



Bayesian methods for event analysis of intracellular currents



Josh Merel^{a,b,*,1}, Ben Shababo^{c,1}, Alex Naka^c, Hillel Adesnik^{c,d}, Liam Paninski^{a,b,e,f}

^a Neurobiology and Behavior Program, Columbia University, United States

^b Center for Theoretical Neuroscience, Columbia University, United States

^c Helen Wills Neuroscience Institute, University of California, Berkeley, United States

^d Department of Molecular and Cellular Biology, University of California, Berkeley, United States

^e Department of Statistics, Columbia University, United States

^f Grossman Center for the Statistics of Mind, Columbia University, United States

HIGHLIGHTS

- We present a Bayesian approach for automatic event analysis.
- The method was designed and validated for voltage-clamp recordings.
- The method outperforms existing methods on simulated and real data.
- We demonstrate extensions useful for synaptic mapping experiments.

ARTICLE INFO

Article history:

Received 22 March 2016

Received in revised form 13 May 2016

Accepted 16 May 2016

Available online 18 May 2016

Keywords:

Event detection
Postsynaptic current
Calcium imaging
Connectivity mapping
Bayesian methods
MCMC

ABSTRACT

Background: Investigation of neural circuit functioning often requires statistical interpretation of events in subthreshold electrophysiological recordings. This problem is non-trivial because recordings may have moderate levels of structured noise and events may have distinct kinetics. In addition, novel experimental designs that combine optical and electrophysiological methods will depend upon statistical tools that combine multimodal data.

New method: We present a Bayesian approach for inferring the timing, strength, and kinetics of post-synaptic currents (PSCs) from voltage-clamp electrophysiological recordings on a per event basis. The simple generative model for a single voltage-clamp recording flexibly extends to include additional structure to enable experiments designed to probe synaptic connectivity.

Results: We validate the approach on simulated and real data. We also demonstrate that extensions of the basic PSC detection algorithm can handle recordings contaminated with optically evoked currents, and we simulate a scenario in which calcium imaging observations, available for a subset of neurons, can be fused with electrophysiological data to achieve higher temporal resolution.

Comparison with existing methods: We apply this approach to simulated and real ground truth data to demonstrate its higher sensitivity in detecting small signal-to-noise events and its increased robustness to noise compared to standard methods for detecting PSCs.

Conclusions: The new Bayesian event analysis approach for electrophysiological recordings should allow for better estimation of physiological parameters under more variable conditions and help support new experimental designs for circuit mapping.

© 2016 Elsevier B.V. All rights reserved.

1. Introduction

Subthreshold neuronal activity provides an unsurpassed richness of information about a single cell's physiological properties.

Access to subthreshold activity allows for inference about intrinsic biophysical properties (e.g. membrane and ion channel parameters), circuit level properties (e.g. synaptic connectivity), neural coding (e.g. receptive fields), and synaptic properties (e.g. quantal properties and plasticity). At present, whole-cell patch-clamp stands alone in its ability to reliably access subthreshold activity owing to excellent signal-to-noise ratio (SNR) and very high temporal precision, as opposed to optical subthreshold measurements. At the same time, optical technologies have advanced to the point where we can observe the suprathreshold activity of hundreds of

* Corresponding author at: Neurobiology and Behavior Program, Columbia University, United States

E-mail address: jmerel@gmail.com (J. Merel).

¹ These authors contributed equally to this work.

individual neurons simultaneously with calcium imaging and stimulate neurons by subtype or spatial location (Rickgauer et al., 2014). However, the limits on temporal resolution and the indirectness of the observations make inferring fine-scale network and cellular parameters difficult. Approaches which combine optical tools with electrophysiology offer unique advantages (Scanziani and Häusser, 2009). In this work, we present new statistical techniques useful for analyzing whole-cell data as well as extensions demonstrating how our approach is particularly well-suited to settings where electrophysiology is combined with optical physiology.

1.1. Our setting and approach

Fundamentally, many of the subthreshold-based analyses mentioned above depend on the interpretation of the recorded time series as a sequence of *events*. In this setting, events are the successful transmission of neurotransmitter onto the recorded cell, and when this occurs, a transient current flows into or out of the cell, known as a postsynaptic current (PSC). The analyses of experiments designed to infer properties of evoked or spontaneous inputs to a cell (e.g. monosynaptic mapping or quantal/mini-PSC analyses) require determining when a postsynaptic event happened and describing that event. Estimating PSC properties is most straightforward when recordings are acquired using the voltage-clamp configuration which employs a feedback circuit to hold the membrane potential at a constant value thus mitigating variability in PSC properties due to the intrinsic biophysics of the cell (though see Bar-Yehuda and Korngreen, 2008).

In this work, we present a Bayesian approach for inferring the timing, strength, and kinetics of postsynaptic currents from voltage-clamp recordings, and we demonstrate on simulated and real data that this method performs better than standard methods for detecting PSCs. The improvement in single-trial accuracy with our method should allow for better estimation of physiological parameters with less data and under more variable conditions (e.g. when the exact timings of stimuli or its effects are unknown). In addition, the quantification of uncertainty over PSC features provided by Bayesian inference enables new experimental designs (e.g. Shababo et al., 2013).

Bayesian approaches are naturally extensible, so the intuitive, generative model and straightforward inference procedure flexibly extend to include structure relevant to the analyses mentioned above. Specifically, we extend the core single-trial model to include types of data obtained in monosynaptic mapping experiments which may involve optical stimulation artifacts or combine voltage-clamp recordings and optical recordings. For this latter extension, we combine the single trace model presented in this work with related work on calcium imaging (Pnevmatikakis et al., 2013) to demonstrate a Bayesian approach to analyzing mapping experiments consisting of simultaneous population calcium imaging and single cell voltage-clamp recordings (Aaron and Yuste, 2006).

1.2. Review of other approaches

To the best of our knowledge, all previous methods for inferring PSCs have relied on first inferring the timing of single events (i.e. event onsets), and then sometimes fitting per-event kinetics given that event time. These methods have tended to fall into two categories. The superficially simpler of the two approaches is to find events by thresholding the trace or its first derivative (i.e. finite difference). In practice, such methods have extra parameters for smoothing, computing the appropriate offset, or post-processing. Implementations tend to over-detect candidate events and then evaluate candidates based on analysis of per event kinetics (Jonas et al., 1993; Ankri et al., 1994; Hwang and Copenhagen, 1999; Kudoh and Taguchi, 2002). For concreteness, consider a two-stage

approach wherein a threshold is used to identify initial candidates, and then a model is fit to the transient dynamics in order to confirm or reject candidate events by comparison of the parameters of the dynamics against pre-determined criteria (Ankri et al., 1994). Even with post-processing, such methods can be non-selective and tend not to exploit all of the available information (i.e. the transient dynamics are not used to detect the events initially).

Threshold methods have been largely superseded by the second class of approaches, template-based methods (Clements and Bekkers, 1997; Pernía-Andrade et al., 2012). In these methods, templates are usually learned by averaging event-responses collected by a simpler method (e.g. thresholding and/or hand-curation). While template methods are straightforward, initial attempts to apply these methods failed when the amplitude of the events varied or where events overlapped – both common scenarios. The first commonly used algorithm for PSC detection that attempted to avoid issues related to amplitude variability introduced the idea of rescaling a fixed template at each time step (Clements and Bekkers, 1997). Following this trend, template-matching approaches have gradually shifted towards deconvolution methods, which are a more well-founded way to use templates (Pernía-Andrade et al., 2012). Deconvolution generally refers to methods that assume the observed trace is the result of convolving a template with unobserved events (of varying amplitude), and such methods invert this model to estimate the times from the template. Both of these template-based methods produce inferred events with different amplitudes and a threshold can then be used to screen out small events (see Guzman et al. (2014) for a Python implementation of Clements and Bekkers (1997) and Pernía-Andrade et al. (2012), and see Richardson and Silberberg (2008) for deconvolution of current clamp traces).

Methods that rely on fixed-shape templates can work very well when the shape of the event is consistent across events, but postsynaptic events can vary in shape and amplitude, especially for events from different pre-synaptic sources due to different dendritic filtering, issues with space-clamp, or different receptor subunit distributions. Indeed, a core rationale behind the initial preference for threshold based approaches was the recognition that events may vary too much for a single template. While it is possible to use approaches that employ multiple templates (Li et al., 2007; Shi et al., 2010), there are still potential issues related to the stage-wise separation between learning the template and subsequent detection causing a sub-optimal use of information.

We take a Bayesian approach, rooted in a probabilistic, generative model. Broadening the taxonomy, this approach is a type of deconvolution method. However, we do not consider a single template (or a handful of templates), but instead a distribution over templates through the use of prior distributions on the kinetics and amplitudes of individual PSCs. Importantly, we also model event timing in continuous time (i.e. without binning), and we incorporate an autocorrelated, $AR(p)$ noise process (Chib and Greenberg, 1994), which provides a more accurate description of the data. This leads to more precise detection of event times and inference that is more robust (i.e. less susceptible to noise). As such, our inference better leverages all available information (i.e. all events and full timecourse of each event). Given this probabilistic formulation of the noise process and the inclusion of priors on the PSC features, we can then perform posterior inference in this model using Markov chain Monte Carlo (MCMC, see Section 2).

