

Fast nonconvex deconvolution of calcium imaging data

SEAN W. JEWELL*

Department of Statistics, University of Washington, Seattle, WA 98195, USA
swjewell@uw.edu

TOBY DYLAN HOCKING

School of Informatics, Computing, and Cyber Systems, Northern Arizona University, Flagstaff, AZ 83011, USA

PAUL FEARNHEAD

Department of Mathematics and Statistics, Lancaster University, Lancaster LA1 4YF, UK

DANIELA M. WITTEN

Department of Statistics, University of Washington, Seattle, WA 98195, USA and Department of Biostatistics, University of Washington, Seattle, WA 98195, USA

SUMMARY

Calcium imaging data promises to transform the field of neuroscience by making it possible to record from large populations of neurons simultaneously. However, determining the exact moment in time at which a neuron spikes, from a calcium imaging data set, amounts to a non-trivial deconvolution problem which is of critical importance for downstream analyses. While a number of formulations have been proposed for this task in the recent literature, in this article, we focus on a formulation recently proposed in Jewell and Witten (2018. Exact spike train inference via ℓ_0 optimization. *The Annals of Applied Statistics* **12**(4), 2457–2482) that can accurately estimate not just the spike rate, but also the specific times at which the neuron spikes. We develop a much faster algorithm that can be used to deconvolve a fluorescence trace of 100 000 timesteps in less than a second. Furthermore, we present a modification to this algorithm that precludes the possibility of a “negative spike”. We demonstrate the performance of this algorithm for spike deconvolution on calcium imaging datasets that were recently released as part of the `spikefinder` challenge (<http://spikefinder.codeneuro.org/>). The algorithm presented in this article was used in the Allen Institute for Brain Science’s “platform paper” to decode neural activity from the Allen Brain Observatory; this is the main scientific paper in which their data resource is presented. Our C++ implementation, along with R and python wrappers, is publicly available. R code is available on CRAN and Github, and python wrappers are available on Github; see <https://github.com/jewellsean/FastLZeroSpikeInference>.

Keywords: Calcium imaging; Changepoint detection; Neuroscience; Nonconvex optimization.

*To whom correspondence should be addressed.

1. INTRODUCTION

Due to recent advances in calcium imaging technology, it has become possible to record from large populations of neurons simultaneously in behaving animals (Dombbeck *and others*, 2007; Ahrens *and others*, 2013; Prevedel *and others*, 2014). These data result in a fluorescence trace for each neuron.

However, most downstream analyses require not a fluorescence trace, but instead a measure of the neuron's activity over time. Consequently, a number of unsupervised and—more recently—supervised methods have been developed to infer neural activity on the basis of the fluorescence trace (Yaksi and Friedrich, 2006; Holekamp *and others*, 2008; Sasaki *and others*, 2008; Dyer *and others*, 2010, 2013; Grewe *and others*, 2010; Vogelstein *and others*, 2009, 2010; Pnevmatikakis *and others*, 2013; Deneux *and others*, 2016; Friedrich and Paninski, 2016; Theis *and others*, 2016; Friedrich *and others*, 2017; Jewell and Witten, 2018).

In this article, we make use of a generative model that connects the observed fluorescence trace y_t to the underlying and unobserved calcium concentration c_t , and the unknown spike times (Vogelstein *and others*, 2010; Friedrich and Paninski, 2016; Friedrich *and others*, 2017). This model assumes that the observed fluorescence is a noisy version of the underlying calcium, which exponentially decays, unless there is a spike, in which case there is an instantaneous increase in the calcium concentration, as follows:

$$\begin{aligned} y_t &= \beta_0 + \beta_1 c_t + \epsilon_t, & \epsilon_t &\sim_{\text{ind.}} (0, \sigma^2), & t &= 1, \dots, T; \\ c_t &= \gamma c_{t-1} + z_t, & & & t &= 2, \dots, T, \end{aligned} \quad (1.1)$$

where $z_t \geq 0$, and where $z_t > 0$ indicates the presence of a spike at the t th timestep. At most timesteps $z_t = 0$, corresponding to no spike, and the calcium will decay exponentially at a rate governed by the parameter γ , which is assumed known. For simplicity, in what follows, we assume that the intercept β_0 is equal to zero. However, this is easy to relax; see Section 2.4 for a straight-forward extension, and Section 3 for practical considerations. Moreover, we set β_1 equal to one, since the problems we will solve are scale-invariant. That is, the value of β_1 does not affect the scientific conclusions.

Under the additional assumption that the errors ϵ_t are normally distributed, model (1.1) suggests estimating the concentration by solving the following constrained ℓ_0 optimization problem

$$\underset{c_1, \dots, c_T, z_2, \dots, z_T}{\text{minimize}} \left\{ \frac{1}{2} \sum_{t=1}^T (y_t - c_t)^2 + \lambda \sum_{t=2}^T 1_{(z_t \neq 0)} \right\} \text{ subject to } z_t = c_t - \gamma c_{t-1} \geq 0, \quad (1.2)$$

where λ is a non-negative tuning parameter that controls the tradeoff between how closely the calcium concentration matches the fluorescence trace, $\sum_{t=1}^T (y_t - c_t)^2$, and the number of non-zero spikes, $\sum_{t=2}^T 1_{(z_t \neq 0)}$. The solution to this optimization problem directly provides an estimate for the spike times; that is, if $\hat{z}_t \neq 0$, then we infer a spike at time t . We note that this problem is over-parameterized, in the sense that knowing c_1, \dots, c_T determines z_2, \dots, z_T .

While (1.2) follows from the biological process described in (1.1), the ℓ_0 penalty makes the problem nonconvex and thus seemingly intractable. Consequently, rather than solving (1.2), prior approaches have solved a convex relaxation to (1.2) (Vogelstein *and others*, 2010; Friedrich and Paninski, 2016; Friedrich *and others*, 2017), where the ℓ_0 penalty is replaced by an ℓ_1 penalty.

