step is much more difficult; analytical approaches are often infeasible. Monte Carlo solutions to this problem require us to obtain samples from the probability distribution over the activity of all hidden neurons. This problem is the focus of the present paper. In the context of the expectation maximization framework, we assume therefore that an estimate of the model parameters θ is available, and the problem is to obtain a sample from $P(\mathbf{n}_{hidden}|\mathbf{n}_{observed};\theta)$ (in the case that the spike trains of a subset of neurons is observed directly), or $P(\mathbf{n}_{hidden}|\mathbf{F}_{observed};\theta)$ (in the case that spike trains are observed indirectly via calcium fluorescence measurements). Here and throughout we use bold notation for vectors, i.e. $\mathbf{n}(t) = \{n_i(t), i = 1, \dots, N\}$ and $\mathbf{n} = \{n_i(t), i = 1, \dots, N; t \in [0, T]\}$; θ denotes the set of all parameters in the model, including $\mathbf{b}, \mathbf{w}, \tau^c$, and \mathbf{A} . In cases where no confusion is possible below we will suppress the dependence on θ .

2.2 A block-wise Gibbs approach for sampling from the distribution over activity of hidden neurons

In (Mishchenko et al., 2010), we noted that in order to sample from the desired joint distribution over the activity of all hidden neurons, $P(\mathbf{n}_{hidden}|\cdot)$, it is sufficient to be able to efficiently sample sequentially from the conditional distribution over one hidden neuron given all of the other hidden neurons, $P(\mathbf{n}_i|\mathbf{n}_{\setminus i};\cdot)$: if this is possible, then a sample from the full joint distribution, $P(\mathbf{n}_{hidden}|\cdot)$, can be obtained using a block-wise Gibbs algorithm. This blockwise Gibbs approach makes sense in this context because the connectivity weights $w_{ij}, i \neq j$, are relatively weak in many neural contexts; for example, synaptic strengths between cortical neurons are typically fairly small (see the discussion in section 2.7 below). However, dependencies between the spike indicator variables $\{n_i(s), n_i(t)\}$ within a single neuron may be quite large if the time gap $|s - t|$ is small, since the self-terms w_{ii} are typically strong over small timescales; for example, as mentioned above, after every spike a neuron will enter a refractory state during which it is unable to spike again for some time period. See (Toyoizumi et al., 2008) for further discussion.

Thus we will focus below exclusively on the single-neuron conditional sampling problem, $\mathbf{n}_i \sim P(\mathbf{n}_i|\mathbf{n}_{\setminus i};\cdot)$. In the next couple sections we will discuss methods for sampling from $P(\mathbf{n}_i|\mathbf{n}_{\setminus i})$; in section 2.6 we will discuss adaptations of these methods for incorporating calcium fluorescence observations in $P(\mathbf{n}_i|\mathbf{n}_{\setminus i}, \mathbf{F}_i)$.

2.3 Exact sampling via the hidden Markov model forward-backward procedure

In some special cases we can solve the single-neuron conditional sampling problem quite explicitly. Specifically, if the support of each of the coupling terms $w_{ij}(t)$ is contained in the interval $[0, K\Delta]$ for some sufficiently small number of time steps K , then we can employ

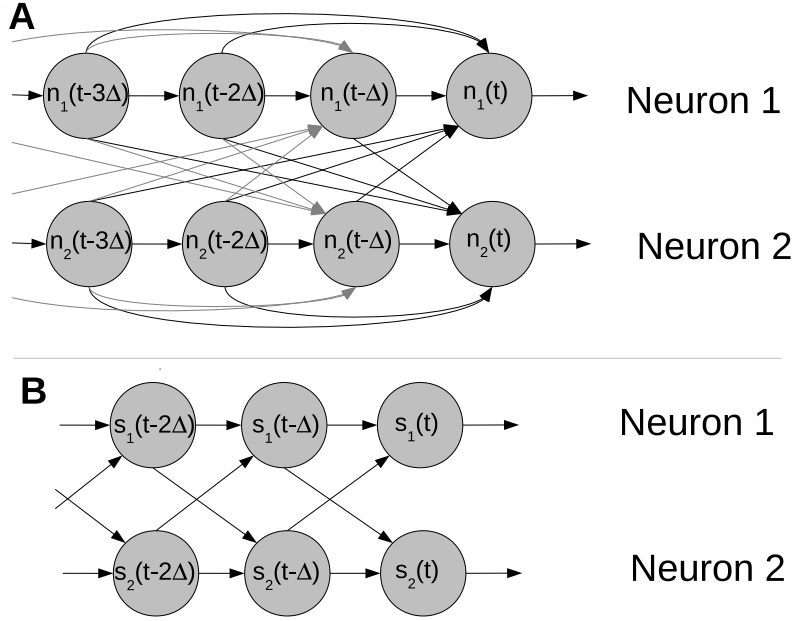


Figure 1: Graphical model describing (A) the K -order Markov model for neural dynamics in terms of the instantaneous spikes, $n_i(t)$, and (B) a simpler Markov model in terms of the grouped states $s_i(t) = \{n_i(t'), t - K\Delta < t' \leq t\}$.

standard hidden Markov model methods to sample from the desired conditional distribution exactly. The key is to note that in this case Eq. (1) defines a K th-order Markov process; thus it is convenient to apply the standard substitution $s_i(t) = \{n_i(t'), t - K\Delta < t' \leq t\}$. Here $s_i(t)$ denotes the “state” of neuron i , which is completely characterized by the binary K -string (e.g. $s_i(t) = \text{“010011100...”}$) describing its spiking over the K most recent time bins. Clearly, the size of the state-space is 2^K . See Fig. 1 for an illustration.

To exploit this Markov structure, note that we may write the conditional distributions $P[s_i(t)|s_i(t-\Delta), \mathbf{s}_{\setminus i}(t-\Delta)]$ and $P[\mathbf{s}_{\setminus i}(t+\Delta)|\mathbf{s}_{\setminus i}(t), s_i(t)]$ explicitly using model (1), and in fact it is easy to see that the system

$$\begin{aligned} s_i(t) &\sim P[s_i(t)|s_i(t-\Delta), \mathbf{s}_{\setminus i}(t-\Delta)], \\ \mathbf{s}_{\setminus i}(t+\Delta) &\sim P[\mathbf{s}_{\setminus i}(t+\Delta)|\mathbf{s}_{\setminus i}(t), s_i(t)] \end{aligned} \quad (4)$$

forms an autoregressive hidden Markov model (Rabiner, 1989), with $\mathbf{s}_{\setminus i}(t+\Delta)$ playing the role of the observation at time t (Fig. 1). Thus, if K is not too large, standard finite hidden Markov model (HMM) techniques can be applied to obtain the desired spike train samples. Specifically, we can use the standard filter forward-sample backward algorithm (Rabiner, 1989; de Gunst et al., 2001) to obtain samples from $P(\mathbf{s}_i|\mathbf{s}_{\setminus i})$, from which we can easily transform back to obtain the desired sample \mathbf{n}_i . This requires $O(T2^K)$ computation time and $O(2^K)$ memory, since the transition matrix $P[s_i(t)|s_i(t-\Delta), \mathbf{s}_{\setminus i}(t-\Delta)]$ is sparse (due to the redundant definition of the state variable $s_i(t)$), so the matrix-vector multiplication required in the transition step scales linearly with the number of states, instead of quadratically, as would be seen in the general case. Note, for clarity, that in the above discussion we are

assuming that the coupling terms $w_{ij}(t)$ are known (or have been estimated in a separate step), and therefore K is also known, directly from the maximal support of the couplings $w_{ij}(t)$; thus we do not apply any model selection or estimation techniques for choosing K here.

2.4 Efficient sampling in the weak-coupling limit

Clearly the HMM strategy described above will become ineffective if K becomes larger than about 10 or so, due to the exponential growth of the Markov state-space. Therefore it is natural to seek more general alternate methods. Recall that the interneuronal coupling terms $w_{ij}, i \neq j$, are small. Thus it seems promising to examine the conditional spike train probability in the limit of weak coupling ($w_{ij} \rightarrow 0$), in the hope that a good sampling method might suggest itself.

The log-likelihood of the hidden spike train \mathbf{n}_i can be written as

$$\log P(\mathbf{n}_i | \mathbf{n}_{\setminus i}) = \log P(\mathbf{n}_i, \mathbf{n}_{\setminus i}) + \text{const.} = \sum_i \sum_t n_i(t) \log f[J_i(t)] - f[J_i(t)]\Delta - n_i(t)! + \text{const.}; \quad (5)$$

for notational simplicity, we have used a Poisson model for $n_i(t)$ given the past spike trains, but it turns out that we can use any exponential family with linear sufficient statistics in the following computations; see the appendix for details. Now, if we make the abbreviation

$$J_i^-(t) = \sum_{s < t} \sum_j w_{ij}(t-s)n_j(s), \quad (6)$$

we can easily expand to first order in w_{ij} :

$$\begin{aligned} \log P(\mathbf{n}_i, \mathbf{n}_{\setminus i}) &= \sum_t n_i(t) \log f[J_i(t)] - f[J_i(t)]\Delta - n_i(t)! \\ &\quad + \sum_{j \neq i} \sum_t n_j(t) \log f[b_j(t) + J_j^-(t)] - f[b_j(t) + J_j^-(t)]\Delta + \text{const.} \\ &= \sum_t n_i(t) \log f[J_i(t)] - f[J_i(t)]\Delta - n_i(t)! \\ &\quad + \sum_{j \neq i} \sum_t n_j(t) \log f[b_j(t)] - f[b_j(t)]\Delta + \left(n_j(t) \frac{f'[b_j(t)]}{f[b_j(t)]} - f'[b_j(t)]\Delta \right) J_j^-(t) + o(w) + \text{const.} \\ &= \sum_t n_i(t) \log f[J_i(t)] - f[J_i(t)]\Delta - n_i(t)! \\ &\quad + \sum_{j \neq i} \sum_t \left(n_j(t) \frac{f'[b_j(t)]}{f[b_j(t)]} - f'[b_j(t)]\Delta \right) \sum_{s < t} \sum_k w_{jk}(t-s)n_k(s) + o(w) + \text{const.} \\ &= \sum_t n_i(t) \log f[J_i(t)] - f[J_i(t)]\Delta - n_i(t)! \\ &\quad + \sum_{j \neq i} \sum_t \sum_{s > t} \left(n_j(s) \frac{f'[b_j(s)]}{f[b_j(s)]} - f'[b_j(s)]\Delta \right) w_{ji}(s-t)n_i(t) + o(w) + \text{const.} \\ &= \sum_t n_i(t) \left\{ \log f[J_i(t)] + \sum_{j \neq i} \sum_{s > t} \frac{f'[b_j(s)]}{f[b_j(s)]} w_{ji}(s-t) [n_j(s) - f[b_j(s)]\Delta] \right\} - n_i(t)! + o(w) + \text{const.}, \end{aligned} \quad (7)$$

where we have used the definition of $J_j^-(t)$ in the third equality, rearranged sums (and changed variables) over t and s in the fourth equality, regathered terms in the final equality, and retained only terms involving a factor of $n_i(t)$ throughout.