A tradeoff is that the proposed approach is more computationally intensive than previous approaches. Nevertheless, we believe the flexibility and robustness that this approach affords makes up for this in many settings. Beyond handling overlapping events and variation in the shape of events, our method inherits advantages of probabilistic modeling. The method is extensible and amenable to serving as a modular component of hierarchical models, as we

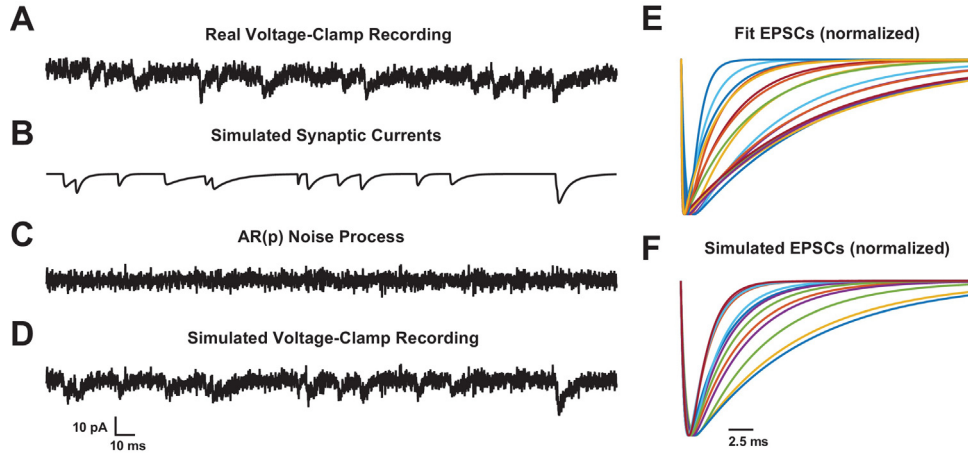


Fig. 1. The correspondence between real data and the generative model. (A) A real voltage-clamp recording. (B) Simulated synaptic currents generated from a fit to the data in (A) (noiseless). (C) AR(p) noise process ($p = 2$) generated from a fit to the data in (A). (D) A model-based voltage-clamp recording simulation (the sum of the data in (B) and (C)), to illustrate that the simulation visually captures core features of the real data. (E) Individual events estimated from the data in (A) (with normalized amplitudes), and panel (F) the individual simulated events used in (B) (with normalized amplitudes).

show. Moreover, while existing methods tend to produce all-or-none results and the precise timing of the event is a secondary consideration, using a probabilistic approach, it is straightforward to consider posterior uncertainty. We essentially get a level of confidence for detection of each event and the level of uncertainty in the precise timing of the event. This posterior uncertainty in event times can translate into posterior uncertainty for other parameters of interest, such as synaptic weights.

1.3. Overview

In the following sections, we will first present the details of the model for single-trial voltage-clamp traces and extensions mentioned previously. We then provide the details of the inference scheme we use for sampling from the posterior distribution. In the results, we compare our approach with the standard template-based approach (Clements and Bekkers, 1997) as well as a Wiener Filtering approach (Wiener, 1949), since these serve as competitive and robust baselines. Note that while we are not aware of the Wiener Filter having been explicitly proposed for this application, the Wiener Filter is a standard deconvolution approach (similar to Pernía-Andrade et al., 2012). Unlike the deconvolution approach in Pernía-Andrade et al. (2012), the Wiener Filter automatically and optimally determines deconvolution parameters from the power spectral density (PSD), and we found this performs better over a range of data than a deconvolution approach with hand-tuned parameters (not shown). We show inference results from our approach on simulated and real data for spontaneous EPSCs and IPSCs across several cell types, PSCs evoked via paired-patching as ground-truth validation, and PSCs evoked optically with one-photon and two-photon stimulation with stimulation artifacts. We also show results on simulated data for a mapping experiment which combines voltage-clamp recordings with calcium imaging, illustrating extensibility.

2. Methods

We draw on tools developed in statistics (Moller and Waagepetersen, 2004) and signal processing (Tan and Goyal, 2008) to decompose a voltage-clamp recording into interpretable elements. In this application, our events are unitary synaptic currents and their features describe the strength and kinetics of each event. In previous work, we have found similar methods useful for

inferring spiking events in calcium imaging data (Pnevmatikakis et al., 2013).

The framework involves (1) specifying a generative model for voltage-clamp recordings, the parameters of which describe event times, features, the noise model, etc., and (2) performing Bayesian inference on event times and features and the model parameters jointly. Theoretically, Bayesian estimators have nice guarantees (under a “true” model, see Lehmann and Casella (1998)). In practice, the generic Bayesian approach can fail if the model is inadequate (e.g. in our case, if the generative model does not capture the true statistics of voltage-clamp traces) or if the inference algorithm performs poorly and the true model posterior is not obtained. Motivated by these legitimate concerns, it is critical to validate that for our application, the model captures the statistics of real voltage-clamp data and that inference performs well on both simulated and real data (see Section 3).

2.1. Model of a single electrophysiological trace

In the simplest version of our model, the observed current trace, y_t , is a discrete time series composed of the sum of n unitary synaptic currents, a baseline (holding current), b , and observation noise ϵ_t (Eq. (1)).

$$y_t = \sum_{i=1}^n a_i f_i(t - t_i) + b + \epsilon_t. \quad (1)$$

$$f_i(t) = (e^{-t/\tau_i^d} - e^{-t/\tau_i^r}) \mathbb{1}(t \geq 0). \quad (2)$$

$$\epsilon_t = \sum_{j=1}^p \phi_j \epsilon_{t-j} + u_t, \quad u_t \sim \mathcal{N}(0, \sigma^2). \quad (3)$$

Each event, indexed by i , is characterized by an event time, $t_i \in \mathbb{R}^+$, which need not be aligned with the sampling time of y_t , its own kinetics determined by $f_i(\cdot)$, and a strength which we define as the amplitude, or peak current, of the event, a_i . For synaptic currents, we use a difference of exponentials for f_i which is parameterized by a rise time constant, τ_i^r , and a decay time constant, τ_i^d . For an example of the model, see Fig. 1B–D. Observe in Fig. 1E that recovered kinetics of individual events do vary significantly, suggesting that a model which captures this structure should perform better than methods which rely on a single (or handful of) template(s).

As opposed to an *i.i.d.* Gaussian noise process, the more general autoregressive, $AR(p)$, process better captures the noise in voltage-clamp recordings. We have found the noise model to be crucial for robust inference (see Section 3). In an $AR(p)$ noise model, the noise has temporal correlations due to direct dependencies between noise values for p timesteps. In this work, we use an AR noise model with $p=2$ (Eq. (3)). Voltage-clamp recordings exhibit correlated noise whose source can be electrical hardware, changes in resistance between the electrode and the interior of the neuron, and other biological non-event contributions to the observation. In practice, it is these forms of temporally correlated noise which lead to many of the false positives since they are more likely to exhibit a similar shape to true events.

With Eqs. (1) and (3), we can write down the likelihood of the observed, noisy data given the parameters ($\Theta \equiv \{\sigma, b, n, \{a_i, t_i, \tau_i^d, \tau_i^r\}_{i=1,\dots,n}, \{\phi_j\}_{j=1,\dots,p}\}$). We begin with the *i.i.d.*, i.e. $AR(0)$, case,

$$p(Y|\Theta) = \prod_{t=1}^T (2\pi\sigma^2)^{-1/2} \exp\left[-\frac{1}{2\sigma^2}(y_t - \hat{y}_t)^2\right], \quad (4)$$

where \hat{y}_t refer to the predicted noiseless trace:

$$\hat{y}_t = \sum_{i=1}^n a_i f_i(t - t_i) + b. \quad (5)$$

In Eq. (8) of Chib and Greenberg (1994), the likelihood is extended to the $AR(p)$ case

$$p(Y|\Theta) = \prod_{t=1}^T (2\pi\sigma^2)^{-1/2} \exp\left[-\frac{1}{2\sigma^2}(y_t - \hat{y}_{t|t-1})^2\right], \quad (6)$$

where $\hat{y}_{t|t-1}$ is (adapted from Eq. (11) of Chib and Greenberg (1994), omitting boundary observations),

$$\hat{y}_{t|t-1} = \hat{y}_t + \sum_{j=1}^p \phi_j (y_{t-j} - \hat{y}_{t-j}). \quad (7)$$

The probabilistic model provides a natural objective function:

$$\mathcal{L}(\Theta|Y) \propto \ln p(Y|\Theta) + \ln p(\Theta). \quad (8)$$

It is possible to optimize this log-posterior directly, or inference can be performed to obtain an estimate of the posterior distribution. $p(\Theta)$ corresponds to the prior probability on the parameters. In a probabilistic formulation, it is worth explicitly keeping in mind that the posterior distribution for a given parameter can only have support where its prior distribution has support, so hard constraints (e.g. a parameter being positive or a minimum amplitude size) are naturally incorporated as prior information. For amplitudes, baselines, time-constants, and event times, we use non-informative, improper uniform priors over either real or bounded real numbers (we selected very broad ranges appropriate for each parameter). The prior on the event count is given by a Poisson distribution which has one free parameter corresponding to the prior expected number of events – this is a parameter that a user would tune depending on their general expectation about the number of events in their data. The prior on the noise level σ^2 is a diffuse inverse gamma distribution (which is the conjugate prior), and we use a diffuse prior distribution over ϕ (i.e. a normal distribution subject to stability constraints, same as Chib and Greenberg (1994)) – see Section 2.3 on inference for more details.

Additionally, we note that more sophisticated priors in a Bayesian model serve as natural ways to extend the core model. For example, distributions over event features could be modulated by clustering onto presynaptic sources or the rate of events in each trace could be time-varied based on stimulation of presynaptic cells

in mapping experiments. Model extensions are explored in the next section.