In recent work, Jewell and Witten (2018) showed that it is possible to efficiently solve the related nonconvex optimization problem

$$\underset{c_1, \dots, c_T, z_2, \dots, z_T}{\text{minimize}} \left\{ \frac{1}{2} \sum_{t=1}^T (y_t - c_t)^2 + \lambda \sum_{t=2}^T 1_{(z_t \neq 0)} \right\} \text{ subject to } z_t = c_t - \gamma c_{t-1}, \quad (1.3)$$

obtained by removing the positivity constraint, $c_t - \gamma c_{t-1} \geq 0$, from (1.2). The positivity constraint enforces the biological property that a firing neuron can only cause the calcium concentration to increase. Nonetheless, despite the slight loss in physical interpretability caused by the omission of the positivity constraint, Jewell and Witten (2018) showed that solving (1.3) leads to improved performance over existing deconvolution approaches that perform a convex relaxation of (1.2). In particular, the method of Jewell and Witten (2018) provides an accurate estimate of the *specific timesteps at which a neuron fires*.

Unfortunately, the algorithm proposed in Jewell and Witten (2018) for solving (1.3) is too slow to conveniently run on large-scale data. For traces of 100 000 timesteps, the implementation runs in a few minutes for a single value of the tuning parameter λ ; in practice the user must apply the algorithm over a fine grid of values of λ , leading potentially to hours of computation time for a single trace. Furthermore, a single experiment could result in hundreds or thousands of fluorescence traces (Ahrens and others, 2013; Vladimirov and others, 2014).

In this article, we develop a fast algorithm for solving problem (1.3); for traces of 100 000 timesteps our implementation runs in less than a second. Furthermore, this new algorithm can easily accommodate the positivity constraint that was omitted from (1.3); in other words, we can directly solve problem (1.2). Additionally, we exploit ideas from Haynes and others (2017) to efficiently “choose” good values of λ ; that is, values of λ where the solution to (1.3) changes.

The algorithm we develop to solve (1.2) was used to obtain the key scientific results in the Allen Institute’s main scientific paper from the Allen Brain Observatory (de Vries and others, 2018). Additionally, the Allen Institute for Brain Science recently released an update to their software development kit that provides users with the output from our algorithm for close to 60 000 neurons during different experimental conditions.

In what follows, we introduce our new algorithm for solving (1.2) and (1.3) in Section 2. We compare its performance in Section 3 to a convex relaxation of (1.2) on a number of calcium imaging datasets that were recently released as part of the `spikefinder` challenge (<http://spikefinder.codeneuro.org/>). We close with a discussion in Section 4.

2. A FAST FUNCTIONAL PRUNING ALGORITHM FOR SOLVING PROBLEMS (1.3) AND (1.2)

2.1. A review of Jewell and Witten (2018)

Jewell and Witten (2018) point out that the ℓ_0 optimization problem (1.3) is equivalent to a changepoint detection problem,

$$\text{minimize}_{0=\tau_0 < \tau_1 < \dots < \tau_k < \tau_{k+1}=T, k} \left\{ \sum_{j=0}^k \mathcal{D}(y_{(\tau_j+1):\tau_{j+1}}) + \lambda k \right\}, \quad (2.4)$$

where

$$\mathcal{D}(y_{a:b}) \equiv \min_{\alpha} \left\{ \frac{1}{2} \sum_{t=a}^b (y_t - \alpha \gamma^{t-b})^2 \right\}. \quad (2.5)$$

In problem (2.4), we select the optimal changepoints τ_1, \dots, τ_k and the number of changepoints k such that the cost of segmenting the data into $k + 1$ exponentially decaying regions is minimal, where (2.5) is the cost associated with the region that spans the a th to b th timesteps. Problems (2.4) and (1.3) are equivalent in the sense that $\hat{z}_{\tau_1+1} \neq 0, \dots, \hat{z}_{\tau_k+1} \neq 0$ and all other $\hat{z}_t = 0$.

To solve the changepoint problem, [Jewell and Witten \(2018\)](#) exploit a simple recursion ([Jackson and others, 2005](#)),

$$F(s) = \min_{0=\tau_0 < \tau_1 < \dots < \tau_k < \tau_{k+1}=s, k} \left\{ \sum_{j=0}^k \mathcal{D}(y_{(\tau_j+1):\tau_{j+1}}) + \lambda k \right\} = \min_{0 \leq \tau < s} \{F(\tau) + \mathcal{D}(y_{(\tau+1):s}) + \lambda\}, \quad (2.6)$$

where $F(s)$ is the optimal cost of segmenting the data $y_{1:s} \equiv [y_1, \dots, y_s]$, and where we define $F(0) \equiv -\lambda$. This results in an algorithm with computational complexity $\mathcal{O}(T^2)$, which can be substantially improved by noticing that the minimization on the right hand side of (2.6) can be performed over a smaller set \mathcal{E}_s without sacrificing the global optimum ([Killick and others, 2012](#)); details are provided in [Jewell and Witten \(2018\)](#). As mentioned in the introduction, the algorithm runs in a few minutes for traces of length 100 000, and yields the global optimum to (1.3). We note that the recursion (2.6) does not naturally lead to an algorithm to solve (1.2); this is discussed in further detail in Section 2.3.

2.2. Functional pruning for solving (1.3)

2.2.1. Motivation for functional pruning. In order to motivate the potential for a much faster algorithm for solving (1.3) than the one proposed in [Jewell and Witten \(2018\)](#), consider Figure 1a.