Thus if we note the resemblance between eq. (7) and the Poisson log-probability, we see that one attractive approach is to sample from a spike train proposal $n_i(t)$ with log-rate equal to the term within brackets,

$$\log f[J_i(t)] + \sum_{j \neq i} \sum_{s > t} \frac{f'[b_j(s)]}{f[b_j(s)]} w_{ji}(s-t) [n_j(s) - f[b_j(s)]\Delta]; \quad (8)$$

since we are conditioning on $\{n_j(t)\}$ (i.e., these terms are assumed fixed), and there are no $n_i(s)$ terms affecting the proposed rate of $n_i(t)$ for $s > t$, it is straightforward to sample $n_i(t)$ recursively forward according to this rate, and then use Metropolis-Hastings to compute the required acceptance probability and obtain a sample from the desired distribution.

We note that this approach is conceptually quite similar to that of (Pillow and Latham, 2007), who proposed sampling recursively from a process of the form

$$\begin{aligned} n_i(t) &\sim \text{Poisson}[f[J_i^{PL}(t)]\Delta], \\ J_i^{PL}(t) &= b_i^{PL} + \sum_j \sum_{s < t} w_{ij}^{PL}(t-s)n_j(s) + \sum_{j \neq i} \sum_{s > t} w_{ij}^{PL}(t-s)n_i(s). \end{aligned} \quad (9)$$

(Pillow and Latham, 2007) discussed a somewhat computationally-intensive procedure in which the parameters b_i^{PL} and w_{ij}^{PL} are reoptimized iteratively via a maximum pseudolikelihood procedure. If we reverse the logic we used to obtain equation (7), it is clear that our approach simply uses an “effective input”

$$\tilde{J}_i(t) = b_i(t) + J_i^-(t) + \sum_{j \neq i} \sum_{s > t} \frac{f[b_i(t)]}{f'[b_i(t)]} \frac{f'[b_j(s)]}{f[b_j(s)]} w_{ji}(s-t) [n_j(s) - f[b_j(s)]\Delta]. \quad (10)$$

in place of $J_i^{PL}(t)$ above. Further, in the special case of an exponential nonlinearity, $f(J) = \exp(J)$, all pre-factors in Eq.(10) cancel out, leading to an expression that resembles eq. (9) quite closely:

$$\tilde{J}_i(t) = b_i(t) + \sum_j \sum_{s < t} w_{ij}(t-s)n_j(s) + \sum_{j \neq i} \sum_{s > t} w_{ji}(s-t) [n_j(s) - f[b_j(s)]\Delta]. \quad (11)$$

Thus, in the weak-coupling limit $w \rightarrow 0$, and ignoring self-interaction terms w_{ii} , we see that the optimal “back-inputs” $w_{ij}^{PL}(t-s)$, $s > t$, from (Pillow and Latham, 2007) may be analytically identified as the time- and index-reversed original forward couplings, $w_{ji}(-t)$.

2.5 Hybrid HMM-Metropolis-Hastings approaches

So far we have developed two methods for sampling from $P(\mathbf{n}_i | \mathbf{n}_{\setminus i})$: the HMM method (section 2.3) is exact but becomes inefficient when K is large, while the weak-coupling method (section 2.4) becomes inefficient when the coupling terms w become large. What is needed is a hybrid approach that combines the strengths of these two methods. The key is that the cross-coupling terms w_{ij} are typically fairly weak (as emphasized above), and the self-coupling terms $w_{ii}(t)$ are large only for a small number of time delays t .

To take advantage of this special structure, we begin by constructing a truncated HMM that retains all coupling terms $w_{ij}(t)$ for t up to some maximal delay t_{max} , where t_{max} is chosen to keep the size of the state-space acceptably small but large enough so that the coupling terms are captured to an acceptable degree. Then we include the discarded coupling terms (i.e., terms at lags longer than t_{max}) via the first-order approximation. More concretely, denote the conditional probability of $\{n_i(t)\}$ under the truncated HMM as

$$\prod_t P_{hmm}(s_i(t)|s_i(t-\Delta); \mathbf{n}_{\setminus i}) \quad (12)$$

(recall the correspondence between the K -string state variable $s_i(t)$ and the binary spiking variable $n_i(t)$). Note that this forms an inhomogeneous Markov chain in the grouped state variables $s_i(t)$, and an inhomogeneous (t_{max}/Δ) -Markov chain in $n_i(t)$, as discussed in section 2.3. Now we want to incorporate the discarded coupling terms up to first order: we simply form the product

$$\prod_t P_{hmm}(s_i(t)|s_i(t-\Delta); \mathbf{s}_{\setminus i}) \exp \left(n_i(t) \sum_{j \neq i} \sum_{t' > t_{max}} w_{ji}(t' - t) \frac{f'(b_j(t'))}{f(b_j(t'))} [n_j(t') - f(b_j(t'))\Delta] \right). \quad (13)$$

We use this simple product form here in an analogy to the HMM case, in which we condition on observations by simply forming the (normalized) product of the prior distribution on state variables and the likelihood of the observations given the state variables; in this case, the first term plays the role of the prior over the state variables $s_i(t)$ and the second term corresponds to the pseudo-observations represented by the weak coupling terms incorporating the observed firing of the other neurons n_j . Although, to be clear, this joint probability does not correspond rigorously to any small-parameter expansion of $P(\mathbf{n}_i|\mathbf{n}_{\setminus i})$, we might expect that this hybrid may outperform alternative strategies such as using only the truncated HMM, given by Eq. (12), or only the weak-coupling correction, given by Eq. (13), as we will see in the Results section below (see especially Fig. 6). Since this product form for the joint probability is structurally equivalent to the conditional probability of an HMM in the state variables $s_i(t)$ — more precisely, the graphical model (Jordan, 1999) corresponding to the above expression is a chain in terms of $s_i(t)$ — we can employ the forward-backward procedure to sample from this proposal, and then compute the Metropolis-Hastings acceptance probability to obtain samples from the desired conditional $P(\mathbf{n}_i|\mathbf{n}_{\setminus i})$.

Note that the MH acceptance probability will decrease as a function of the dimensionality T of the desired spike train and of the size of the discarded weights $w_{ij}(t), t > t_{max}$, since for large w the proposal density discussed above will approximate the target density less accurately. In general, it is necessary to employ a blockwise approach: i.e., we update the spike train $\{n_i(t)\}$ in blocks whose length is chosen to be small enough that the MH acceptance probability is sufficiently high, but large enough so that the total correlations between blocks are small and the overall chain mixes quickly. We will examine these tradeoffs in more quantitative detail in the Results section below.

2.6 Efficient sampling given calcium fluorescence imaging observations

In this section we turn our attention to the problem of sampling from the fluorescent-conditional distribution $\mathbf{n}_i \sim P(\mathbf{n}_i|\mathbf{n}_{\setminus i}, \mathbf{F}_i)$. In principle, we could apply the same basic approach as before, exploiting the HMM structure of equations (1-2) in the variables

$\{s_i(t), C_i(t)\}$; however, the state-space for the calcium variable $C_i(t)$ is continuous (instead of discrete), requiring us to adapt our methods somewhat.

We will briefly mention two alternative methods before introducing the novel approach that is the focus of this section. First, in (Mishchenko et al., 2010) we introduced a sampler based on a technique from (Neal et al., 2003). This method requires drawing a rather large auxiliary sample from the continuous $C_i(t)$ state space and then employing a modified forward-backward approach to obtain spike train samples. The techniques we will discuss below do not require such an auxiliary sample, and are therefore significantly more computationally efficient; we will not compare these methods further here. Second, standard pointwise Gibbs sampling is fairly straightforward in this setting: we write $P(\mathbf{F}_i, \mathbf{n}) = P(\mathbf{F}_i|\mathbf{n}_i)P(\mathbf{n})$, and then note that both $P(n_i(t)|\mathbf{n}_{\setminus i,t})$ and $P(\mathbf{F}_i|\mathbf{n}_i)$ can be computed easily as a function of $n_i(t)$. (To compute the latter quantity, note from eq. (2) that that $n_i(t)$ only affects the values of $F_i(t')$ appreciably for $t < t' < n_c\tau_i^c$, for a suitably large number of time constants n_c .) We will discuss the performance of the Gibbs approach further below.

Now we turn to the main theme of this section: to adapt the MH-based blockwise approach discussed above, it is essential to develop a proposal density that efficiently incorporates the observed fluorescence data, in order to ensure reasonable acceptance rates. We use a filter-backwards-and-sample-forward approach. The basic idea is to compute, via a backwards recursion, the conditional future observation density $P(F_i(t : T)|s_i(t), C_i(t); \mathbf{n}_{\setminus i})$ given the current state $(s_i(t), C_i(t))$, for all $t \leq T$. Then, we may easily sample from a proposal of the form

$$\begin{aligned} s_i(t) &\sim P\left(F_i(t : T)|s_i(t), C_i(t); \mathbf{n}_{\setminus i}\right) \times P\left(s_i(t)|s_i(t - \Delta); \mathbf{n}_{\setminus i}\right), \\ C_i(t) &= C_i(t - \Delta) - \Delta/\tau_i^c(C_i(t - \Delta) - C_i^b) + A_i n_i(t), \end{aligned} \quad (14)$$

and by appending the samples $s_i(t)$, $0 < t \leq T$, we obtain a sample from the desired density $P(\mathbf{n}_i|\mathbf{n}_{\setminus i}, \mathbf{F}_i)$. Here the spiking term $P(s_i(t)|s_i(t - \Delta); \mathbf{n}_{\setminus i})$ may include weak-coupling terms or truncations to keep the state-space tractably bounded, as discussed in the previous section. (Again, recall the correspondence between the K -string state variable $s_i(t)$ and the binary spiking variable $n_i(t)$.)