2.2. Model extensions

2.2.1. Including optical currents

In this extension, we consider an active mapping experiment where we have some level of spatially structured optical stimulation (via optogenetics Fennó et al., 2011 or neurotransmitter uncaging Callaway and Yuste, 2002) of presynaptic cells while holding a postsynaptic cell in voltage-clamp (Shepherd et al., 2005; Katzel et al., 2011). In this setting, detections of PSCs coincident with stimulations can be used to infer connectivity between neurons. However, in some protocols, the postsynaptic cell also responds to the stimulation.² If this is the case, we will see a direct optically evoked current in our voltage-clamp recording when we attempt to stimulate cells near the patched cell. In order to remove artifacts of this sort, we can incorporate a parameterized, additive term in the generative model and then perform inference jointly with respect to these parameters. The choice of an additive term is more reasonable for optogenetic stimulation because the currents are carried through different channels whereas with neurotransmitter uncaging the direct stimulation current and the synaptic current may be competing for the same channels.

For example, it is straightforward to include a parameterized kernel for the optical response of the neuron, $h(\cdot)$, and then to convolve that response with the known optical input to the neuron (e.g. the laser or LED power),

$$y_t = \sum_{i=1}^n a_i f_i(t - t_i) + \sum_j^{n_s} a_j h(\theta_j) * d_j(t) + b + \epsilon_t, \quad (9)$$

where we have n_s optical inputs each with known timecourse $d_j(t)$ and the response kinetics are modeled with a convolution which is parameterized by θ_j (e.g. a set of time constants). Each stimulation will have its own gain, a_s , which depends on the density of the corresponding channels at the location of that stimulation. We have found this approach to be useful under certain conditions with optogenetics (Fig. 5D). However, for some optogenetic currents, the multi-state kinetics of opsins can make it difficult to design a parameterized $h(\cdot)$ which can be sampled efficiently. In this case, the shape of the optical current could be measured empirically and then modeled as

$$y_t = \sum_{i=1}^n a_i f_i(t - t_i) + \sum_{j=1}^{n_s} a_j h(t - t_j^{(s)}) + b + \epsilon_t, \quad (10)$$

where $h(\cdot)$ is now the stereotyped shape of the current and we know the set of times $\{t_j^{(s)}\}_{j=1,\dots,n_s}$ at which we have stimulated. Since the currents are filtered in the dendrites, Eq. (10) breaks down slightly when handling stimulations at locations at varying distances from the soma of the patched cell (Fig. 5D). Nonetheless, it still provides a good trade off between accuracy and computational tractability. In addition, one could create several optical current templates based on the distance of the stimulation site from the soma (Mena et al., 2015).

² This is seen with glutamate uncaging at locations near the patched cell and could also occur with optogenetics when mapping connections between cells of the same transcriptional identity in the same location or when mapping many heterogeneous populations of cells (i.e. when a pan-neuronal promoter is used).

2.2.2. Mapping with calcium imaging and voltage-clamp recordings

Next we extend the model to show how the inferred events could be identified with presynaptic sources when the local population is partially observed via calcium imaging (Aaron and Yuste, 2006) (alternatively, a similar approach could combine information from voltage imaging). In this setting, our observations will consist of the fluorescence traces of the imaged cells and the voltage-clamp recording. Both the presynaptic fluorescence traces, c_t^k where k indexes the neurons observed via imaging, and the postsynaptic electrophysiological recording, y_t , can be interpreted as a sum of events (Pnevmatikakis et al., 2013),

$$c_t^k = \sum_{i=1}^{n_k} a_k^{(c)} g_k(t - t_{ki}) + b_k + \eta_t^k, \quad (11)$$

$$y_t = \sum_{k=1}^K \sum_{i=1}^{n_k} a_k^{(y)} f_k(t - t_{ki}) + \sum_{j=1}^{n_\theta} a_j^{(y)} f_j(t - t_j) + b + \epsilon_t, \quad (12)$$

with ϵ_t as before and $\eta_t^k \sim \mathcal{N}(0, \sigma_k^2)$. Like $f_k(\cdot)$, $g_k(\cdot)$ is also a sum of exponentials but the timescale of the kinetics is much slower. $\{t_j\}_{j=1, \dots, n_\theta}$ are the times of PSCs with no corresponding imaged, presynaptic cell. We want to point out that because we model the process in continuous time, the observed calcium traces need not have the same sampling rate or observation times as each other or the voltage-clamp recording (which will be sampled many orders of magnitude faster).

In this demonstration, we have chosen to assume the calcium events for the same cell are all of equal amplitude – this would be appropriate if the spikes occur sparsely or if calcium transients sum roughly linearly. If the summation is known to be nonlinear with some biophysically plausible nonlinearity (e.g. see Vogelstein et al., 2009), this could be straightforwardly accommodated by modifying the model and similar inference methods may still be applied. Alternatively, it is also straightforward for the calcium event amplitudes to vary across events.

We have similarly chosen for all postsynaptic events to be of the same amplitude and shape when they follow from a specific presynaptic cell, but in practice one can extend the clustering to impose cell-specific priors on these amplitudes and shapes. Some subset of events observed in the electrophysiological recording will arise from unobserved presynaptic inputs so we allow these to be explained by events with no observed presynaptic cause and with independent shape and amplitude per event (second summation term indexed $i = 1, \dots, n_\theta$). We could also incorporate a fixed delay between presynaptic events and postsynaptic events (this would be an additional parameter, pre-specified or inferred, in the postsynaptic transient response $f_k(\cdot)$).

Since we now observe multiple traces, our new objective function is simply the sum of the log probabilities of the various traces. The multiple traces are conditionally independent given event times $\{t_{ki}\}_{k=1, \dots, K, i=1, \dots, n_k}$ and $\{t_j\}_{j=1, \dots, n_\theta}$, so we have an overall objective corresponding to the log-posterior:

$$\mathcal{L}(\Theta | C^k, Y) \propto \ln p(Y | \Theta) + \sum_{k=1}^K \ln p(C^k | \Theta) + \ln p(\Theta), \quad (13)$$

again with $p(\Theta)$ corresponding to the prior probability on the full set of parameters (see Section 3.4 for results).

2.3. Inference

For this work, we perform inference using Markov chain Monte Carlo (MCMC) (Neal, 1993; Gelman et al., 2014). MCMC techniques allow us to obtain samples from the posterior distribution over

all unknown variables in the model and thereby approximate the posterior by a histogram of such samples. Specifically, we perform Gibbs sampling over all the parameters (Gelman et al., 2014). This means that for each parameter, we hold all other parameters fixed and conditionally update the focal parameter by sampling it from its conditional distribution. A “sweep” consists of an update of all parameters. While it is possible to compute conditional distributions analytically for sufficiently simple models, it is quite simple to use generic sampling methods to update parameters for any model for which one can compute the likelihood. For example, we use random-walk Metropolis (RWM) to update many parameters – this consists of updating parameters by proposing updates from a distribution centered on the current value and accepting or rejecting proposed updates such that the resulting set of samples are consistent with the conditional distribution (Gelman et al., 2014). Alternatively, more powerful samplers such as Hamiltonian Monte Carlo could be used (Gelman et al., 2014) but simple RWM sufficed here.

In the single-trial, voltage-clamp case, we update $\{t_i, a_i, b, \tau_i^d, \tau_i^f\}_{i=1, \dots, n}$ by RWM. Inclusion of a direct optical current, in the simplest case, only contributes one additional amplitude parameter for each stimulation, and these can be inferred similarly. Through sampling, we also estimate the posterior distribution of the number of events in each trace, given a Poisson prior (note that this could be generalized to an inhomogeneous Poisson process prior so that the event rate can vary based on inputs such as optical stimulation). To add and remove events (i.e. perform inference over n), we use birth-death moves which consist of the proposal of a new event time and the removal of an existing event time respectively (Moller and Waagepetersen, 2004).

The noise process parameters $\phi_{1, \dots, p}$ are sampled by rejection sampling from the constrained conditional distribution and σ is sampled from its conditional distribution. Specifically, we reproduce the updates for the $\phi_{1, \dots, p}$ and σ , which are provided in Section 4.1 of Chib and Greenberg (1994). $\phi_{1, \dots, p}$ is shown to be conditionally normal with a mean and posterior that depend on $\hat{e}_t = y_t - \hat{y}_t$ and σ^2 , but also with the constraint that the AR(p) process defined by $\phi_{1, \dots, p}$ is stable (note that stability here means that the AR process remains bounded on bounded intervals if run autonomously, which can be quickly checked by examining the magnitude of the roots of the associated characteristic polynomial Lay, 2012). Omitting boundary observations:

$$E = [\hat{e}_{t-1}, \dots, \hat{e}_{t-p}], \quad (14)$$

$$\Phi_n = \Phi_0 + \sigma^{-2}(E'E), \quad (15)$$

$$\hat{\phi} = \Phi_n^{-1}(\Phi_0\phi_0 + \sigma^{-2}E'e), \quad (16)$$

$$\phi \sim \mathcal{N}(\hat{\phi}, \Phi_n^{-1}) \mathbb{1}_{S\phi}, \quad (17)$$

where $\mathbb{1}_{S\phi}$ is an indicator over the set of stable ϕ values and ϕ_0 and Φ_0 are set to weakly informative prior values. Following Chib and Greenberg (1994), we place an inverse-gamma prior on σ^2 , which is the conjugate prior given the likelihood (Eq. (6)) so that we can sample directly from the conditional posterior,

$$\delta_1 = \sum_{t=1}^T (y_t - \hat{y}_{t|t-1})^2, \quad (18)$$

$$\sigma^2 \sim \mathcal{IG}\left(\frac{1}{2}(T + \nu_0), \frac{1}{(1/2)(\delta_1 + \delta_0)}\right), \quad (19)$$

where ν_0 and δ_0 are also set to weakly informative prior values.