In this figure, we are interested in determining the optimal cost of segmenting the data up to time 40, that is, calculating $F(40)$ in (2.6). Instead of directly applying the recursion (2.6), we consider a slightly different question: What is the optimal most recent changepoint before the 40th timestep, conditional on, the unknown calcium concentration c_{40} ? Given the previously stored values $F(0), \dots, F(39)$, and the data $y_{1:40}$, it is straightforward to calculate the best most recent changepoint $\tau^*(c_{40})$, as $\tau^*(c_{40}) = \operatorname{argmin}_{0 \leq \tau < 40} \left\{ F(\tau) + \frac{1}{2} \sum_{t=\tau+1}^{40} (y_t - \gamma^{t-40} c_{40})^2 + \lambda \right\}$, for any value of the calcium concentration c_{40} .

Figure 1a displays the most recent changepoint $\tau^*(c_{40})$ as a function of c_{40} . We observe that regardless of the value of the calcium at the current timestep—and consequently, regardless of the fluorescence values $y_{41}, y_{42}, y_{43}, \dots, y_T$ —the only possible times for the most recent changepoint before the 40th timestep are 20, 37, and 39; that is, $\tau^*(c_{40}) \in \{20, 37, 39\}$ for all possible c_{40} .

However, the algorithm proposed in [Jewell and Witten \(2018\)](#) does not exploit the fact that 20, 37, and 39 are the only possible times for the most recent changepoint before the 40th timestep: the minimization in (2.6) is performed over the set $\{0, \dots, 39\}$, or else over a slightly smaller set $\{18, \dots, 39\}$ using ideas from [Killick and others \(2012\)](#). This suggests that by viewing the cost of segmenting the data up until the s th timestep as a function of the calcium at the s th timestep, we could potentially develop an algorithm that is much faster than the one in [Jewell and Witten \(2018\)](#) in that it would only require performing the minimization in (2.6) over $\{20, 37, 39\}$. The idea of using this type of conditioning was first suggested by [Rigail \(2015\)](#) and [Maidstone and others \(2017\)](#), albeit to speed up algorithms for detecting changepoints in a different class of models.

2.2.2. The functional pruning algorithm. To begin, we substitute the cost function $\mathcal{D}(y_{(\tau+1):s})$ into the recursion (2.6), in order to obtain

$$\begin{aligned} F(s) &= \min_{0 \leq \tau < s} \{F(\tau) + \mathcal{D}(y_{(\tau+1):s}) + \lambda\} \\ &= \min_{0 \leq \tau < s} \left\{ F(\tau) + \min_{\alpha} \left\{ \frac{1}{2} \sum_{t=\tau+1}^s (y_t - \alpha \gamma^{t-s})^2 \right\} + \lambda \right\} \end{aligned}$$