To calculate $P(F_i(t : T)|s_i(t), C_i(t))$, we can use the standard backward HMM recursion (Rabiner, 1989),

$$P(F_i(t : T)|s_i(t), C_i(t)) = P(F_i(t)|C_i(t))P(F_i(t + \Delta : T)|s_i(t), C_i(t)), \quad (15)$$

with

$$\begin{aligned} P(F_i(t + \Delta : T)|s_i(t), C_i(t)) &= \sum_{s_i(t+\Delta)} \int dC_i(t + \Delta) P(F_i(t + \Delta : T)|s_i(t + \Delta), C_i(t + \Delta)) \times \\ &\quad \times P(s_i(t + \Delta)|s_i(t); \mathbf{n}_{\setminus i}) P(C_i(t + \Delta)|s_i(t + \Delta), C_i(t)). \end{aligned} \quad (16)$$

We have already discussed the transition probability $P(s_i(t+\Delta)|s_i(t); \mathbf{n}_{\setminus i})$. $P(F_i(t : T)|s_i(t), C_i(t))$ is a continuous function of $C_i(t)$. We could solve this backwards recursion directly by breaking the $C_i(t)$ axis into a large number of discrete intervals and employing standard numerical integration methods to compute the required integrals at each time step. However, this approach is computationally expensive in the high SNR regime (i.e., where the ratio of the spike-driven calcium bump A_i is large relative to the fluorescence noise scale $\sqrt{V(C_i(t))}$), where a fine discretization becomes necessary. (Vogelstein et al., 2009) introduced a more efficient approximate recursion for this density that we adapt here. The first step is to approximate $P(F_i(t : T)|s_i(t), C_i(t))$ with a mixture of Gaussians; this approximation is exact

in the limit of linear and Gaussian fluorescence observations (i.e., in the case that the $S(\cdot)$ is linear and $V(\cdot)$ is constant in eq. (2)), and works reasonably in practice (see (Vogelstein et al., 2009) for further details).

It is straightforward to see in this setting that we require $2^K 2^{T-t}$ mixture components to represent $P(F_i(t:T)|s_i(t), C_i(t))$ exactly; each term in the mixture corresponds to a distinct sequences of spikes $\mathbf{n}_i(t:T)$ and initial conditions $s_i(t)$. For large T , such a mixture of course cannot be computed explicitly. However, as (Vogelstein et al., 2009) pointed out, we may further approximate the intractable $2^K 2^{T-t}$ mixture with a smaller $2^K(T-t+1)$ mixture, parametrized by the total number of spikes $\sum_{t'=t}^T n_i(t')$ instead of the full spike sequence $\mathbf{n}_i(t:T)$ (since mixture components with the same total number of spikes overlap substantially, due to the long timescale τ_i^c of the calcium decay); thus we may avoid any catastrophic exponential growth in the complexity of the representation.

In order to calculate and update this approximate mixture we proceed as follows. At each timestep t we represent $P(F_i(t:T)|s_i(t), C_i(t))$ as a mixture of $2^K(T-t+1)$ Gaussian components with weights $p_{s,k}$, means $m_{s,k}$, and variances $v_{s,k}$, for 2^K distinct states s and $k = 1, \dots, T-t+1$. In order to update this mixture backwards one timestep, from time t to time $t-\Delta$, we first integrate over $C_i(t+\Delta)$ in Eq.(16). Since $C_i(t)$ evolves deterministically given $n_i(t)$, this reduces to updating the means and the variances in each Gaussian component,

$$\begin{aligned} m_{s,k} &\rightarrow [m_{s,k} - \Delta/\tau_i^c C_b^i - A_i n_i(t)] / [1 - \Delta/\tau_i^c] \\ v_{s,k} &\rightarrow v_{s,k} / (1 - \Delta/\tau_i^c)^2 \\ p_{s,k} &\rightarrow p_{s,k} / (1 - \Delta/\tau_i^c) \end{aligned} \quad (17)$$

Second, we perform the multiplication with the observation density $P(F_i(t)|C_i(t))$ in Eq.(15). In the case that $F_i(t)$ is linear and Gaussian in $C_i(t)$, each such product is again Gaussian, and the means and the variances are updated as follows ($M_c(t)$ and $V_c(t)$ are the mean and the variance for $P(F_i(t)|C_i(t))$, respectively):

$$\begin{aligned} m_{s,k} &\rightarrow (M_c(t-\Delta)v_{s,k} + m_{s,k}V_c(t-\Delta))/(v_{s,k} + V_c(t-\Delta)) \\ v_{s,k} &\rightarrow V_c(t-\Delta)v_{s,k}/(v_{s,k} + V_c(t-\Delta)) \\ p_{s,k} &\rightarrow p_{s,k} \times (2\pi(v_{s,k} + V_c(t-\Delta)))^{-1/2} \exp \left[-\frac{M_c(t-\Delta)^2}{2V_c(t-\Delta)} - \frac{m_{s,k}^2}{2v_{s,k}} + \frac{(M_c(t-\Delta)v_{s,k} + m_{s,k}V_c(t-\Delta))^2}{2(v_{s,k} + V_c(t-\Delta))V_c(t-\Delta)v_{s,k}} \right]. \end{aligned} \quad (18)$$

More generally (i.e., in the case of a nonlinear $S(\cdot)$ or nonconstant $V(\cdot)$ in eq. (2)), standard Gaussian approximations may be used to arrive at a similar update rule; again, see (Vogelstein et al., 2009) for details.

Third, we perform the summation over $s_i(t+\Delta)$ in Eq. (16), and reorganize the obtained mixture to reduce the number of components from $2(T-t+1)$ to $(T-t+2)$ for each $s_i(t)$. Eq.(16) doubles each mixture component into two new Gaussians, corresponding to the two terms in the sum for $n_i(t+\Delta) = 1$ or $n_i(t+\Delta) = 0$. For brevity, we denote these terms as $w^+ = P(s^+(t+\Delta)|s_i(t))$ and $w^- = P(s^-(t+\Delta)|s_i(t))$, where $s^+(t+\Delta)$ stands for the state $s_i(t+\Delta)$ describing a spike at time $t+\Delta$, and $s^-(t+\Delta)$ denotes the absence of a spike at time $t+\Delta$. We group all new components in pairs such that each pair corresponds to a given number of spikes on the interval $t:T$. Each such pair of Gaussian components is then merged into a single Gaussian with equivalent weight, mean, and variance:

$$\begin{aligned} p_{s,k} &\rightarrow w^+ + w^- \\ m_{s,k} &\rightarrow (w^+ m_{s^+,k-1} + w^- m_{s^-,k}) / (w^+ + w^-) \\ v_{s,k} &\rightarrow v_{s,k}^m + (w^+ v_{s^+,k-1} + w^- v_{s^-,k}) / (w^+ + w^-), \end{aligned} \quad (19)$$

and the $v_{s,k}^m$ term corresponds to the variance of the means,

$$v_{s,k}^m = (w^+(m_{s^+,k-1}(t) - m_{s,k}(t - \Delta))^2 + w^-(m_{s^-,k}(t) - m_{s,k}(t - \Delta))^2)/(w^+ + w^-). \quad (20)$$

See Fig. 8 below for an illustration.

2.7 Simulating populations of spiking neurons

To test the performance of different sampling algorithms we simulated a population of $N = 50 - 800$ neurons, following the approach described in (Mishchenko et al., 2010). Briefly, we simulated a spontaneously active randomly connected neural network, with each neuron described by model Eq. (1), and connectivity and functional parameters of individual neurons chosen randomly from distributions based on the experimental data available for cortical networks in the literature (Sayer et al., 1990; Braitenberg and Schuz, 1998; Gomez-Urquijo et al., 2000; Lefort et al., 2009). Networks consisted of 80% excitatory and 20% inhibitory neurons (Braitenberg and Schuz, 1998; Gomez-Urquijo et al., 2000). Neurons were connected to each other in a sparse, spatially homogeneous manner: the probability that any two neurons i and j were connected (i.e., that either w_{ij} or w_{ji} was nonzero) was 0.1 (Braitenberg and Schuz, 1998; Lefort et al., 2009). The scale of the connectivity weights w_{ij} was matched to results from the cortical literature, as cited above, and the overall average firing rate of the networks was set to be about 5 Hz. The connectivity waveforms $w_{ij}(t)$ were modeled as exponential functions with time constant fixed for all neurons at 10 *msec*; for the self-coupling terms $w_{ii}(t)$, neurons strongly inhibited themselves over short time scales (an absolute refractory effect of 2 ms) and weakly inhibited themselves with an exponentially-decaying weight over a timescale of 10 ms. Finally, we used an exponential nonlinearity, $f(\cdot) = \exp(\cdot)$, for simplicity. Again, see (Mishchenko et al., 2010) for full details and further discussion.

3 Results

The efficiency of any Metropolis-Hastings sampler is determined by the ease with which we can sample from the proposal density (and compute the acceptance probability) versus the degree to which the proposal approximates the true target $P(\mathbf{n}_i | \mathbf{n}_{observed})$. In this section we will numerically compare the efficiency of the proposal densities we have discussed above. In particular, we will examine the following proposal densities, listed in rough order of complexity:

- Time-homogeneous Poisson process.
- Point process with log-conditional intensity function given by the delayed input $J_i^-(t)$, Eq. (6).
- Point process with log-rate determined by full effective input in weak-coupling approximation $\tilde{J}_i(t)$, Eq. (10).
- Hybrid truncated HMM proposal including weak-coupling terms (eq. (13)).

As a benchmark, we also compare these samplers against a simple pointwise Gibbs sampler, in which we draw from $P(n_i(t) | \mathbf{n}_{i,\setminus t}, \mathbf{n}_i)$ sequentially over t .

We begin by inspecting a simpler toy model of neural spiking. In this model both refractory and interneuronal coupling effects are assumed to be short, i.e. on the time scale of ≈ 20 msec. This simplification allows us to characterize the neural state fully by $K \approx 10$ past time bins ($\Delta = 2$ ms in these simulations), with the state variable $s_i(t) = \{n_i(t'), t - K\Delta < t' \leq t\}$. In this case, Eq. (1) describes a hidden Markov model with a state space that is small enough to sample from directly, using the forward-backward procedure detailed in section 2.3. Thus, in this case we can obtain the probability distribution $P(\mathbf{n}_i | \mathbf{n}_{observed})$ explicitly, and compare different MH algorithms against a ground truth.