For the full mapping case, inference is similar. However, additional moves need to be incorporated to improve mixing. That is, in addition to the single-variable updates, additional custom moves are performed each sweep. The additional moves correspond to

proposing a swap for the presynaptic identity for a PSC event. That is, an MCMC step is proposed wherein an event associated with one presynaptic source is eliminated and an event at the same time is considered for another source. To implement such moves, we simply propose to drop an event at a given time for one presynaptic source and to add an event at the same time for another presynaptic source and evaluate the combined acceptance or rejection of these proposed moves using the Metropolis-Hastings ratio (Gelman et al., 2014).

2.4. Specific implementation details

Code implementing our inference routine is available on github.³ When applying our inference method to data, we can run from a cold start, or we can initialize event times with those found via a simpler, faster method (e.g. Clements and Bekkers, 1997 or Pernía-Andrade et al., 2012). In this work, to establish initial performance of our algorithm, we present results on simulated and real data using cold starts. MCMC methods ideally will be run long enough that a proper posterior distribution can be estimated. A standard convergence check for MCMC methods is to run multiple chains and assess that they provide similar estimates. We did not run multiple Markov chains in this case for all analyses. Rather, for ease of use, we followed a simple heuristic of initially tuning the total number of sweeps by running multiple Markov chains and then consistently performing a conservatively large number of sweeps to robustly achieve good performance (i.e. as measured relative to other methods).

For applications, it is pragmatic to adjust the total number of samples to trade off reliability (i.e. high probability of convergence) against computation time. For small timeseries (roughly 1 s), inference consisting of 2000 sweeps of the sampler tends to provide an accurate posterior (see Section 3, simulated data) and requires anywhere from a few seconds to a few minutes on a contemporary computer depending on the number of inferred events. It is standard to discard the first fraction of the sweeps as *burn-in* (we discard between one-fifth to two-fifths of the total number of sweeps, depending on the event rate and its prior). For longer timeseries or many timeseries, we run smaller sections of the timeseries in parallel on a computing cluster (on a single CPU, the computation time would scale linearly with the number of requested samples and scale linearly in the number of inferred events, which is assumed to be proportional to the length of the trace). We also note that RWM requires a proposal width hyperparameter – this can be automatically tuned early in the sampling process by adjusting the proposal variance such that the accept rate is reasonable (Rosenthal, 2011). We perform this automatic tuning in an *ad hoc* adaptive fashion, incrementally increasing or decreasing the proposal variance when either too many or too few moves are being accepted.

We also note that in the implementation, we frequently rely on the log-likelihood, $\ln(p(Y|\Theta))$, to determine whether to accept or reject moves. In order to minimize redundant operations, it proves useful to store $\hat{e}_t = y_t - \hat{y}_t$ and we perform evaluations of $y_t - \hat{y}_{t-1}$ by observing that:

$$y_t - \hat{y}_{t-1} = y_t - \left(\hat{y}_t + \sum_{j=1}^p \phi_j (y_{t-j} - \hat{y}_{t-j}) \right) \quad (20)$$

$$= y_t - \hat{y}_t - \sum_{j=1}^p \phi_j (y_{t-j} - \hat{y}_{t-j}) \quad (21)$$

$$= \hat{e}_t - \sum_{j=1}^p \phi_j \hat{e}_{t-j}. \quad (22)$$

Finally, note that key hyperparameter settings are presented in [Supplementary Table 1](#).

2.5. Experimental methods

All experiments were performed in accordance with the guidelines and regulations of the Animal Care and Use Committee of the University of California, Berkeley. Mice used for experiments in this paper were either wild type (ICR white strain, Charles River), som-IRES-Cre (JAX stock #018973); Ai9 Rosa-LSL-tdTomato (JAX stock #007909), PV-Cre (JAX stock #008069); Ai9 Rosa-LSL-tdTomato, or emx1-IRES-Cre (JAX stock #005628).

Viral infection: Neonatal emx1-cre mice were injected with AAV9-CAG-flexed-Chr2-tdTomato (for 1-photon experiments) or AAV9-syn-ChrimsonR-tdTomato (for 2-photon experiments) at P0–P4. Viruses were acquired from the University of Pennsylvania Vector Core. Undiluted viral aliquots were loaded into a Drummond Nanoject injector. Neonates were briefly cryo-anesthetized and placed in a head mold. With respect to the lambda suture coordinates for S1 were: 2.0 mm AP; 3.0 mm L; 0.3 mm DV.

2-Photon optogenetic stimulation: 1040 nm light (femtoTrain, Spectra-Physics) was delivered to the sample using a VIVO 2-Photon workstation (3i) based on a Sutter Moveable Objective Microscope (Sutter, Novato, CA) and the hologram was created using a Phasor 2-Photon computer-generated holography system (3i). Light was delivered for 10 ms at 100 mW power on sample. The hologram for these data was a disc of radius 15 μ m.

For more detailed methods on brain slicing, *in vitro* and *in vivo* electrophysiology, and 1-photon optogenetic stimulation, see [Pluta et al. \(2015\)](#). All electrophysiology was analyzed at 20 kHz, or equivalently, with timebins of 0.05 ms.

3. Results

3.1. AR noise model validation

We have found the choice of noise model to be critical when analyzing voltage-clamp data. It is common to use *i.i.d* Gaussian noise for neurophysiological time-series ([Paninski et al., 2012](#); [Richardson and Silberberg, 2008](#)), but very often the noise can exhibit temporal correlations. To obtain a recording of a voltage-clamp noise process which contains no events, we recorded from a neuron exposed to an excitatory synaptic blocker (Kynurenic Acid, 4 mM) while holding the cell near the inhibitory reversal potential (−70 mV). Under these conditions, the recording should be relatively event free; nonetheless the recording shows clear temporal correlations ([Fig. 2A](#)). In the context of deconvolving these data, it is primarily this correlated noise that drives false positives because it can have similar features to PSCs.

To better capture the structure of voltage-clamp noise, we used the more general AR process which we found was sufficiently flexible and expressive to represent the types of noise we encountered. Specifically, an AR(2) model balanced model expressiveness and computational cost. [Fig. 2A](#) and [B](#) shows that the extra structure in the AR(2) model does in fact provide a better description of the data.

For a systematic validation, we performed a comparison between the AR(0) and AR(2) inference for many simulated traces (10 × 1 s traces per noise level). Specifically, we can simulate traces with random event times and with various levels of AR(2) noise added to the traces (i.e., varying SNR). In these simulations, events are naturalistic in that they have variability in their distribution of

³ <https://github.com/jsmerel/joint-calcium-ephys-mapping>.