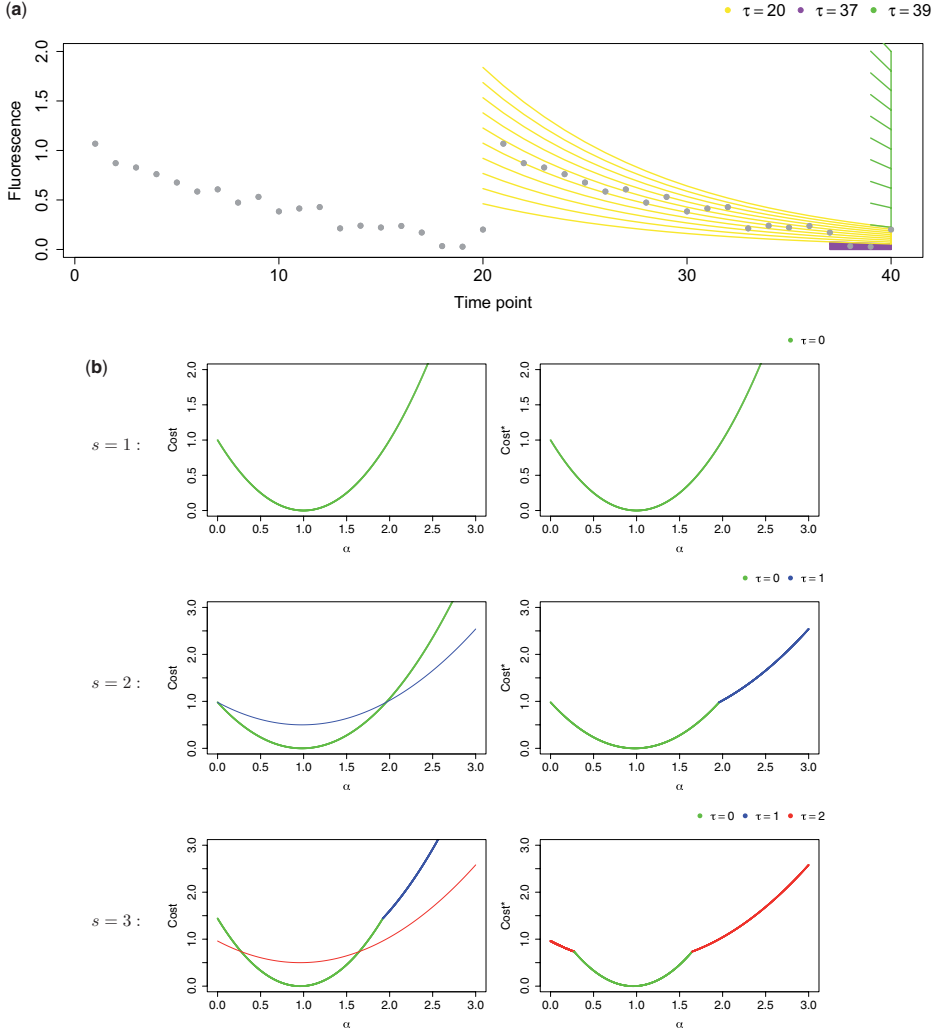


Fig. 1. Algorithm 1: (a) motivation and (b) example. (a) A simple example to show that there are only a few possible values for the most recent changepoint before timestep 40. We consider solving for the most recent changepoint, given data $y_{1:40}$, for each possible value of the calcium concentration at the 40th timestep, c_{40} . For each possible value of c_{40} , we display the estimated calcium concentration going back in time to the most recent changepoint before timestep 40. The colors indicate the time of the most recent changepoint. In this example, there are only three possibilities for the most recent changepoint: $\{20, 37, 39\}$. For example, $\tau^*(0.001) = 37$, $\tau^*(0.02) = 20$, and $\tau^*(1) = 39$. (b) Evolution of Cost_s^{τ} and $\text{Cost}_s^*(\alpha)$ for Example 1. The left-hand panels display the functions $\text{Cost}_{s-1}^*(\alpha/\gamma) + \frac{1}{2}(y_s - \alpha)^2$ and $\min_{\alpha'} \text{Cost}_{s-1}^*(\alpha') + \lambda + \frac{1}{2}(y_s - \alpha)^2$, and the right-hand panels show the function $\text{Cost}_s^*(\alpha)$, which is the minimum of those two functions. Rows index the timesteps, $s = 1, 2, 3$. The functions are colored based on the timestep of the most recent changepoint, that is, the value of τ corresponding to \mathcal{R}_s^{τ} . *Top*: When $s = 1$, $\text{Cost}_1^*(\alpha) = \frac{1}{2}(y_1 - \alpha)^2$; this corresponds to the region $\mathcal{R}_1^0 = [0, \infty)$. *Center*: When $s = 2$, $\text{Cost}_2^*(\alpha)$ is the minimum of two quantities: $\text{Cost}_1^*(\alpha/\gamma) + \frac{1}{2}(y_2 - \alpha)^2$, which corresponds to the most recent changepoint being at timestep zero, and $\min_{\alpha'} \text{Cost}_1^*(\alpha') + \lambda + \frac{1}{2}(y_2 - \alpha)^2$, which corresponds to the most recent changepoint being at timestep one. These two functions are shown on the left-hand side, and $\text{Cost}_2^*(\alpha)$ is shown on the right-hand side. *Bottom*: When $s = 3$, $\text{Cost}_3^*(\alpha)$ is calculated similarly; see Example 1 for additional details.

$$\begin{aligned}
&= \min_{\alpha} \min_{0 \leq \tau < s} \left\{ F(\tau) + \left\{ \frac{1}{2} \sum_{t=\tau+1}^s (y_t - \alpha \gamma^{t-s})^2 \right\} + \lambda \right\} \\
&= \min_{\alpha} \min_{0 \leq \tau < s} \text{Cost}_s^{\tau}(\alpha) \\
&= \min_{\alpha} \text{Cost}_s^{*}(\alpha), \tag{2.7}
\end{aligned}$$

where

$$\text{Cost}_s^{\tau}(\alpha) \equiv F(\tau) + \frac{1}{2} \sum_{t=\tau+1}^s (y_t - \alpha \gamma^{t-s})^2 + \lambda, \tag{2.8}$$

and

$$\text{Cost}_s^{*}(\alpha) = \min_{0 \leq \tau < s} \text{Cost}_s^{\tau}(\alpha). \tag{2.9}$$

In words, $\text{Cost}_s^{\tau}(\alpha)$ is the cost of partitioning the data up until time s , given that the most recent changepoint was at time τ , and the calcium at the s th timestep equals α . $\text{Cost}_s^{*}(\alpha)$ is the optimal cost of partitioning the data up until time s , given that the calcium at the s th timestep equals α .

The following proposition will prove useful in what follows.

Proposition 1 For $\text{Cost}_s^{*}(\alpha)$ defined in (2.9), the following recursion holds:

$$\text{Cost}_s^{*}(\alpha) = \min \left\{ \text{Cost}_{s-1}^{*}(\alpha/\gamma), \min_{\alpha'} \text{Cost}_{s-1}^{*}(\alpha') + \lambda \right\} + \frac{1}{2}(y_s - \alpha)^2. \tag{2.10}$$

The proof of Proposition 1 is in [Appendix S1 of the supplementary material](#) available at *Biostatistics* online. The recursion in (2.10) encompasses two possibilities: either there is a changepoint at the $(s-1)$ st timestep, and we must determine the optimal cost up to that time, $\min_{\alpha'} \text{Cost}_{s-1}^{*}(\alpha') + \lambda + \frac{1}{2}(y_s - \alpha)^2$, or there is no changepoint at the $(s-1)$ st timestep, $\text{Cost}_{s-1}^{*}(\alpha/\gamma) + \frac{1}{2}(y_s - \alpha)^2$. The recursion in (2.10) is reminiscent of (2.6), and raises the following question: can we use (2.10) as the basis for a recursive algorithm for solving the problem of interest, (1.3)? At first, it appears almost hopeless, since the recursion (2.10) involves a *function* of α , a real-valued parameter. However, as we will see, it turns out that $\text{Cost}_s^{*}(\alpha)$ and $\text{Cost}_s^{\tau}(\alpha)$ are simple functions of α that are easy to analytically manipulate.

Observe that, by definition (2.9), the optimal cost $\text{Cost}_s^{*}(\alpha)$ takes the form

$$\text{Cost}_s^{*}(\alpha) = \begin{cases} \text{Cost}_s^0(\alpha), & \alpha \in \mathcal{R}_s^0, \\ \vdots & \vdots \\ \text{Cost}_s^{s-1}(\alpha), & \alpha \in \mathcal{R}_s^{s-1}, \end{cases} \tag{2.11}$$

where $\mathcal{R}_s^{\tau} \equiv \left\{ \alpha : \min_{0 \leq \tau' < s} \text{Cost}_s^{\tau'}(\alpha) = \text{Cost}_s^{\tau}(\alpha) \right\}$; this is the set of values for the calcium at the s th timestep such that the most recent changepoint occurred at time τ . Furthermore, by inspection of (2.8), we see that $\text{Cost}_s^{\tau}(\alpha)$ is itself a quadratic function of α for all τ . Thus, $\text{Cost}_s^{*}(\alpha)$ is in fact *piecewise quadratic*. This means that in order to efficiently store the function $\text{Cost}_s^{*}(\alpha)$, we must simply keep track of the regions $\mathcal{R}_s^0, \dots, \mathcal{R}_s^{s-1}$, as well as the three coefficients (constant, linear, quadratic) that define the

quadratic function corresponding to each region. We will now present a small toy example illustrating how the recursion (2.10) can be used to build up optimal cost functions, each of which is piecewise quadratic.