In Fig. 2, we inspect the true instantaneous posterior spiking rate $r_i(t) = P(n_i(t) | \mathbf{n}_{observed}) / \Delta$ for the hidden neuron in this toy model, calculated exactly with the forward-backward procedure and estimated using different effective rates for a few of the MH proposals discussed above. We see that the true rate $r_i(t)$ varies widely around its mean value, implying that the simplest proposal density (the time-homogeneous Poisson process) will result in rather low acceptance rates, as indeed we will see below. Similarly, a naive approximation for $r_i(t)$ using only the delayed input $J_i^-(t)$ fails to capture much of the structure of $r_i(t)$. Incorporating both the past and the future spiking activity of the observed neurons via the weak-coupling input $\tilde{J}_i(t)$ leads to a much more accurate approximation of the true rate $r_i(t)$.

Similar results are obtained when we apply the MH algorithm using these proposal densities (Figures 3 and 4). We use MH with $M = 5000$ samples for each proposal density, with a burn-in period of 1000 samples, to obtain both the autocorrelation functions (Fig. 3) and estimates for the instantaneous spiking rate $r_i(t)$ (Fig. 4). We observe that the homogeneous Poisson proposal density leads to a very long autocorrelation scale, implying inefficient mixing, i.e., long runs are necessary to obtain accurate estimates for quantities of interest such as $r_i(t)$. Indeed, we see in Fig. 4 that 5000 samples are insufficient to accurately reconstruct the desired rate $r_i(t)$. Similarly, for the proposal density based on the delayed inputs $J^-(t)$, the autocorrelation scale is shorter, but still in the range of 10 – 20 samples. The weak-coupling proposal, which incorporates information from both past and future observed spiking activity, mixes quite well, with an autocorrelation scale on the order of a single sample, and leads to an accurate reconstruction of the true rate in Fig. 4. Interestingly, the simplest Gibbs sampling algorithm also performs well, with an autocorrelation length similar to that of the best MH sampler shown here.

We also study the performance of the weak-coupling proposal as a function of the strength of the interneuronal interactions in the network, and as a function of the number of neurons N in the population and the length T of the desired spike train (Fig. 5). The acceptance rate falls at a rate approximately inverse to the coupling strength, which seems sensible, since this proposal is based on the approximation that the coupling strength is weak. We observe similar behavior with respect to the size of the neural population. In particular, the acceptance rate drops below $R \sim 0.1$ when $N \sim 600 - 1000$. On the other hand, increasing the spike train length T affects performance to only a moderate degree: even when the length of the sampled spike train is increased from 10 sec to 160 sec, the acceptance rate remains above $R \sim 0.4$.

One would expect that the performance of the MH algorithm in the case of more strongly-coupled neural networks can be improved by including strong short-term interaction and refractory effects explicitly into the proposal density. We discussed such an approach in section 2.5. For weakly coupled networks, we expect the performance of this hybrid algorithm to be similar to that of our original proposal density; however, for more strongly coupled neural networks we expect a better performance. This expectation is borne out in the simulations shown in Fig. 6: the hybrid sampler (solid black line) performs significantly better than

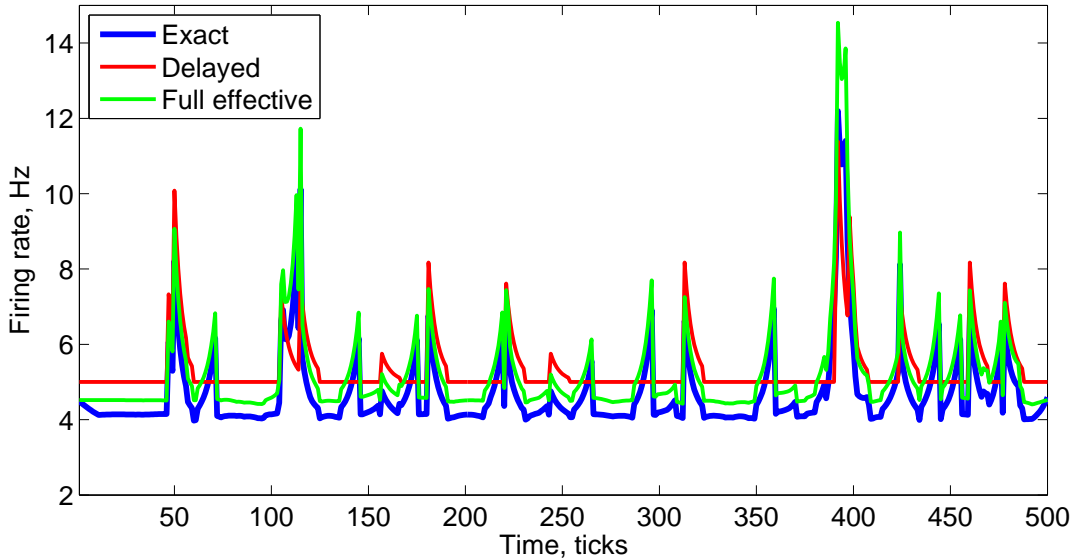


Figure 2: Quantifying the approximation accuracy of two spike train proposal densities. Blue trace: the true spiking rate $r_i(t)$ of one neuron conditioned on the spiking activity of all other neurons in a population of $N = 50$ neurons; the rate is computed exactly via forward-backward procedure described in section 2.3. A population of spontaneously spiking neurons with neurophysiologically feasible parameters (section 2.7) was simulated for a total of 1 sec at a time resolution of $\Delta = 2$ msec. Approximate spiking rates obtained using just the delayed input, $J^-(t)$ (Eq. 6), or full weak-coupling input, $\tilde{J}(t)$ (Eq. 11), are shown in red and green, respectively. (In particular, the delayed input $J_i^-(t)$ and the full effective input $\tilde{J}_i(t)$ are computed as described in section 2.4, and then we approximate the rates using $r_i(t) = f[J_i^-(t)]$ or $f[\tilde{J}_i(t)]$.) While the proposal based on the delayed inputs $J^-(t)$ reproduces the true spiking rate somewhat poorly, the weak-coupling $\tilde{J}(t)$ approximation is significantly more accurate.

the original MH algorithm (dashed black line) and uniformly better than alternative hybrid strategies, e.g., where only the short-scale truncated HMM is used, or where all interneural interactions are accounted for via the weak-coupling approximation, while the self-terms w_{ii} are accounted for via a truncated HMM (gray lines). Similar results are obtained when we vary the size of the neural population N (Fig. 6).

In the simulation settings described above (timestep $\Delta = 2$ msec, with coupling currents w_{ij} lasting up to 50 ms) the MH algorithm (coded in Matlab without any particular algorithmic optimization) took $\approx 300 - 400$ sec to produce 1000 samples of 1000 time-ticks each on a PC laptop (Intel Core Duo 2GHz), dominated by the time necessary to produce 1000 spike train proposals using the forward recursion. The hybrid algorithm had a similar computational complexity ($t_{max} = 10$ ms in the truncated HMM). The Gibbs algorithm had a somewhat higher computational cost, due largely to the fact that updates of the currents $J_i(t)$ for up to 50 ms needed to be performed during each step, to decide whether to flip the state of each spike variable $n_i(t)$. This resulted in running times for the Gibbs sampler which were 5 – 20 times slower than for the MH sampler. The Gibbs updates are even slower in the calcium imaging setting (discussed at more length in the next section), due to the fact

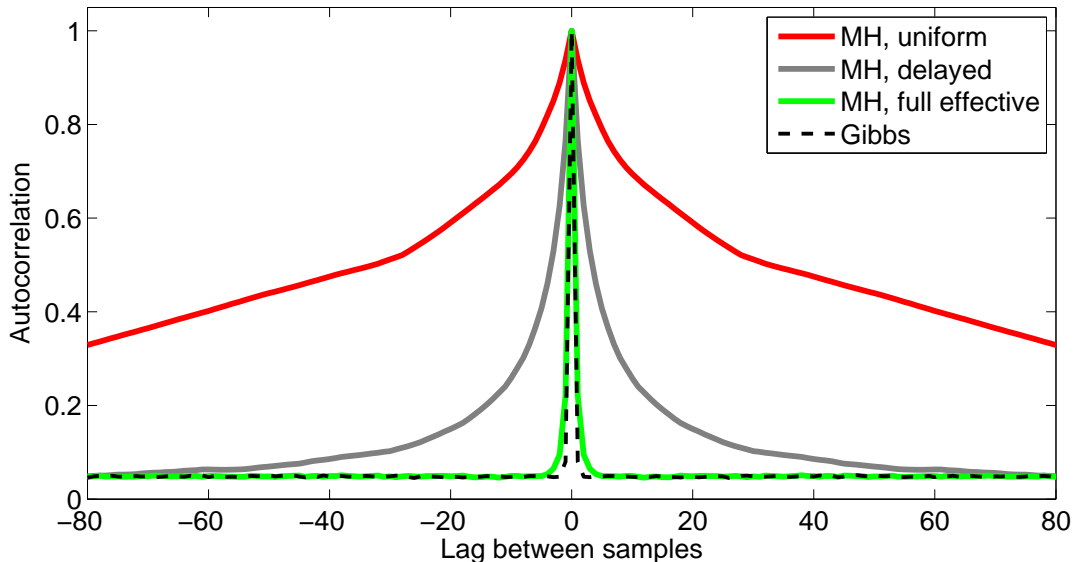


Figure 3: Autocorrelation functions for Gibbs and MH samplers. (The autocorrelation function is averaged over all time bins in the spike train.) Echoing the results of Fig. 2, we see that the homogeneous (“uniform”) Poisson MH algorithm mixes slowly; the sampler based on the full weak-coupling input $\tilde{J}_i(t)$ mixes significantly more quickly than the sampler based on the delayed inputs $J_i^-(t)$. The standard Gibbs sampler performed about as well as the weak-coupling MH sampler. Simulation details are as Fig. 2, except a 10 sec spike train was simulated here to collect a sufficient amount of data to distinguish the different algorithms.

that each update requires us to update $C_i(t)$ over $n_c\tau_i^c$ timesteps (recall section 2.6) for each proposed flip of $n_i(t)$.