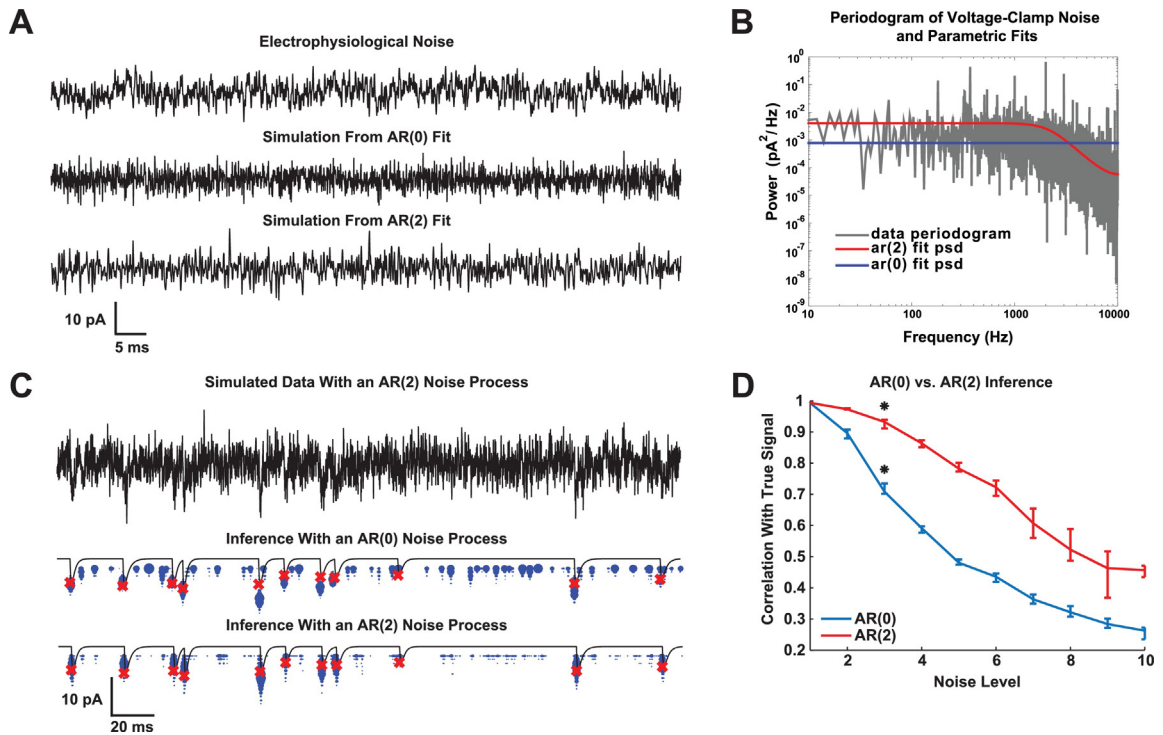


Fig. 2. Validation of the $AR(p)$ noise model and inference. (A) Top: a real voltage-clamp recording under “event-free” conditions such that the recording is dominated by noise. Middle: simulated noise from an $AR(0)$ fit to the real example. Bottom: simulated noise from an $AR(2)$ fit to the real example. The $AR(2)$ example better captures the structure of real noise. (B) Gray: periodogram estimate of the power spectral density (PSD) of an “event-free” recording. Blue: a parametric fit to the PSD with an $AR(0)$ model. Red: a parametric fit to the PSD with an $AR(2)$ model. (C) Top: a simulated voltage-clamp recording with $AR(2)$ noise. Middle, black: the true current for the top trace. Red: the true time-amplitude coordinate for each PSC. Blue: a bivariate histogram reflecting the estimated time-amplitude posterior distribution using an $AR(0)$ noise model (uses a small, but non-zero minimum event size threshold). Bottom: same as in the middle trace but inference is performed with an $AR(2)$ noise model. (D) Curves depict accuracy of inference as a function of SNR level (ranging an order of magnitude, with 1 indicating low noise relative to size of events and 10 indicating noise that has marginal variance larger than the signal) for $AR(0)$ vs $AR(2)$ model-based inference. The measure of accuracy is correlation coefficient between true (simulated) trace and estimate of posterior mean. The asterisk indicates a biologically realistic SNR level equal to the example in (C). For each point on the curve, we simulate timeseries with random event-times and amplitudes (traces are median and inter-quartile range over 10 repeats). Both algorithms do very well in the high SNR regime. As soon as noise level begins to make inference difficult, both algorithms begin to lose accuracy. However the $AR(2)$ model inference degrades much more slowly. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

amplitudes and time constants, and events may overlap (2C). For inference with either noise model, performance is similar when there is low levels of noise (or for specific $AR(2)$ noise process parameters that result in only weak noise autocorrelation, not shown), but the inference results diverge dramatically when the noise is larger in magnitude. To summarize inference, we examine the correlation between the posterior mean trace and the “true” simulated, noiseless event trace (2D). For biologically realistic $AR(2)$ noise structure and magnitude (indicated by an asterisk in 2D), inference with the $AR(2)$ model indeed performs considerably better than with the $AR(0)$ model. Note that in the high-SNR limit, simple methods like template matching algorithms, greedy optimization, or other direct optimization of the model likelihood can perform well enough, so sample-based inference would be computationally excessive. However, this validation demonstrates that in the biologically realistic noise regime, noise is sufficiently large and structured for proper inference to be useful.

3.2. Comparison with other methods on spontaneous and evoked PSCs

We also compared our method to two common PSC detection algorithms: a standard template-based approach (Clements and Bekkers, 1997) and a deconvolution approach (a Wiener Filter Wiener, 1949, which we found tended to improve upon the slightly simpler inverse filter used in Pernía-Andrade et al. (2012)). Although Clements and Bekkers (1997) is nearly twenty years

old and is often outperformed by deconvolution approaches, it is implemented in two commonly used software packages for the analysis of electrophysiological data: Axograph and pClamp. Therefore it is still regularly used when PSC detection is performed (Mardinly et al., 2016; Medelin et al., 2016). When testing these algorithms, we attempted to give each its best chance to perform well. Specifically, that means that the template-based and deconvolution methods were given the average of the true underlying events as a template. We gave the Wiener Filter the true noise power spectral density for simulated data. For real data with high enough spontaneous rates, it was difficult to find a “quiet” section of the trace to estimate the noise power spectral density (PSD); therefore we provided an AR fit to the noise from Bayesian inference for the Wiener Filter’s noise PSD. For simulated data with Bayesian inference, we provided the true priors on τ^r and τ^d that were used to generate the data.

First, we simulated a test set of recordings with realistic levels of noise and with relatively low SNR PSCs. Each of the three algorithms produces a time series equal in length to the input recording that is something like a score that an event is happening at that point in time. For the template-based method the output represents the goodness-of-fit to the event template at that point in time, and for the deconvolution the output is an estimate of the event amplitude at that point in time. For the Bayesian approach the output time series is the marginal posterior of an event at each sample. Importantly, in practice the output for the template-matching and deconvolution methods only have an empirically reasonable interpretation in the high SNR case. To obtain estimates of event times

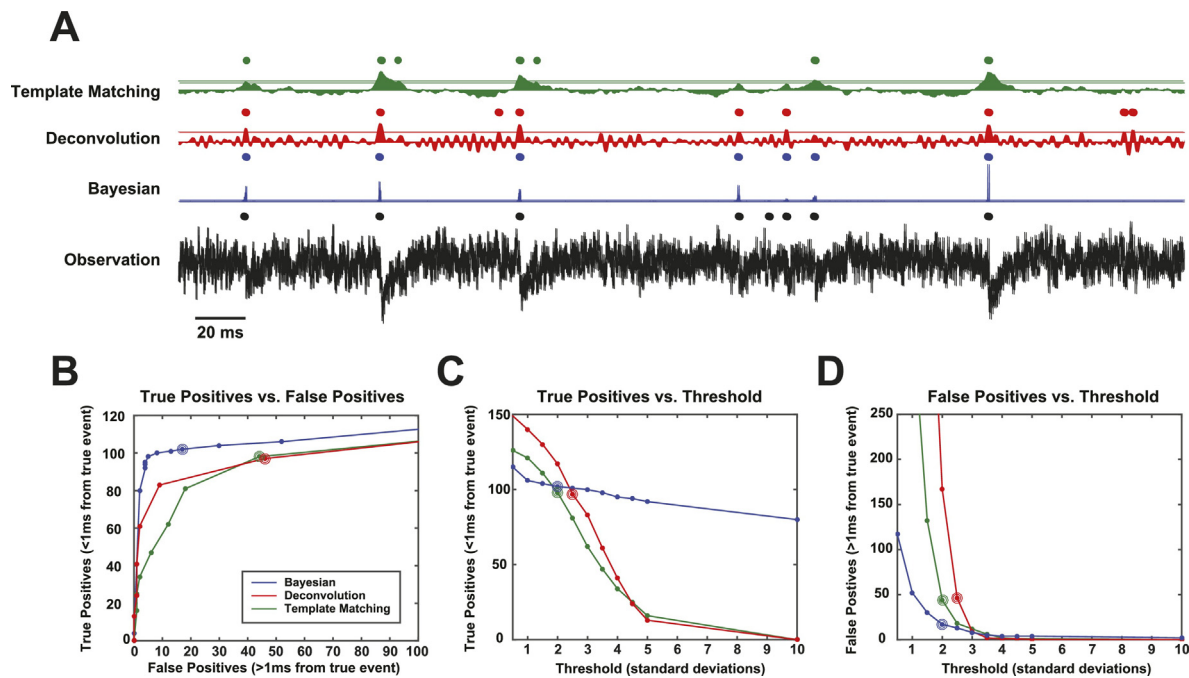


Fig. 3. Comparison of methods on simulated data. (A) An example of inference/detection results for each method on a simulated voltage-clamp recording. For each method we show the output time-series for that method (i.e. the score template matching, deconvolved trace for deconvolution, and the posterior of event times for the Bayesian approach), estimated event times, and the threshold used to determine those event times (for the Bayesian trace the threshold is so low that it cannot be seen above the baseline posterior). (B) Parametric curve showing the number of true positives and false positives as a function of peak detection threshold. The results are from detection across 10 simulated traces with a total of 180 events ranging from 0.5 to 10 pA. The highlighted point in each line corresponds to the threshold used in (A). (C) Same as in (B) except showing the true positive count as a function of the threshold. (D) Same as in (C) except showing the false positive count.

for each simulated trace for each method, we detected peaks above some threshold in these time series as a function of their standard deviations. Fig. 3A shows an example of results for each algorithm on simulated data. The threshold for each method's output is shown as a horizontal line, and inferred events are shown as dots. (The threshold for the Bayesian approach is difficult to see as it is very close to zero.)

By varying the threshold and counting the number of true positives and false positives, we can get a sense of how noisy the output is from each approach. The Bayesian approach is able to accurately detect more events while accumulating fewer false positives than the other methods (Fig. 3B). Importantly, it also maintains similar levels of true and false positives as the threshold is varied (Fig. 3C and D), indicating that our approach is more robust to the choice of threshold. This is achieved because of the more precise and less noisy representation of PSC timing produced by the Bayesian method which can be seen in the sharpness of the peaks in the Bayesian trace in Fig. 3A as compared to the other methods. The average error for true positives was also smallest for the Bayesian method (see Supplemental Fig. 1). While there exists a threshold for the template-matching and deconvolution methods that performs quite well, small deviations from this value lead to vastly more false positives or less true positives (Fig. 3C and D). When applying these methods to real data, it may not be possible to finely tune the threshold parameter on a per dataset basis since ground truth information is not available. This can be especially troublesome when online or closed-loop analysis is desired.

We next compared methods on real voltage-clamp recordings in which we could modulate the SNR of an event physiologically. To achieve this, we made paired patch recordings of a PV+ cell and a layer V pyramidal cell until we found an inhibitory connection from the fast-spiking cell onto the pyramidal cell with a probability of a postsynaptic event extremely close to 1.0. We then moved the holding potential for the postsynaptic cell towards the reversal potential for the inhibitory current. In this way, we could obtain

ground truth data in which we had direct control over the SNR (Fig. 4A–C, top traces).

We ran 50 trials each at three holding potentials representing relatively high, medium, and low SNR regimes and detected events using all three methods. Only the Bayesian approach was sensitive enough to detect events reliably in the low SNR case (Fig. 4A–C, rasters). Similar to the simulated data, the Bayesian approach was also the only method which was robust to the thresholding parameter, indicating that the posterior over event times is highly peaked. As expected, as the SNR decreases, the ability to accurately detect the timing of the event decreases.