EXAMPLE 1 Consider the simple dataset $y = [1.00, 0.98, 0.96, \dots]$ with $\lambda = \frac{1}{2}$ and $\gamma = 0.98$. We start with $\text{Cost}_1^*(\alpha)$, which is just a quadratic centered around y_1 ,

$$\text{Cost}_1^*(\alpha) = \text{Cost}_1^0(\alpha) = \frac{1}{2}(y_1 - \alpha)^2 = \frac{1}{2}(1.00 - \alpha)^2, \quad \alpha \in \mathcal{R}_1^0 \equiv [0, \infty).$$

Then, at the next time point, we form $\text{Cost}_2^*(\alpha)$ based on (2.10),

$$\begin{aligned} \text{Cost}_2^*(\alpha) &= \min \left\{ \text{Cost}_1^*(\alpha/\gamma), \min_{\alpha'} \text{Cost}_1^*(\alpha') + \lambda \right\} + \frac{1}{2}(y_2 - \alpha)^2 \\ &= \min \left\{ \frac{1}{2}(1 - \alpha/\gamma)^2, 0 + \frac{1}{2} \right\} + \frac{1}{2}(0.98 - \alpha)^2 \\ &= \begin{cases} \frac{1}{2}(1 - \alpha/\gamma)^2 + \frac{1}{2}(0.98 - \alpha)^2, & \alpha \in \mathcal{R}_2^0 \equiv [0, 2\gamma) \\ \frac{1}{2} + \frac{1}{2}(0.98 - \alpha)^2, & \alpha \in \mathcal{R}_2^1 \equiv [2\gamma, \infty) \end{cases} \end{aligned}$$

Again, using the recursion (2.10) we obtain the next optimal cost function,

$$\begin{aligned} \text{Cost}_3^*(\alpha) &= \min \left\{ \text{Cost}_2^*(\alpha/\gamma), \min_{\alpha'} \text{Cost}_2^*(\alpha') + \lambda \right\} + \frac{1}{2}(y_3 - \alpha)^2 \\ &= \min \left\{ \text{Cost}_2^*(\alpha/\gamma), \frac{1}{2} \right\} + \frac{1}{2}(0.96 - \alpha)^2 \\ &= \begin{cases} \frac{1}{2} + \frac{1}{2}(0.96 - \alpha)^2, & \alpha \in \mathcal{R}_3^2 \equiv \gamma^2 \left\{ \left[0, 1 - \frac{1}{\sqrt{1+\gamma^2}} \right) \cup \left[1 + \frac{1}{\sqrt{1+\gamma^2}}, \infty \right) \right\} \\ \frac{1}{2}(1 - \alpha/\gamma^2)^2 + \frac{1}{2}(0.98 - \alpha/\gamma)^2 + \frac{1}{2}(0.96 - \alpha)^2, & \alpha \in \mathcal{R}_3^0 \equiv \gamma^2 \left[1 - \frac{1}{\sqrt{1+\gamma^2}}, 1 + \frac{1}{\sqrt{1+\gamma^2}} \right) \end{cases} \end{aligned}$$

We note that $\text{Cost}_3^*(\alpha)$ is defined over just \mathcal{R}_3^0 and \mathcal{R}_3^2 . This example is displayed in Figure 1b.

Although we have shown how to efficiently build optimal cost functions $\text{Cost}_s^*(\alpha)$ from $s = 1, \dots, T$, it remains to establish that these cost functions can be used to determine the optimal changepoints, that is, the values of τ_1, \dots, τ_k that solve (1.3). These can be obtained by finding the value of τ that satisfies

$$\tau^*(s) = \{ \tau : \min_{\alpha} \text{Cost}_s^{\tau}(\alpha) = \min_{\alpha} \text{Cost}_s^*(\alpha) \} \quad (2.12)$$

for $\tau^*(T), \tau^*(\tau^*(T)), \dots$ until 0 is obtained. Full details are provided in Algorithm 1. To summarize, we have developed a recursive algorithm for solving (1.3) using the recursions in Proposition 1.

EXAMPLE 2 “Example 1 revisited”

We return to Example 1 to illustrate how (2.12) can be used to determine the optimal changepoints. In the interest of simplicity, we assume that $T = 3$; in other words, we have observed all of the data. Then, $\tau^*(3) = \{ \tau : \min_{\alpha} \text{Cost}_3^{\tau}(\alpha) = \min_{\alpha} \text{Cost}_3^*(\alpha) \}$, where

$$\min_{\alpha} \text{Cost}_3^{\tau}(\alpha) = \begin{cases} \min_{\alpha} \text{Cost}_3^2(\alpha) = 0.73, & \alpha \in \mathcal{R}_3^2 \\ \min_{\alpha} \text{Cost}_3^0(\alpha) = 5.4 \times 10^{-8}, & \alpha \in \mathcal{R}_3^0 \end{cases}.$$

Algorithm 1: A functional pruning algorithm for solving (1.3)

Initialize: Compute $\text{Cost}_1^*(\alpha) := \text{Cost}_1^0(\alpha) = \frac{1}{2}(y_1 - \alpha)^2$, and set $\mathcal{R}_1^0 = [0, \infty)$