3.1 Incorporating calcium fluorescence imaging observations

Next we examine the performance of the method we developed in section 2.6 for sampling from $P(\mathbf{n}_i|\mathbf{n}_{\setminus i}, \mathbf{F}_i)$. We find that this sampler performs quite well in moderate and high SNR settings. More precisely, for values of eSNR greater than ≈ 5 (recall the definition of eSNR in eq. (3)), the conditional distribution $P(\mathbf{n}_i|\mathbf{F}_i)$ is localized near the true spike train quite effectively, leading to a high MH acceptance rate (Figures 7-9). Indeed, we find (as in (Vogelstein et al., 2009)) that it is possible to achieve a sort of “super-resolution” in the sense that the MH algorithm can successfully return $P(\mathbf{n}_i|\mathbf{F}_i)$ on the time-scale of $\Delta = 2$ msec even when fluorescence observations are obtained at a much lower frame-rate (here $FR = 50$ Hz).

Conversely, for low values of eSNR, e.g., $eSNR \approx 1$, the backwards density $P(\mathbf{F}_i(t : T)|n_i(t), C_i(t))$ becomes non-informative (i.e., relatively flat as a function of $n_i(t)$ and $C_i(t)$), and the acceptance rate of our MH algorithm reverts to the rate obtained under the conditions of no calcium imaging data. However, in the low-to-moderate SNR regime (e.g., $eSNR \approx 2$), the performance of the MH sampler can drop substantially. This is primarily due to deviations in the shape of $P(F_i(t)|C_i(t))$ from Gaussian at low SNR. Recall that we made several approximations in computing the backwards density: first, this density is truly a mixture of $\sim 2^{T-t}$ components, whereas we approximate it with a mixture of only $\sim T -$

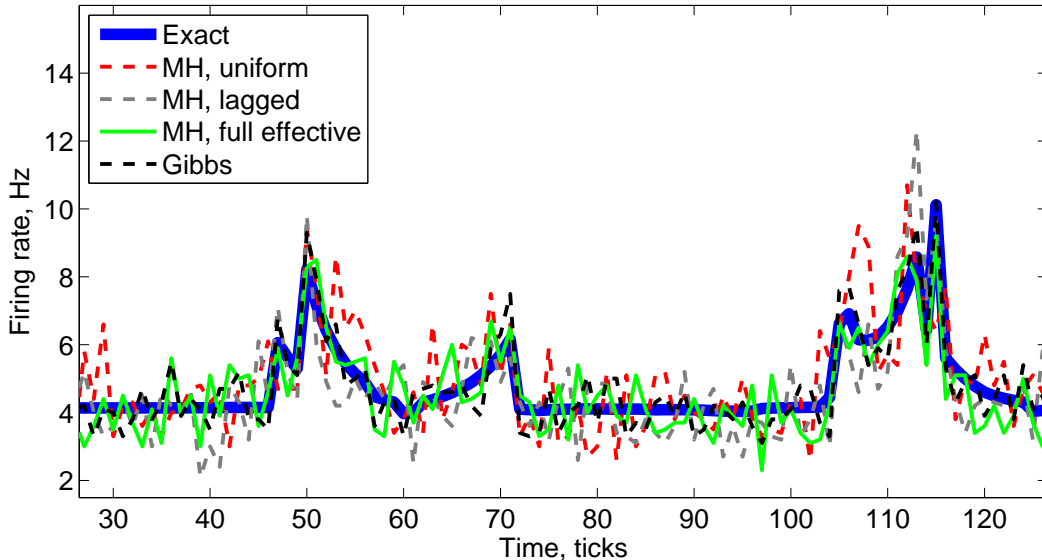


Figure 4: Comparing the accuracy of MH and Gibbs samplers. Each trace indicates the posterior spiking rate $r_i(t)$ for one hidden neuron, estimated from 5000 samples using the Gibbs sampler and the three MH samplers compared in the preceding two figures. The results are similar: the homogeneous Poisson proposal performs badly (in the sense that the $r_i(t)$ computed based on these 5000 samples approximates the true $r_i(t)$ poorly), while the weak-coupling and Gibbs samplers outperform the sampler based on the delayed inputs $J_i^-(t)$ (in terms of variance around the true rate, computed via the full forward-backward HMM method, shown in blue). MH acceptance rates, R , varied from $R \approx 0.75$ for MH using the homogeneous Poisson proposal to $R \approx 0.98$ for MH using the proposal with full effective input $\tilde{J}_i(t)$. Simulation details are as Fig. 2. A shorter time interval of the simulated spike train is shown for greater clarity here.

t components. Second, we assume that each mixture component is Gaussian. Although the fluorescence $F_i(t)$ is described by normal statistics given the calcium variable $C_i(t)$, the relationship between the fluorescence mean and calcium transient can be nonlinear, and the variance may depend on $C_i(t)$ (recall eq. 2). This makes the conditional distribution of $C_i(t)$ non-Gaussian in general, particularly at low-to-moderate levels of SNR (where the likelihood term is informative but not sufficiently sharp to justify a simple Laplace approximation). We tested the impact of each of these approximations individually by constructing a set of toy models where these different approximations were made exact; we found that the non-Gaussianity of $P(C_i(t)|F_i(t))$, due to nonlinear dependence of $F_i(t)$ on $C_i(t)$, was the primary factor responsible for the drop in performance. Thus we expect that in cases where the saturating function $S(\cdot)$ is close to linear, the sampler should perform well across the full SNR range (Fig. 10); in highly nonlinear settings, more sophisticated approximations (based on numerical integration techniques such as expectation propagation (Minka, 2001)) may be necessary, as discussed in the conclusion section below. The MH sampler here took about twice as long as in the fully-observed spike train case discussed in the previous section, largely due to the increased complexity of the backward recursion described in section 2.6.

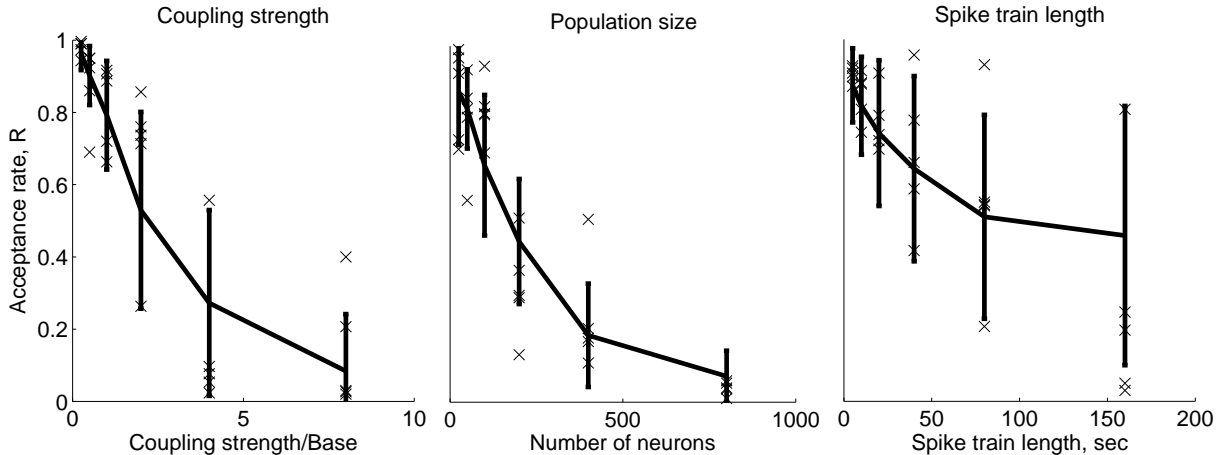


Figure 5: Acceptance rate of the weak-coupling MH algorithm as a function of coupling strength, C , neural population size, N , and spike length, T . We simulated sparsely connected networks of inhibitory and excitatory neurons as described in Section 2.7, with refractory effects up to 10 *msec* long and interneural interactions up to 50 *msec* long. Here a coupling strength value of $C = 1$ corresponds to the neurobiologically motivated set of parameters in section 2.7 (also in Figs. 2-4); other values of C corresponded to scaling $f(J(t)) \rightarrow f(C \cdot J(t))$ in eq. (1). 10 *sec* of neural activity was simulated at a time resolution of $\Delta = 2$ *msec*, with the neural population spiking at $\approx 4 - 5$ Hz. 64 trials of 500 samples were simulated, with a new random neural network generated in each trial, from which we estimated $P(\mathbf{n}_i | \mathbf{n}_{\setminus i}; \mathbf{w})$. Average acceptance rate R and standard deviation are shown for such trials, as well as examples of several individual trials (“x”). MH algorithm performance degraded significantly as C or N increased. On the other hand, for larger values of T performance degraded much less substantially, and even for the largest T we examined (160 *sec*) the acceptance rate remained above $R \sim 0.4$.

Given the good performance of Gibbs sampling noted in the calcium-free setting discussed above, we also examined the Gibbs approach here. However, we found that the Gibbs algorithm was not able to procure the samples successfully given calcium imaging observations. In many cases we found that the Gibbs sampler converged rapidly to a particular spike train close to the truth, but would then become “stuck.” If initialized again, the sampler would often converge to a different spike train, close to the truth, only to become stuck again. This behavior is due to the fact that the conditional distributions $P(n_i(t) | \mathbf{n}_{i, \setminus t}, \mathbf{F}_i)$ can be quite sharply concentrated, leading to poor mixing between spike train configurations with high posterior probability; of course, this is a common problem with the Gibbs sampler (and indeed, this well-understood poor mixing behavior of the standard Gibbs chain is what led us to develop the more involved methods presented here in the first place). Roughly speaking, the extra constraints imposed by the fluorescence observations in $P(n_i(t) | \mathbf{n}_{i, \setminus t}, \mathbf{F}_i)$ relative to $P(n_i(t) | \mathbf{n}_{i, \setminus t})$ make the former distribution more “frustrated,” in physics language, making it harder for the Gibbs sampler to reach nearby states with high posterior probability and leading to the relatively slower Gibbs mixing rate in the calcium-imaging setting.