Any approach will have hyperparameters that must be selected, and the Bayesian approach allows for tuning of hyperparameters corresponding to prior distributions on model parameters. We show that a single set of prior distribution hyperparameters can perform well across several physiological regimes by running inference on traces from different cell types and under different recording conditions while holding the hyperparameters constant. In Fig. 4G we show that with a single set of prior parameters our method can detect EPSCs across three different cell types: a layer II/III pyramidal cell, a SOM+ interneuron, and PV+ interneuron. Despite the differing statistics in each of these traces (event features and noise), event detection performs well. In 4H we show results for spontaneous IPSCs in a layer II/III pyramidal cell. For these results, we used the same prior parameters as in Fig. 4G except that we increased the adjusted the bounds for the time constants to account for the slower kinetics of IPSCs. Finally, in Fig. 4I we show results on an FS cell recorded *in vivo*. The prior parameter settings here are similar to those in Fig. 4G except that we had to increase the rate prior on the number of events. We stress that there is no correct prior setting, but that the priors allow the user to trade off between objectives like true versus false positives (e.g., through the rate parameter or minimum event amplitude). For reference, we include the hyperparameter settings for all inference results in Supplementary Table 1.

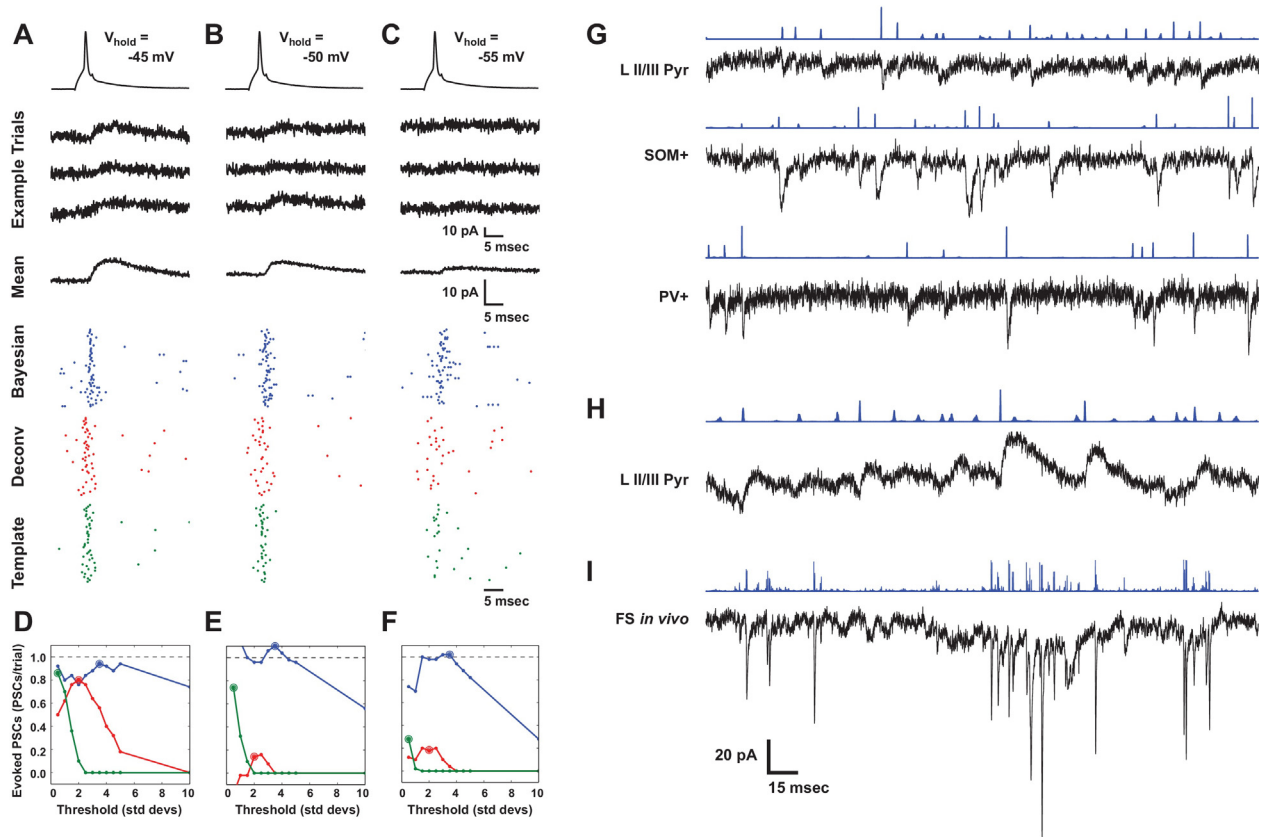


Fig. 4. Results on real voltage-clamp recordings. (A–F) Comparison to other detection methods on real data while modulating the SNR. (A) Top: a single trial of the spiking presynaptic cell firing from direct current injection. Middle traces: three example trials of the connected postsynaptic cell responding with an IPSC. The postsynaptic cell is being clamped at -45 mV. Below the examples is the mean trace over all 50 trials. Bottom rasters: estimated events for 50 trials for each detection method. (B and C) Same as in (A) except with the postsynaptic cell held at -50 mV (B) and -55 mV (C). The scales for the example traces are all the same across (A–C) and likewise for the mean traces. (D–F) The average number of evoked events counted in a window around the spike time (1.0 ms before to 5.0 ms after the spike) above a baseline rate of detected events per trial as a function of threshold for each method. Colors and symbols as in Fig. 2B and C. The dotted horizontal line represents perfect performance of one extra event detected per trial. (G–I) Examples of results on several types of real data. (G) Spontaneous EPSCs detected in a layer II/III pyramidal cell, top, a SOM+ cell, middle, and a PV+ cell, bottom. (H) Spontaneous IPSCs in a layer II/III pyramidal cell. (I) Spontaneous EPSCs and IPSCs detected in an *in vivo* recording from a FS cell.

3.3. Extension 1: direct optical stimulation artifact

For certain experiments, it may be productive to actively drive cells to fire in order to study circuit properties. In particular, we may want to combine voltage-clamp recordings with spatiotemporally structured optical stimulation of a putative presynaptic population of neurons. This can induce optically evoked artifacts from direct optical currents in the voltage-clamp recording. As an example, one could express an optogenetic channel pan-neuronally which would allow the simultaneous mapping EPSCs and IPSCs onto one cell. Under these conditions, the cell under voltage-clamp will also express opsin and respond to any stimulating light. Similarly, in neurotransmitter uncaging mapping experiments, there will be a direct stimulation of the postsynaptic cell at most stimulation sites close to the cell.

We show that our approach is able to decompose a simulated trace with a direct optical current in Fig. 5. Specifically, we simulated a trace consisting of many EPSCs with an additive direct optical stimulation current, consistent with Eq. (10). On this simulated data, inference performs very well in terms of extracting events simultaneously with artifact isolation.

In Fig. 5B–D, we present results on real data obtained by combining voltage-clamp recording with one-photon stimulation via a digital micromirror device (DMD) as well as two-photon, holographic stimulation of neurons. In Fig. 5B, we show direct currents at many different stimulation locations with one-photon

excitation via a DMD. Under these conditions, the shape of the current is not well characterized by a parameterized template function so we must use an empirically derived template. When normalized, all of the currents have roughly similar shapes, validating our approximation of the artifact as a scaled template. In Fig. 5C, we show PSC inference on three trials from a single stimulation location that has putative evoked PSCs overlapping the direct stimulation artifact. In Fig. 5D, we show similar results except the stimulation is performed using two-photon excitation with a spatial-light modulator (SLM). We found that under these conditions, we were able to use the model in Eq. (9) and fit two time constants to the optical current. This allowed for a more flexible fit that could account for the differences in the direct current shape as the stimulation location varied. The two sets of traces and results in Fig. 5D are for three trials at two different locations. These locations were chosen because they contained putative evoked EPSCs.

3.4. Extension 2: passive mapping experiment (simulation)

In addition to electrical recordings from a post-synaptic cell, we may also have calcium indicator available in pre-synaptic cells (Aaron and Yuste, 2006). This allows for a passive mapping experiment from many pre-synaptic cells (optically imaged) to a single post-synaptic cell (patched). Here we provide an example of the usage of the joint inference procedure (Fig. 6A). We have simulated a population of presynaptic cells which are observed via calcium