- 1 **foreach** *timestep* $s = 2, \dots, T$ **do**
- 2 Calculate and store $\text{Cost}_s^*(\alpha) := \min\{\text{Cost}_{s-1}^*(\alpha/\gamma), \min_{\alpha'} \text{Cost}_{s-1}^*(\alpha') + \lambda\} + \frac{1}{2}(y_s - \alpha)^2$
- 3 Set $\mathcal{R}_s^{s-1} = \{\alpha : \text{Cost}_s^*(\alpha) = \min_{\alpha'} \text{Cost}_{s-1}^*(\alpha') + \lambda + \frac{1}{2}(y_s - \alpha)^2\}$
- 4 **foreach** $\tau = 0, \dots, s-1$ **do**
- 5 $\mathcal{R}_s^\tau = (\gamma \mathcal{R}_{s-1}^\tau) \cap (\mathcal{R}_s^{s-1})^c$
- 6 **end**
- 7 **end**
- 8 Initialize list of changepoints $cp := (T)$
- 9 Set the current changepoint $\tau^{cur} := T$
- 10 Initialize list of estimated calcium concentrations $c := ()$
- 11 **while** $\tau^{cur} > 0$ **do**
- 12 $\tau^{prev} := \tau^{cur}$
- 13 Determine the most recent changepoint $\tau^{cur} := \left\{ \tau : \underset{\alpha}{\text{argmin}} \{ \text{Cost}_{\tau^{prev}}^*(\alpha) \} \in \mathcal{R}_{\tau^{prev}}^\tau \right\}$
- 14 Determine the calcium concentration at τ^{prev} , $\alpha^* := \underset{\alpha \in \mathcal{R}_{\tau^{prev}}^{\tau^{cur}}}{\text{argmin}} \{ \text{Cost}_{\tau^{prev}}^*(\alpha) \}$
- 15 Update list of changepoints $cp := (\tau^{cur}, cp)$
- 16 Update list of calcium concentrations, $c := (\alpha^*, c)$
- 17 **foreach** *timestep* $s = (\tau^{prev} - 1), \dots, (\tau^{cur} + 1)$ **do**
- 18 Calculate calcium concentration, α^*/γ , and then append to list, $c := (\alpha^*/\gamma, c)$
- 19 Scale $\alpha^* := \alpha^*/\gamma$
- 20 **end**
- 21 **end**

Output : Set of changepoints cp , number of changepoints $k := \text{card}(cp)$, and estimated calcium concentrations c .

Therefore, the most recent changepoint is $\tau^*(3) = 0$. In fact, since the most recent changepoint is at timestep 0, we say that there are no changepoints.

Algorithm 1 is an instance of the class of functional pruning algorithms proposed in [Maidstone and others \(2017\)](#).

2.2.3. Computational time of functional pruning. We saw in Example 1 that Proposition 1 can lead to a recursive algorithm for solving the problem of interest (1.3). At first glance, since $\text{Cost}_s^*(\alpha)$ is piecewise quadratic with s regions (2.11), and our recursive algorithm requires computing $\text{Cost}_1^*(\alpha), \dots, \text{Cost}_T^*(\alpha)$, it appears that a total of $1 + 2 + \dots + T = \mathcal{O}(T^2)$ operations must be performed in order to deconvolve a fluorescence trace of length T . Critically, however, this is not the case. This is because, in practice, $\text{Cost}_s^*(\alpha)$ is piecewise quadratic with *substantially fewer than s regions*, as we saw in Figure 1a. To see this, recall from (2.11) that the τ th region up to timestep s is defined as $\mathcal{R}_s^\tau \equiv \left\{ \alpha : \min_{0 \leq \tau' < s} \text{Cost}_{\tau'}^*(\alpha) = \text{Cost}_s^\tau(\alpha) \right\}$. However, if \mathcal{R}_s^τ is the empty set—that is, if there is no α such that $\min_{0 \leq \tau' < s} \text{Cost}_{\tau'}^*(\alpha) = \text{Cost}_s^\tau(\alpha)$ —then $\text{Cost}_s^*(\alpha)$ is, in fact, not a function of the τ th region.




Fig. 2. Timing comparisons between three algorithms for solving (1.2) and (1.3) with $\lambda = 1$. Functional pruning approach used in Algorithm 1 and Algorithm S8.1 of the supplementary material available at *Biostatistics* online (purple), and two algorithms from Jewell and Witten (2018): one based on recursion (2.6) (Optimal partitioning), and one based on an improvement to (2.6) that makes use of ideas from Killick and others (2012) (Inequality pruning). Fifty sample datasets are simulated according to (1.1) with coefficient $\beta_0 = 0$, decay parameter $\gamma = 0.998$, normal errors $\epsilon_t \stackrel{\text{ind}}{\sim} N(0, \sigma = 0.15)$, Poisson distributed spikes $z_t \stackrel{\text{ind}}{\sim} \text{Pois}(\theta)$ where $\theta \in \{0.1, 0.01, 0.001\}$, and initial calcium value $c_1 \sim \text{Pois}(\theta)$. Standard errors are on average $< 0.1\%$ of the average computation time. Panels correspond to different values of θ . Timing results were obtained on an Intel Xeon E5-2620 2.0 GHz processor.

In practice, \mathcal{R}_s^r will often be the empty set. For instance, see Figure 1a. We note that in this example, at timestep $s = 40$, the optimal cost function is only a function of three regions,

$$\text{Cost}_{40}^*(\alpha) = \begin{cases} 1.88\alpha^2 - 0.17\alpha + 2.08, & \alpha \in \mathcal{R}_{40}^{37} \equiv [0, 0.06] \\ 142.08\alpha^2 - 39.60\alpha + 3.85, & \alpha \in \mathcal{R}_{40}^{20} \equiv [0.06, 0.22] \\ 0.50\alpha^2 - 0.10\alpha + 2.10, & \alpha \in \mathcal{R}_{40}^{39} \equiv [0.22, \infty) \end{cases}$$

In a similar way, in Example 1, we saw that $\text{Cost}_3^*(\alpha)$ was a function of two regions.