Finally, we applied the hybrid sampler to a sample of real calcium fluorescence imaging data (Fig. 11), in which $F(t)$ and the true spike times $n(t)$ were recorded simultaneously. Thus we have access to ground truth for $n(t)$ here, though of course only the observed fluorescence

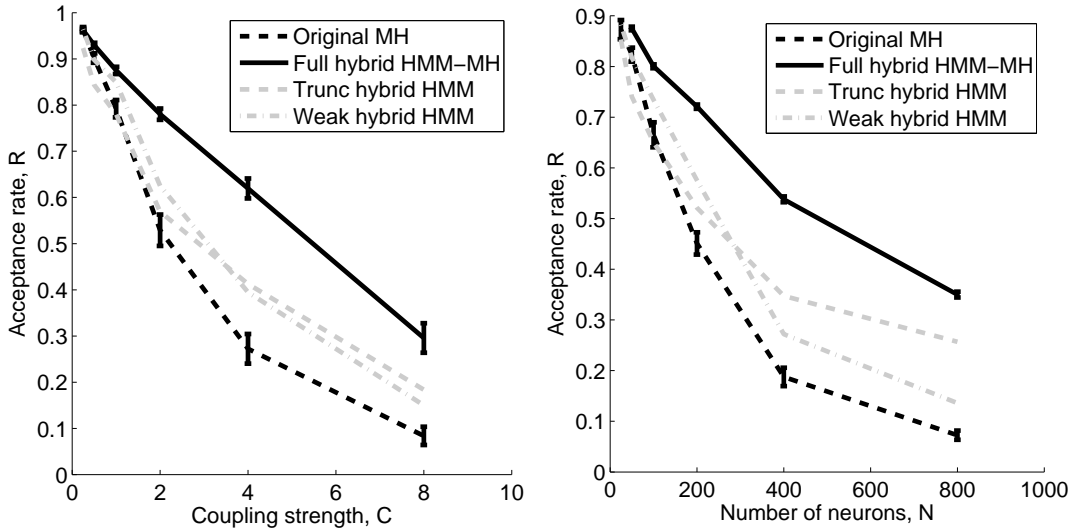


Figure 6: Acceptance rate of hybrid HMM-MH algorithm (Eq. 13) as a function of the overall coupling strength C and neural population size N . The hybrid algorithm (solid black line) performs substantially better than the original MH algorithm described in section 2.4 (dashed black line), and uniformly better than alternative hybrid strategies, e.g., where only a short-scale truncated HMM is used or where all interneural interactions are accounted for via weak-coupling approximation (gray lines). Simulations as in Fig. 5.

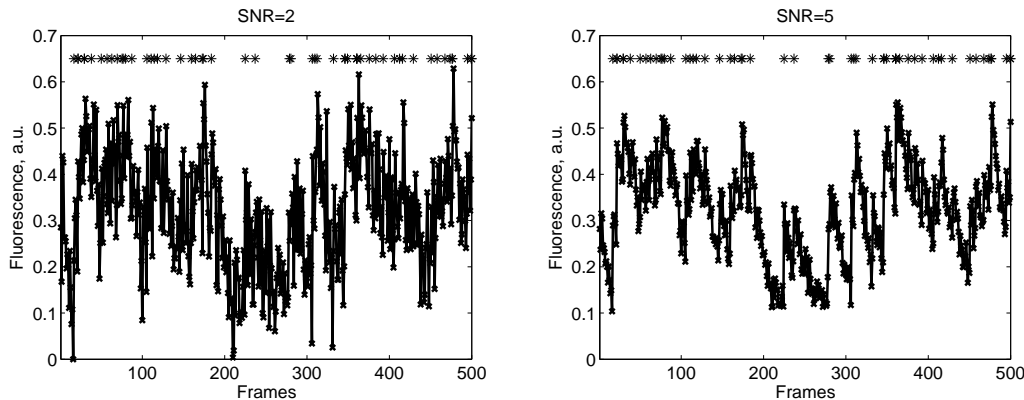


Figure 7: Example of observed fluorescence for low $SNR \approx 2$ and relatively high $SNR \approx 5$. $\Delta = 2$ msec, and calcium imaging frame-rate $FR = 50$ Hz. 10 sec spike trains were simulated, with the neural population spiking at $\approx 4 - 5$ Hz. Actual spikes of the target neuron are indicated with stars.

\mathbf{F} is used to infer $p(n(t)|\mathbf{F})$. The model parameters were estimated using the EM method discussed in (Vogelstein et al., 2009); again, only the observed \mathbf{F} was used to infer the model parameters, not the true spike times $n(t)$. About 50 spikes' worth (5000 fluorescence frames) of data was sufficient to adequately constrain the parameters. Then we applied the hybrid sampler using these parameters to the subset of data shown in Figure 11. The sampler does a good job of recovering the spike rate $n(t)$ from the observed fluorescence data \mathbf{F} , and seems to do a reasonable job of recovering the corresponding jumps in $p(C(t)|\mathbf{F})$ that occur at spike

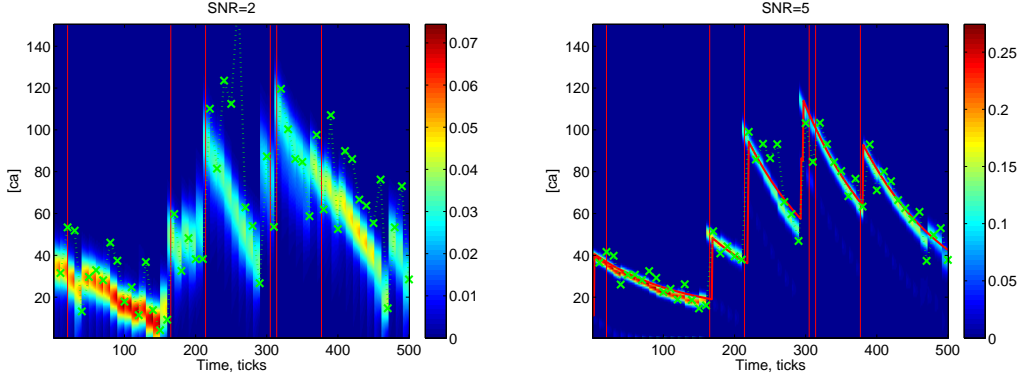


Figure 8: Example of $P(\mathbf{F}_i(t:T)|C_i(t))$ calculated as a T -mixture of Gaussians, as a function of time (x-axis) and $C_i(t)$ (y-axis); colorbar indicates the probability density at time t . For reference, true calcium concentration is shown in red, and “observed” calcium concentration at the imaging frames (i.e. $S^{-1}(F_i(t))$) is shown with green “x”. Simulation details are as in Fig. 7, first second (50 frames) is shown for clarity. Time ticks correspond to $\Delta = 2 \text{ msec}$. Actual spike times of the target neuron are shown with red lines.

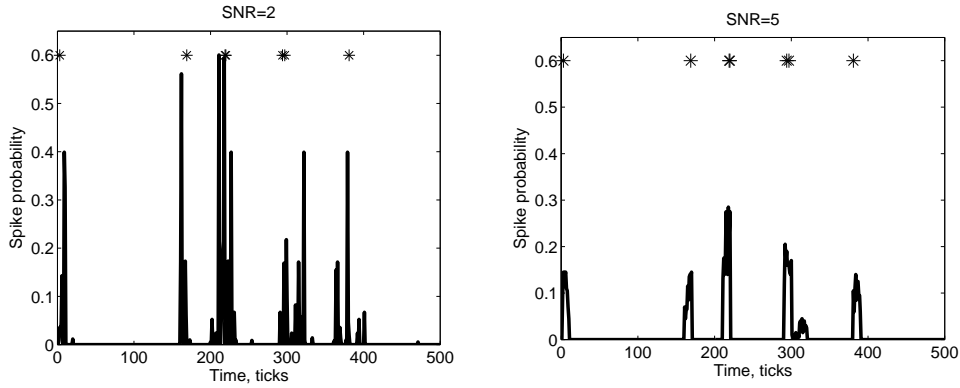


Figure 9: Posterior spike probability, $P(n_i(t)|\mathbf{n}_{\setminus i}, \mathbf{F}_i)$, estimated using MH algorithm. While MH algorithm performs well for $SNR = 5$ and above ($R \approx 0.8$), for low $SNR = 2$ acceptance rate is only $R \approx 0.01$. Simulation details are as in Fig. 7, first second (50 frames) is shown for clarity. Time ticks correspond to $\Delta = 2 \text{ msec}$. Actual spikes of the target neuron are shown with asterisks. Note that in the high-SNR case the sampler successfully recovers the smoothly-varying $P(n_i(t)|\mathbf{n}_{\setminus i}, \mathbf{F}_i)$ on a finer time-scale ($\Delta \approx 2 \text{ msec}$) than the original fluorescence imaging data provided ($\Delta_{FR} \approx 20 \text{ msec}$). However, in the low-SNR case the recovered firing rate is overly spiky and variable due to the slow mixing speed of the M-H chain; recall that similar behavior is visible in Fig. 4.

times. Note that we do not have access to the true intracellular calcium concentration $C(t)$, and therefore no ground truth comparisons are possible for this variable.

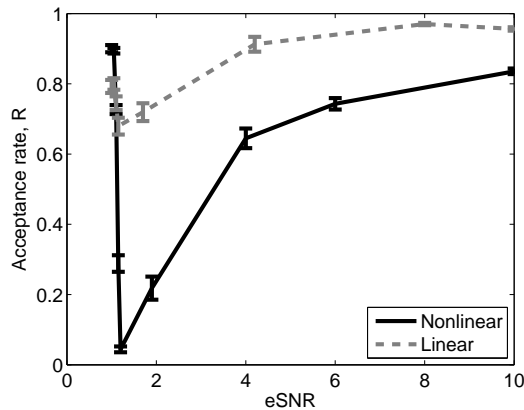


Figure 10: Acceptance rate of MH algorithms as a function of the calcium imaging effective SNR, Eq. 3. Calculations for two calcium signal models, with a toy linear (gray dashed) and realistic Hill (solid black) transfer function, $S(C_i(t))$, are shown. (See (Vogelstein et al., 2009; Mishchenko et al., 2010) for further details on the precise form of the nonlinear function $S(C_i(t))$; in this case, three or four spikes were sufficient to drive the calcium variable into a regime where the fluorescence signal was significantly saturated.) Performance of MH algorithm is good both for high and low eSNR, but can suffer for intermediate $eSNR \approx 1.1 - 3$. This performance drop is primarily due to deviations of the conditional distribution $P(C_i(t)|F_i(t))$ from the Gaussian shape, as exemplified by the much better performance in the model where $S(\cdot)$ is linear and $P(C_i(t)|F_i(t))$ is thus Gaussian. Simulation details are as in Fig. 7. For each eSNR 50 simulations for different neural populations were performed, and the average and standard error of these simulations are shown.