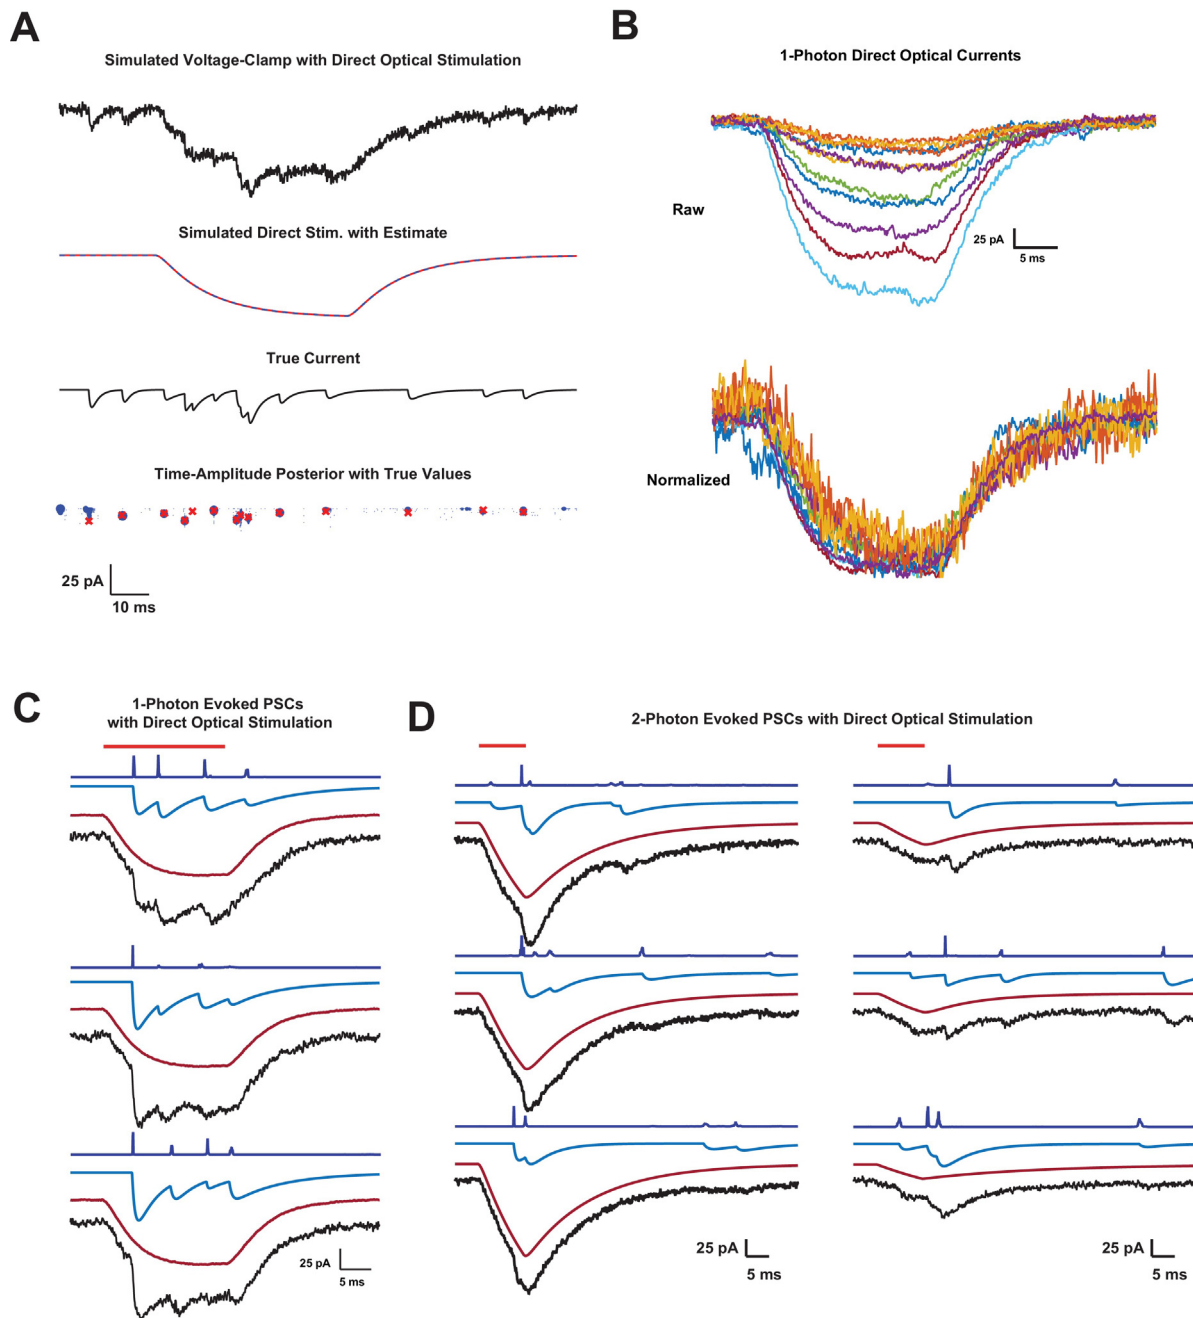


Fig. 5. Removing direct stimulation artifacts from active mapping data. (A) Top: a simulated voltage-clamp recording that is contaminated by a large direct, optical current. Top middle: the true direct current (blue) with an estimate of the inferred direct current (dashed, red). Bottom middle: the true current used in the top trace. Bottom, the inferred amplitude-time posterior (blue) with true amplitude-time coordinates overlaid (red X's). (B) Top: direct optical currents evoked in a single cell by stimulating different locations with one-photon excitation and a DMD. Bottom: same as above but each trace is normalized to have the same amplitude. (C and D) PSC detection with direct optical stimulation on real mapping data. (C) Inference results for one-photon, DMD-based mapping data with direct stimulation contamination. All three trials are for a single stimulation site. Red line shows when the stimulating laser is on. Dark blue trace shows the posterior over event times. Light blue trace shows the MAP estimate for the synaptic currents. Maroon trace, MAP estimate for the direct stimulation. Black, raw voltage-clamp observation. (D) Inference results for two-photon, SLM-based mapping data with inference. Left, three trials from a particular stimulation location which shows putative evoked EPSCs riding on top of direct stimulation with inference. Right, same as the results to the left but for a stimulation location further from the cell. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

imaging (SNR for these cells is plausible for a well-tuned setting, and bin size is 35 ms). These cells drive post-synaptic events in the simulated patched cell (noise is biologically plausible, and bin size is 0.05 ms). In Fig. 6A, 4 of 6 candidate presynaptic neurons are observed. Postsynaptic events evoked by unobserved neurons are inferred on a per event basis whereas events co-occurring with presynaptic events arise from a single variable (see model Eq. (12)).

We see that we can recover events jointly from the calcium imaging and electrophysiology, with event identity successfully linked across the two modalities. In addition, the histogram in the right panel of Fig. 6B indicates that combining electrophysiology and calcium imaging tends to yield increased temporal precision of event times compared to inference from calcium imaging alone (this intuitively follows from the fact that the electrophysiology has much higher temporal resolution and the Bayesian inference

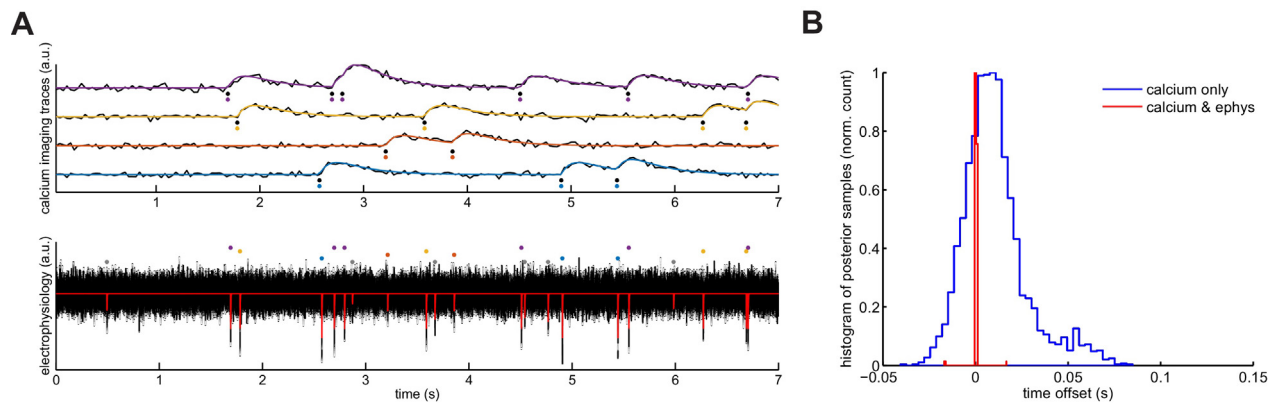


Fig. 6. Joint inference over simulated calcium and voltage data. (A) Observed data and true/inferred event times for four simultaneously observed calcium traces (top) and (for simplicity) just a single voltage trace (bottom). Colored dots indicate to which trace the algorithm matched each event; gray dots in bottom correspond to voltage events that were detected but (correctly) not matched to any detected calcium events. In this example the algorithm correctly matched all calcium-observed events. (B) Histograms of temporal error in estimated event times. Given calcium only, the variance of inferred event times is large; incorporating voltage information drastically reduces this variance, indicating significant potential gains in temporal resolution from this Bayesian data fusion approach. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

can combine information across modalities). This simulation serves as a proof-of-concept that this probabilistic approach can provide meaningful automatic analysis for passive mapping experiments.

4. Discussion

In this work, we have presented a probabilistic formulation of the event detection problem for electrophysiological recordings with an emphasis on PSC detection. This method works on simulated and real data and its performance compares favorably relative to existing methods. We have also shown how the probabilistic model can be incorporated into two extensions applicable for mapping experiments which make use either of optical stimulation or multi-modal fusion of electrophysiology with calcium imaging.

The PSC detection problem explored in this paper is similar to the more general set of problems in neuroscience (and signal processing more broadly) having to do with inferring events in noisy timeseries. In neuroscience, recent work has proposed probabilistic, MCMC-based tools for analysis of calcium imaging data (Vogelstein et al., 2009; Pnevmatikakis et al., 2013). However, we are not aware of similar tools yet being leveraged for analysis of electrophysiology data. The models developed in this paper are formally similar to those used in the calcium imaging setting, with appropriate modifications to allow for structured AR noise and distinct per event kinetics, as well as detailing the extensions related to the mapping setting.