Therefore, though its worst-case performance is upper-bounded by $\mathcal{O}(T^2)$, in practice, Algorithm 1 is typically *much* faster than this. In Appendix S6 of the supplementary material available at *Biostatistics* online, we show that the maximum number of regions, $\max_{s=0, \dots, T} |\{j : \mathcal{R}_s^j \neq \emptyset, 0 \leq j \leq s-1\}|$, is a small fraction of T ; for $T = 100\,000$, fewer than 30 regions are required.

Furthermore, by slightly modifying Theorem 6.1 of Maidstone and others (2017), we can show that Algorithm 1 is no worse than the algorithm proposed in Jewell and Witten (2018). In fact, as shown in Figure 2, Algorithm 1 is typically up to a thousand times faster than that of Jewell and Witten (2018) on a fluorescence trace of length 100 000. In simulations, our C++ implementation of Algorithm 1 runs in less than one second on traces of length 100 000.

2.3. An efficient algorithm to solve the constrained problem (1.2)

As stated in the introduction, our main interest is to solve (1.2) for the global optimum. Problem (1.2) differs from problem (1.3) in that there is an additional constraint that enforces biological reality: firing neurons can only cause an increase, but not a decrease, in the calcium concentration. The algorithm in Jewell and Witten (2018) cannot be used to solve (1.2), because it relies on the recursion in (2.6), which does not allow for any dependence in the calcium concentration before and after a changepoint. Thus, at the time of this writing, there are no algorithms available to efficiently solve (1.2) for the global optimum.

In this section, we utilize a simple modification, due to Hocking and others (2017), to the functional recursion (2.10) that ensures that the constraint $c_t - \gamma c_{t-1} \geq 0$ is satisfied. First, recall from (2.10) that

$\text{Cost}_s^*(\alpha) = \min \left\{ \text{Cost}_{s-1}^*(\alpha/\gamma), \min_{\alpha'} \text{Cost}_{s-1}^*(\alpha') + \lambda \right\} + \frac{1}{2}(y_s - \alpha)^2$, where we take the minimum over two terms, which result from adding an additional point y_s to the current segment, $\text{Cost}_{s-1}^*(\alpha/\gamma) + \frac{1}{2}(y_s - \alpha)^2$, and adding a new candidate changepoint at $s - 1$ and starting a new segment at the s th timestep, $\min_{\alpha'} \text{Cost}_{s-1}^*(\alpha') + \lambda + \frac{1}{2}(y_s - \alpha)^2$.

In the latter case, if there is a spike at the s th timestep, then in order to enforce the positivity constraint, $z_s = c_s - \gamma c_{s-1} \geq 0$, the term $\min_{\alpha'} \text{Cost}_{s-1}^*(\alpha') + \lambda$ in (2.10) needs to be modified to $\min_{\alpha': \alpha \geq \alpha'} \text{Cost}_{s-1}^*(\alpha'/\gamma) + \lambda$. Therefore, we replace (2.10) with

$$\text{Cost}_s^*(\alpha) = \min \left\{ \text{Cost}_{s-1}^*(\alpha/\gamma), \min_{\alpha': \alpha \geq \alpha'} \text{Cost}_{s-1}^*(\alpha'/\gamma) + \lambda \right\} + \frac{1}{2}(y_s - \alpha)^2, \quad (2.13)$$

and we replace (2.8) with

$$\text{Cost}_s^\tau(\alpha) \equiv \min_{\alpha': \alpha' \leq \gamma^{\tau-s}\alpha} \left[\text{Cost}_\tau^*(\alpha') + \frac{1}{2} \sum_{t=\tau+1}^s (y_t - \alpha \gamma^{t-s})^2 + \lambda \right]. \quad (2.14)$$

We note that this is a slight abuse of notation since $\text{Cost}_s^*(\alpha)$ and $\text{Cost}_s^\tau(\alpha)$ take on different definitions depending on the optimization problem ((1.2) or (1.3)). Equations (2.13) and (2.14) can be used to develop an efficient recursive algorithm to solve problem (1.2). Details of the algorithm itself are included in [Appendix S2 of the supplementary material](#) available at *Biostatistics* online. A continuation of Example 1 that solves (1.2) is included in [Appendix S5 of the supplementary material](#) available at *Biostatistics* online. Figure 2 shows the running time of solving (1.2).

2.4. Solving (1.1) for non-zero intercept β_0

Thus far, we have considered (1.1) with $\beta_0 = 0$. To accommodate the possibility of nonzero baseline calcium, we consider the problem

$$\text{minimize}_{c_1, \dots, c_T, z_2, \dots, z_T, \beta_0} \left\{ \frac{1}{2} \sum_{t=1}^T (y_t - (\beta_0 + c_t))^2 + \lambda \sum_{t=2}^T 1_{(z_t \neq 0)} \right\} \text{ subject to } z_t \geq c_t - \gamma c_{t-1}. \quad (2.15)$$

Instead of directly solving (2.15) with respect to $(c_1, \dots, c_T, z_2, \dots, z_T, \beta_0)$, we consider a fine grid of values for β_0 , and solve (1.2) with $y - \beta_0$ using [Algorithm S8.1 of the supplementary material](#) available at *Biostatistics* online, for each value of β_0 considered. The solution to (2.15) is the set $\{\hat{c}_1, \dots, \hat{c}_T, \hat{z}_2, \dots, \hat{z}_T, \beta_0\}$ corresponding to the value of β_0 that led to the smallest value of the objective, over all values of β_0 considered.