4 Conclusion

In this work we developed several Metropolis-Hastings approaches for sampling from the conditional distribution of neuronal spike trains, given either the activity of other neurons in the network or calcium-sensitive imaging observations. The most effective approach was the hybrid method described in section 2.5, which takes advantage of the fact that strong short-term temporal dependencies within a single spike train may be handled via forward-backward hidden Markov model methods, while weaker long-term dependencies between neurons may be handled with the weak-coupling expansion developed in section 2.4. In each case, to sample efficiently from the spike train at time t it is important to incorporate not only past but also future information (i.e., spiking observations from times both before and after t); (Pillow and Latham, 2007) made a similar point. In the appendix we show that these methods may be extended rather easily to other exponential families (not just the Bernoulli and Poisson cases of most interest in the neuroscience setting); further applications to weakly-coupled Markov chains in non-neural settings seem worth exploring.

Two major avenues are open for future work. First, as noted in Fig. 8, the proposed sampler suffers somewhat in the case of strongly nonlinear fluorescence observations, largely because in this case our mixture-of-Gaussians approximation of the backwards density $P[\mathbf{F}_i(t : T)|C_i(t)]$ can break down. More sophisticated methods for approximating this density are available, and should be explored more thoroughly. Second, as discussed in (Vogelstein et al., 2010), applications of these methods to real data are ongoing, via Monte Carlo-Expectation-

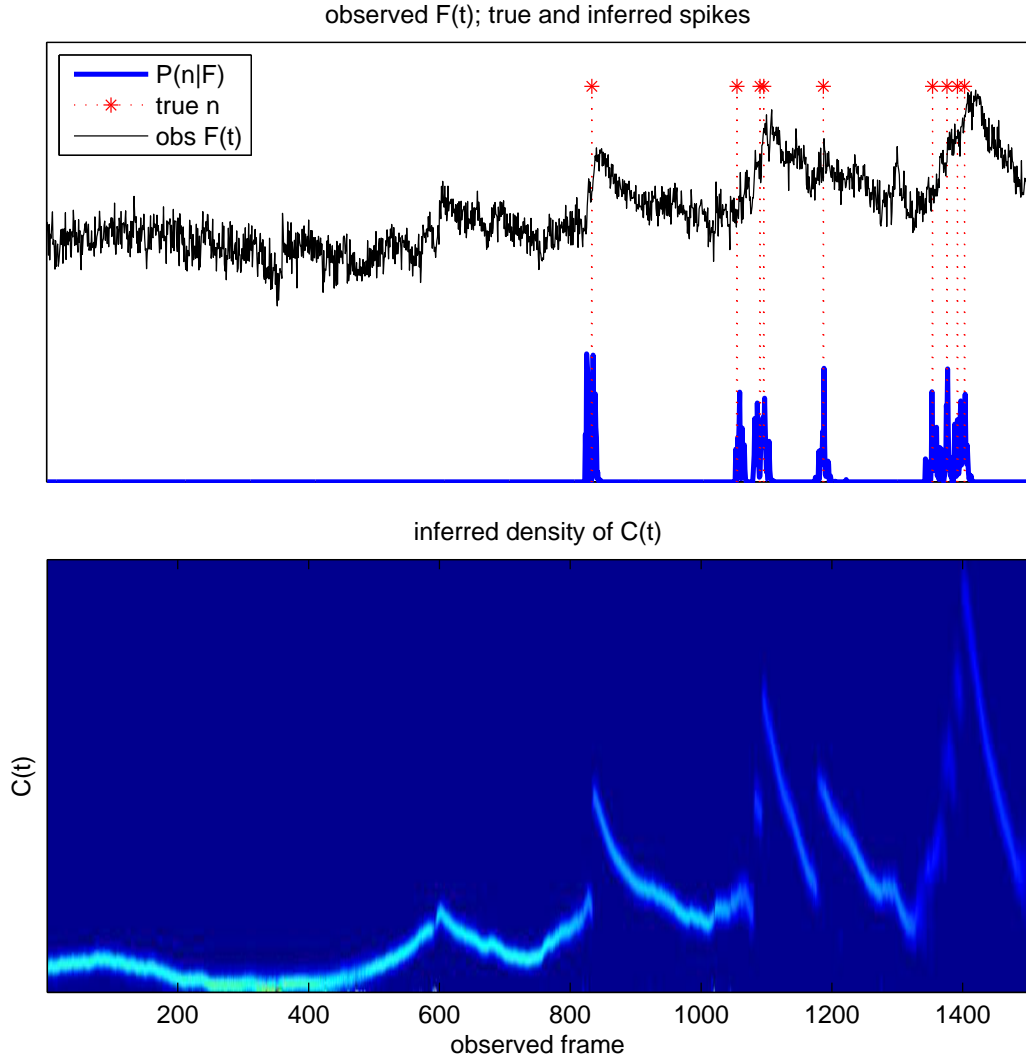


Figure 11: Application to real data (sample data courtesy of T. Sippy, R. Yuste, J. Vogelstein). Top: observed $F(t)$ (black) and true spike times (red) were recorded from a single neuron via simultaneous fluorescence imaging and intracellular patch-clamp electrophysiological recording; see (Vogelstein et al., 2009) for further experimental details. Blue trace indicates posterior $P(n(t)|\mathbf{F})$ computed by the hybrid sampler. Bottom: $p(C(t)|\mathbf{F})$ computed by the hybrid sampler. Y-axis units are arbitrary in this case and have been suppressed. Note that the sampler infers spikes and jumps in $C(t)$ at the correct times.

Maximization methods similar to those discussed in (Vogelstein et al., 2009; Mishchenko et al., 2010), with the fast sampler introduced here replacing the slower Monte Carlo approaches discussed in (Mishchenko et al., 2010). Calcium-fluorescence imaging methods have exploded in popularity over the last several years, and we hope the methods presented here will prove useful in quantifying the cross-correlations and effective connectivity in neural populations observed via fluorescence imaging and multielectrode recording methods.

Appendix: sampling from a weakly-coupled exponential family

As mentioned in section 2.4, it is straightforward to develop a first-order proposal density in the weak-coupling limit more generally, in the case that the variables of interest are drawn from an exponential family distribution with linear sufficient statistics. We begin by writing down our exponential family model, using slightly more compact notation than in sec. 2.4:

$$\log p(\{n_{it}\}) = \sum_{it} f(J_{it})k(n_{it}) + g(J_{it}) + h(n_{it}), \quad (21)$$

with the coupling introduced via

$$J_{it} = b_{it} + \sum_{s>0,j} w_s^{ij} n_{j,t-s}. \quad (22)$$

As before, we may easily expand the joint log-density to the first order:

$$\log p(\{n_{it}\}) = \sum_t f(J_{it})k(n_{it}) + g(J_{it}) + h(n_{it}) \quad (23)$$

$$+ \sum_{t,i \neq j} [f'(b_{jt})k(n_{jt}) + g'(b_{jt})] \sum_{s>0} w_s^{ji} n_{i,t-s} + \text{const}(n_{it}) + o(w). \quad (24)$$

Now if we introduce the assumption that the sufficient statistic is linear, i.e., $k(n) = n$, then after rearranging the double sum over s and t we obtain

$$\log p(\{n_{it}\}) = \sum_t \left[f(J_{it}) + \sum_{s>0,i \neq j} w_s^{ji} [f'(b_{j,t+s})n_{j,t+s} + g'(b_{j,t+s})] \right] n_{it} + h(n_{it}) + \text{const}(n_{it}) + o(w),$$

where again we have suppressed terms (such as $g(J_{it})$) which do not involve n_{it} ; thus, to first order, the conditional distribution of $\{n_{it}\}$ given $\{n_{jt}\}, i \neq j$, remains within the same exponential family, but with a parameter shift

$$f(J_{it}) \rightarrow f(J_{it}) + \sum_{s>0,i \neq j} w_s^{ji} [f'(b_{j,t+s})n_{j,t+s} + g'(b_{j,t+s})]. \quad (25)$$

In the canonical parameterization, $f(J) = J$, standard exponential family theory (Casella and Berger, 2001) shows that $g'(b_{j,t+s}) = -E(n_{j,t+s}|b_{j,t+s})$, and the parameter shift simplifies to

$$J_{it} \rightarrow J_{it} + \sum_{s>0,i \neq j} w_s^{ij} [n_{j,t+s} - E(n_{j,t+s}|b_{j,t+s})]. \quad (26)$$

See (Beck et al., 2007) for a discussion of some related results.

As a concrete example, consider the Gaussian case:

$$\log p(\{n_{it}\}) = - \sum_{it} \frac{1}{2\sigma_i^2} (n_{it} - J_{it})^2 + \text{const.}, \quad (27)$$

with J_{it} as above. In this case we may define $f_i(J_{it}) = J_{it}/\sigma_i^2$ and $g_i(J_{it}) = -J_{it}^2/2\sigma_i^2$; thus we find that the parameter shift in this case is

$$f(J_{it}) = \frac{J_{it}}{\sigma_i^2} \rightarrow \frac{J_{it}}{\sigma_i^2} + \sum_{s>0,i \neq j} \frac{w_s^{ij}}{\sigma_j^2} [n_{j,t+s} - b_{j,t+s}]. \quad (28)$$

In this linear-Gaussian case, we may compute the exact conditional distribution of n_i via the usual Gaussian conditioning formula; the necessary covariance and inverse covariance matrices may be obtained via standard AR model computations. It is straightforward to check that this exact formula agrees with eq. (28) up to $o(w)$ terms in the small- w limit.

Acknowledgments

Thanks to T. Machado, M. Nikitchenko, J. Pillow, and J. Vogelstein for many helpful conversations, and again to T. Sippy and R. Yuste for the data example shown in Fig. 11. This work was supported by a McKnight Scholar and NSF CAREER award, as well as grant IIS-0904353 from the NSF. We also gratefully acknowledge the use of the Hotfoot shared cluster computer at Columbia University.