While there are other discriminative methods that have been proposed for deconvolution (e.g. see Theis et al. (2016) for alternative event-detection methods for calcium imaging), these approaches lack the clear probabilistic generative semantics which facilitate modularity and extension to hierarchical models such as those useful for mapping experiments. Our approach also complements related work inferring synaptic inputs in a probabilistic fashion using particle filtering (Paninski et al., 2012) from voltage traces. The present approach focuses on current traces, allows for per event kinetics, is perhaps simpler to implement, and does not require temporal discretization.

The inference approach presented in this work makes most sense when PSC rates are relatively low such that overlap is limited. For very high levels of overlap (i.e. the high-rate case), we do not expect there is always hope to resolve precise timing of single events, and it is likely that approaches more like Vogelstein et al. (2010) or Paninski et al. (2012) may be appropriate. That said, we

have found that our method can still give sensible results even when event rates are high (for example, see Fig. 4I).

Aside from the extensions we considered, other sophisticated prior structure can be incorporated by making the model hierarchical. For example, in the single trace setting we might expect the events to cluster by pre-synaptic cell identity or cell type. Even without the additional observed traces employed in our mapping model, it would be conceivable to specify and infer latent source clusters based on structure in the distribution of shape or amplitude of post-synaptic events, depending on what were of most interest scientifically. In the simplest case, this would result in a mixture model for the events, with each event associated with a latent pre-synaptic source. We plan to pursue these model extensions in future work.

Acknowledgements

Funding for this research was provided by Global Brain Research Award 325398 from the Simons Foundation, ARO MURI W911NF-12-1-0594, and DARPA N66001-15-C-4032 (SIMPLEX); in addition, this work was supported by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior/Interior Business Center (DoI/IBC) contract number D16PC00003. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. *Disclaimer:* The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DoI/IBC, or the U.S. Government. B.S. is supported by a Fannie and John Hertz Foundation Fellowship and an NSF Graduate Research Fellowship.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at <http://dx.doi.org/10.1016/j.jneumeth.2016.05.015>.

References

- Aaron, G., Yuste, R., 2006. Reverse optical probing (roping) of neocortical circuits. *Synapse* 60 (6), 437.
- Ankri, N., Legendre, P., Faber, D., Korn, H., 1994. Automatic detection of spontaneous synaptic responses in central neurons. *J. Neurosci. Methods* 52 (1), 87–100.

- Bar-Yehuda, D., Korngreen, A., 2008. Space-clamp problems when voltage clamping neurons expressing voltage-gated conductances. *J. Neurophysiol.* 99 (3), 1127–1136.
- Callaway, E.M., Yuste, R., 2002. Stimulating neurons with light. *Curr. Opin. Neurobiol.* 12 (5), 587–592.
- Chib, S., Greenberg, E., 1994. Bayes inference in regression models with arma (p, q) errors. *J. Econom.* 64 (1), 183–206.
- Clements, J., Bekkers, J., 1997. Detection of spontaneous synaptic events with an optimally scaled template. *Biophys. J.* 73 (1), 220.
- Fenno, L., Yizhar, O., Deisseroth, K., 2011. The development and application of optogenetics. *Annu. Rev. Neurosci.* 34 (1), 389–412.
- Gelman, A., Carlin, J.B., Stern, H.S., Rubin, D.B., 2014. *Bayesian Data Analysis*, vol. 2. Taylor & Francis.
- Guzman, S.J., Schlögl, A., Schmidt-Hieber, C., 2014. Stimfit: quantifying electrophysiological data with python. *Front. Neuroinform.* 8 (16), online.
- Hwang, T.N., Copenhagen, D.R., 1999. Automatic detection, characterization, and discrimination of kinetically distinct spontaneous synaptic events. *J. Neurosci. Methods* 92 (1), 65–73.
- Jonas, P., Major, G., Sakmann, B., 1993. Quantal components of unitary EPSCs at the mossy fibre synapse on CA3 pyramidal cells of rat hippocampus. *J. Physiol.* 472 (1), 615–663.
- Katzel, D., Zemelman, B.V., Buettner, C., Wolfel, M., Miesenböck, G., 2011. The columnar and laminar organization of inhibitory connections to neocortical excitatory cells. *Nat. Neurosci.* 14 (1), 100–107.
- Kudoh, S.N., Taguchi, T., 2002. A simple exploratory algorithm for the accurate and fast detection of spontaneous synaptic events. *Biosens. Bioelectron.* 17 (9), 773–782.
- Lay, D.C., 2012. *Linear Algebra and its Applications*, 4th ed. Pearson.
- Lehmann, E.L., Casella, G., 1998. *Theory of Point Estimation*, vol. 31. Springer Science & Business Media.
- Li, G.-H., Jackson, M.F., MacDonald, J.F., 2007. Weighted least squares fitting with multiple templates for detection of small spontaneous signals. *J. Neurosci. Methods* 164 (1), 139–148.
- Mardinly, A.R., Spiegel, I., Patrizi, A., Centofante, E., Bazinet, J.E., Tzeng, C.P., Mandel-Brehm, C., Harmin, D.A., Adesnik, H., Fagioli, M., Greenberg, M.E., 2016. Sensory experience regulates cortical inhibition by inducing IGF1 in VIP neurons. *Nature* 531 (7594), 371–375.
- Medelin, M., Rancic, V., Cellot, G., Laishram, J., Veeraraghavan, P., Rossi, C., Muzio, L., Sivilotti, L., Ballerini, L., 2016. Altered development in GABA co-release shapes glycinergic synaptic currents in cultured spinal slices of the SOD1G93A mouse model of ALS. *J. Physiol.*
- Mena, G., Grosberg, L., Kellison-Linn, F., Chichilnisky, E., Paninski, L., 2015. Large-scale multi electrode array spike sorting algorithm introducing concurrent recording and stimulation. In: *NIPS Workshop on Statistical Methods for Understanding Neural Systems*.
- Moller, J., Waagepetersen, R.P., 2004. *Statistical Inference and Simulation for Spatial Point Processes*. CRC Press.
- Neal, R.M., 1993. *Probabilistic Inference using Markov Chain Monte Carlo Methods*. Department of Computer Science, University of Toronto, Toronto, Canada.
- Paninski, L., Vidne, M., DePasquale, B., Ferreira, D.G., 2012. Inferring synaptic inputs given a noisy voltage trace via sequential Monte Carlo methods. *J. Comput. Neurosci.* 33 (1), 1–19.
- Pernia-Andrade, A.J., Goswami, S.P., Stickler, Y., Fröbe, U., Schlögl, A., Jonas, P., 2012. A deconvolution-based method with high sensitivity and temporal resolution for detection of spontaneous synaptic currents in vitro and in vivo. *Biophys. J.* 103 (7), 1429–1439.
- Pluta, S., Naka, A., Veit, J., Telian, G., Yao, L., Hakim, R., Taylor, D., Adesnik, H., 2015. A direct transaminar inhibitory circuit tunes cortical output. *Nat. Neurosci.* 18 (11), 1631–1640.
- Pnevmatikakis, E.A., Merel, J., Pakman, A., Paninski, L., 2013. Bayesian spike inference from calcium imaging data. In: *Asilomar Conference on Signals, Systems and Computers*. IEEE, pp. 349–353.
- Richardson, M.J., Silberberg, G., 2008. Measurement and analysis of postsynaptic potentials using a novel voltage-deconvolution method. *J. Neurophysiol.* 99 (2), 1020–1031.
- Rickgauer, J.P., Deisseroth, K., Tank, D.W., 2014. Simultaneous cellular-resolution optical perturbation and imaging of place cell firing fields. *Nat. Neurosci.* 17 (12), 1816–1824.
- Rosenthal, J.S., 2011. Optimal proposal distributions and adaptive MCMC. In: *Handbook of Markov Chain Monte Carlo*, pp. 93–112.
- Scanziani, M., Häusser, M., 2009. Electrophysiology in the age of light. *Nature* 461 (7266), 930–939.
- Shababo, B., Paige, B., Pakman, A., Paninski, L., 2013. Bayesian inference and online experimental design for mapping neural microcircuits. *Adv. Neural Inf. Process. Syst. (NIPS)*, 1304–1312.
- Shepherd, G.M.G., Stepanyants, A., Bureau, I., Chklovskii, D., Svoboda, K., 2005. Geometric and functional organization of cortical circuits. *Nat. Neurosci.* 8 (6), 782–790.
- Shi, Y., Nenadic, Z., Xu, X., 2010. Novel use of matched filtering for synaptic event detection and extraction. *PLoS ONE* 5 (11), e15517.
- Tan, V., Goyal, V., 2008. Estimating signals with finite rate of innovation from noisy samples: a stochastic algorithm. *IEEE Trans. Signal Process.* 56 (10), 5135–5146.
- Theis, L., Berens, P., Froudarakis, E., Reimer, J., Roman-Roson, M., Baden, T., Euler, T., Tolias, A.S., Bethge, M., 2016. Benchmarking spike rate inference in population calcium imaging. *Neuron* 90 (3), 471–482.
- Vogelstein, J.T., Watson, B.O., Packer, A.M., Yuste, R., Jerny, B., Paninski, L., 2009. Spike inference from calcium imaging using sequential Monte Carlo methods. *Biophys. J.* 97 (2), 636–655.
- Vogelstein, J.T., Packer, A.M., Machado, T.A., Sippy, T., Babadi, B., Yuste, R., Paninski, L., 2010. Fast nonnegative deconvolution for spike train inference from population calcium imaging. *J. Neurophysiol.* 104 (6), 3691–3704.
- Wiener, N., 1949. *Extrapolation, Interpolation, and Smoothing of Stationary Time Series: With Engineering Applications*, Technology Press Books in Science and Engineering. Technology Press of the Massachusetts Institute of Technology.