References

- Abeles, M. (1991). *Corticonics*. Cambridge University Press.
- Beck, J., Ma, W., Latham, P., and Pouget, A. (2007). Probabilistic population codes and the exponential family of distributions. *Prog Brain Res*, 165:509–519.
- Bohland, J. et al (2009). A proposal for a coordinated effort for the determination of brainwide neuroanatomical connectivity in model organisms at a mesoscopic scale. *arXiv*, page 0901.4598.
- Braitenberg, V. and Schuz, A. (1998). *Cortex: statistics and geometry of neuronal connectivity*. Springer, Berlin.
- Briggman, K. L. and Denk, W. (2006). Towards neural circuit reconstruction with volume electron microscopy techniques. *Current Opinions in Neurobiology*, 16:562.
- Brillinger, D. (1988). Maximum likelihood analysis of spike trains of interacting nerve cells. *Biological Cybernetics*, 59:189–200.
- Brillinger, D. (1992). Nerve cell spike train data analysis: a progression of technique. *Journal of the American Statistical Association*, 87:260–271.
- Bureau, I., Shepherd, G. M., and Svoboda, K. (2004). Precise development of functional and anatomical columns in the neocortex. *Neuron*, 42(5):789–801.
- Casella, G. and Berger, R. (2001). *Statistical Inference*. Duxbury Press.
- Chornoboy, E., Schramm, L., and Karr, A. (1988). Maximum likelihood identification of neural point process systems. *Biological Cybernetics*, 59:265–275.
- Cossart, R., Aronov, D., and Yuste, R. (2003). Attractor dynamics of network up states in the neocortex. *Nature*, 423:283–288.
- de Gunst, M., Kunsch, M., and Schouten, J. (2001). Statistical analysis of ion channel data using Hidden Markov Models with correlated state-dependent noise and filtering. *JASA*, 96:805–815.
- Escola, S. and Paninski, L. (2008). Hidden Markov models applied toward the inference of neural states and the improved estimation of linear receptive fields. *In press, Neural Computation*.

- Gomez-Urquijo, S. M., Reblet, C., Bueno-Lopez, J. L., and Gutierrez-Ibarluzea, I. (2000). Gabaergic neurons in the rabbit visual cortex: percentage, distribution and cortical projections. *Brain Res*, 862:171–9.
- Hagmann, P., Cammoun, L., Gigandet, X., Meuli, R., Honey, C., Wedeen, V., and Sporns, O. (2008). Mapping the structural core of human cerebral cortex. *PLoS Biology*, 6(7):e159.
- Hagmann, P., Kuran, M., Gigandet, X., Thiran, P., Wedeen, V., Meuli, R., and Thiran, J.-P. (2007). Mapping human whole-brain structural networks with diffusion mri. *PLoS ONE*, 2(7):e597.
- Harris, K., Csicsvari, J., Hirase, H., Dragoi, G., and Buzsaki, G. (2003). Organization of cell assemblies in the hippocampus. *Nature*, 424:552–556.
- Hatsopoulos, N., Ojakangas, C., Paninski, L., and Donoghue, J. (1998). Information about movement direction obtained by synchronous activity of motor cortical neurons. *PNAS*, 95:15706–15711.
- Helmstaedter, M., Briggman, K. L., and Denk, W. (2009). 3d structural imaging of the brain with photons and electrons. *Current Opinions in Neurobiology*, page Epub.
- Jordan, M. I., editor (1999). *Learning in graphical models*. MIT Press, Cambridge, MA, USA.
- Kulkarni, J. and Paninski, L. (2007). Common-input models for multiple neural spike-train data. *Network: Computation in Neural Systems*, 18:375–407.
- Lefort, S., Tómm, C., Floyd Sarria, J.-C., and Petersen, C. C. H. (2009). The excitatory neuronal network of the c2 barrel column in mouse primary somatosensory cortex. *Neuron*, 61:301–16.
- Luczak, A., Bartho, P., Marguet, S., Buzsaki, G., and Harris, K. (2007). Sequential structure of neocortical spontaneous activity in vivo. *PNAS*, 104:347–352.
- Luo, L., Callaway, E. M., and Svoboda, K. (2008). Genetic dissection of neural circuits. *Neuron*, 57:634–60.
- Minka, T. (2001). *A Family of Algorithms for Approximate Bayesian Inference*. PhD thesis, MIT.
- Mishchenko, Y., Vogelstein, J., and Paninski, L. (2010). A bayesian approach for inferring neuronal connectivity from calcium fluorescent imaging data. *Annals of Applied Statistics*.
- Neal, R., Beal, M., and Roweis, S. (2003). Inferring state sequences for non-linear systems with embedded hidden Markov models. *NIPS*, 16.
- Nykamp, D. (2003). Reconstructing stimulus-driven neural networks from spike times. *NIPS*, 15:309–316.
- Nykamp, D. (2005a). Revealing pairwise coupling in linear-nonlinear networks. *SIAM Journal on Applied Mathematics*, 65:2005–2032.

- Nykamp, D. (2007). A mathematical framework for inferring connectivity in probabilistic neuronal networks. *Mathematical Biosciences*, 205:204–251.
- Nykamp, D. Q. (2005b). Revealing pairwise coupling in linear-nonlinear networks. *SIAM J Applied Mathematics*, 65(6):2005–2032.
- Ohki, K., Chung, S., Ch’ng, Y., Kara, P., and Reid, C. (2005). Functional imaging with cellular resolution reveals precise micro-architecture in visual cortex. *Nature*, 433:597–603.
- Okatan, M., Wilson, M., and Brown, E. (2005). Analyzing functional connectivity using a network likelihood model of ensemble neural spiking activity. *Neural Computation*, 17:1927–1961.
- Paninski, L. (2004). Maximum likelihood estimation of cascade point-process neural encoding models. *Network: Computation in Neural Systems*, 15:243–262.
- Paninski, L., Fellows, M., Shoham, S., Hatsopoulos, N., and Donoghue, J. (2004). Superlinear population encoding of dynamic hand trajectory in primary motor cortex. *J. Neurosci.*, 24:8551–8561.
- Petreaunu, L., Mao, T., Sternson, S., , and Svoboda, K. (2009). The subcellular organization of neocortical excitatory connections. *Nature*, 457:1142–1145.
- Pillow, J. and Latham, P. (2007). Neural characterization in partially observed populations of spiking neurons. *NIPS*.
- Pillow, J., Shlens, J., Paninski, L., Sher, A., Litke, A., Chichilnisky, E., and Simoncelli, E. (2008). Spatiotemporal correlations and visual signaling in a complete neuronal population. *Nature*, 454:995–999.
- Plesser, H. and Gerstner, W. (2000). Noise in integrate-and-fire neurons: From stochastic input to escape rates. *Neural Computation*, 12:367–384.
- Rabiner, L. R. (1989). A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition. *Proceedings of the IEEE*, 72(2):257–286.
- Rigat, F., de Gunst, M., and van Pelt, J. (2006). Bayesian modelling and analysis of spatio-temporal neuronal networks. *Bayesian Analysis*, 1:733–764.
- Santhanam, G., Ryu, S. I., Yu, B. M., Afshar, A., and Shenoy, K. V. (2006). A high-performance brain-computer interface. *Nature*, 442:195–198.
- Sato, T. R., Gray, N. W., Mainen, Z. F., and Svoboda, K. (2007). The functional microarchitecture of the mouse barrel cortex. *PLoS Biol*, 5(7):e189.
- Sayer, R. J., Friedlander, M. J., and Redman, S. J. (1990). The time course and amplitude of epsps evoked at synapses between pairs of CA3/CA1 neurons in the hippocampal slice. *J. Neurosci.*, 10:826–36.
- Shepherd, G. M., Pologruto, T. A., and Svoboda, K. (2003). Circuit analysis of experience-dependent plasticity in the developing rat barrel cortex. *Neuron*, 38(2):277–89.

- Smith, A. and Brown, E. (2003). Estimating a state-space model from point process observations. *Neural Computation*, 15:965–991.
- Smith, S. J. (2007). Circuit reconstruction tools today. *Curr Opin Neurobiol*, 17(5):601–8. Journal Article Review England.
- Song, S., Sjöström, P. J., Reigl, M., Nelson, S., and Chklovskii, D. B. (2005). Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS Biology*, 3:e68.
- Srinivasan, L., Eden, U., Willsky, A., and Brown, E. (2006). A state-space analysis for reconstruction of goal-directed movements using neural signals. *Neural Computation*, 18:2465–2494.
- Stein, R. B., Weber, D. J., Aoyagi, Y., Prochazka, A., Wagenaar, J. B. M., Shoham, S., and Normann, R. A. (2004). Coding of position by simultaneously recorded sensory neurones in the cat dorsal root ganglion. *J Physiol (Lond)*, 560(3):883–896.
- Stevenson, I. H., Rebesco, J. M., Hatsopoulos, N. G., Haga, Z., Miller, L. E., and Kording, K. P. (2009). Bayesian inference of functional connectivity and network structure from spikes. *IEEE Trans. Neural Systems and Rehab.*, 17:203–13.
- Toyoizumi, T., Rahnema Rad, K., and Paninski, L. (2008). Mean-field approximations for coupled populations of generalized linear model spiking neurons with Markov refractoriness. *Neural Computation*, In press.
- Truccolo, W., Eden, U., Fellows, M., Donoghue, J., and Brown, E. (2005). A point process framework for relating neural spiking activity to spiking history, neural ensemble and extrinsic covariate effects. *Journal of Neurophysiology*, 93:1074–1089.
- Tsien, R. Y. (1989). Fluorescent probes of cell signaling. *Ann. Rev. Neurosci.*, 12:227–253.
- Vakorin, V. A., Krakovska, O. A., and McIntosh, A. R. (2009). Confounding effects of indirect connections on causality estimation. *Journal of Neuroscience Methods*, In press.
- Vidne, M., Kulkarni, J., Ahmadian, Y., Pillow, J., Shlens, J., Chichilnisky, E., Simoncelli, E., and Paninski, L. (2009). Inferring functional connectivity in an ensemble of retinal ganglion cells sharing a common input. *COSYNE*.
- Vogelstein, J., Machado, T., Mishchenko, Y., Packer, A., Yuste, R., and Paninski, L. (2010). Methods for neural circuit inference from population calcium imaging data. *Computational and Systems Neuroscience Abstracts*, doi:10.3389/conf.fnins.2010.03.00176.
- Vogelstein, J., Watson, B., Packer, A., Jedynak, B., Yuste, R., and Paninski, L. (2009). Spike inference from calcium imaging using sequential monte carlo methods. *Biophysical Journal*, 97:636–655.
- Yasuda, R., Nimchinsky, E. A., Scheuss, V., Pologruto, T. A., Oertner, T. G., Sabatini, B. L., and Svoboda, K. (2004). Imaging calcium concentration dynamics in small neuronal compartments. *Sci STKE*, 219:p15.
- Yu, B., Afshar, A., Santhanam, G., Ryu, S., Shenoy, K., and Sahani, M. (2006). Extracting dynamical structure embedded in neural activity. *NIPS*.

Yuste, R., Konnerth, A., Masters, B., et al. (2006). *Imaging in Neuroscience and Development, A Laboratory Manual.*