2.5. Solving (1.1) with additional spike constraints

The methods used to solve (1.2) and (1.3) can also be used to solve the related nonconvex problem

$$\text{minimize}_{c_1, \dots, c_T} \left\{ \frac{1}{2} \sum_{t=1}^T (y_t - c_t)^2 \right\} \text{ subject to } c_t - \gamma c_{t-1} \geq z_{\min} \text{ or } c_t - \gamma c_{t-1} = 0, \quad (2.16)$$

proposed in [Friedrich and others \(2017\)](#). In [Appendix S4 of the supplementary material](#) available at *Biostatistics* online, we examine this proposal more closely. Remarkably, we show that [Algorithm S8.1](#)

of the [supplementary material](#) available at *Biostatistics* online can be generalized to solve

$$\underset{c_1, \dots, c_T}{\text{minimize}} \left\{ \frac{1}{2} \sum_{t=1}^T (y_t - c_t)^2 + \lambda \sum_{t=2}^T 1_{\{c_t - \gamma c_{t-1} \neq 0\}} \right\} \text{ subject to } c_t - \gamma c_{t-1} \geq z_{\min} \text{ or } c_t - \gamma c_{t-1} = 0 \quad (2.17)$$

exactly! We note that this is equivalent to (2.16) by taking $\lambda = 0$.

3. REAL DATA EXPERIMENTS

In this section, we illustrate the performance of the solution to (1.2) for spike deconvolution across a number of datasets, which were aggregated as part of the recent `spikefinder` challenge (<http://spikefinder.codeneuro.org/>). Each dataset consists of both calcium and electrophysiological recordings for a single cell. As part of the `spikefinder` challenge, all data recordings were standardized by resampling to 100 Hz and linear trends were removed from the calcium trace via preprocessing steps described in [Theis and others \(2016\)](#).

Throughout this section, due to computation considerations, the solutions to (1.2) and (1.3) are obtained using slight modifications of [Algorithm 1](#) and [Algorithm S8.1](#) of the [supplementary material](#) available at *Biostatistics* online. These modifications are described in [Appendix S3.1](#) of the [supplementary material](#) available at *Biostatistics* online. Additionally, since the `spikefinder` data removed linear trends from the raw calcium trace, we do not estimate β_0 in (1.1). Instead, we set $\beta_0 = 0$; our empirical results suggest that estimation of β_0 may not be necessary.

In our experiments, we will treat the spikes ascertained using electrophysiological recording as the “ground truth”, and will quantify the ability of spike deconvolution algorithms to recover these ground truth spikes on the basis of the calcium recordings. The data sets differ in terms of the choice of calcium indicator (GCaMP5, GCaMP6, jRCaMP, jRGECO, OGB), scanning technology (AOD, galvo, and resonant), and circuit under investigation (V1 and retina).

Throughout this section, we compare our proposal (1.2) to a recent approach from the literature that employs an ℓ_1 (convex) relaxation to (1.2),

$$\underset{c_1, \dots, c_T, z_2, \dots, z_T}{\text{minimize}} \left\{ \frac{1}{2} \sum_{t=1}^T (y_t - c_t)^2 + \lambda |c_1| + \lambda \sum_{t=2}^T |z_t| \right\} \text{ subject to } z_t = c_t - \gamma c_{t-1} \geq 0, \quad (3.18)$$

proposed by [Friedrich and Paninski \(2016\)](#) and [Friedrich and others \(2017\)](#). [Friedrich and others \(2017\)](#) developed a very fast algorithm to solve (3.18); in simulated examples their algorithm solves (3.18) approximately 40–60× faster than [Algorithm 1](#) and 40–900× faster than [Algorithm S8.1](#) of the [supplementary material](#) available at *Biostatistics* online. This is not surprising, since (3.18) is a convex problem, whereas (1.2) and (1.3) are nonconvex problems. Moreover, in practical applications, [Algorithm 1](#) and [Algorithm S8.1](#) of the [supplementary material](#) available at *Biostatistics* online are often fast enough. Indeed, [de Vries and others \(2018\)](#) uses [Algorithm S8.1](#) of the [supplementary material](#) available at *Biostatistics* online to deconvolve traces from nearly 60 000 neurons.

Since the solution to (3.18) often results in many small non-zero elements of \hat{z}_t , we consider post-thresholding. That is, given $\hat{z}_2, \dots, \hat{z}_T$ that solve (3.18), and a threshold $L > 0$, we set $\tilde{z}_t = \hat{z}_t 1_{\{\hat{z}_t > L\}}$; in other words, we conclude that a spike is present only if $\hat{z}_t > L$.

In [Section 3.1](#), we compare our proposed approach (1.2) to (3.18) on data from the `spikefinder` challenge. We describe our experimental approach in [Section 3.1.1](#). [Section 3.1.2](#) illustrates these methods for a single cell, and in [Section 3.1.3](#), we examine results for all datasets considered in the `spikefinder`

challenge. In Section 3.2, we illustrate on a real-data example that solving (1.2) gives superior estimates than solving (1.3). In Section 3.3, we compare the estimated increase in calcium due to a spike (using (1.2)) to the actual number of recorded spikes (based on the ground truth electrophysiological recordings).

R code to reproduce all experiments is available on GitHub at <https://github.com/jewellsean/fast-nonconvex-experiments>.

3.1. Comparison of (1.2) to (3.18) on data from the *spikefinder* challenge

3.1.1. *Description of methods for Sections 3.1.2–3.1.3.* We now describe the methods that will be used in Sections 3.1.2–3.1.3. Our main objective is to accurately estimate the times at which spikes occur. Thus, we use two measures that directly compare two spike trains, both of which have been used extensively in the neuroscience literature (Quiroga and Panzeri, 2009; Reinagel and Reid, 2000; Gerstner and others, 2014): (i) van Rossum distance with timescale parameter $\tau = 0.1$ (van Rossum, 2001; Houghton and Kreuz, 2012) and (ii) Victor-Purpura distance with cost parameter 10 (Victor and Purpura, 1997, 1996). We also use an additional measure: (iii) the correlation between two downsampled spike trains; details of this measure are provided in Theis and others (2016). As we will see, measures (i) and (ii) are sensitive to the timing of spikes, whereas measure (iii) is somewhat insensitive to the timing of the spikes, and instead quantifies the similarity between the spike rates.

To analyze the performances of the proposals (1.2) and (3.18) over a single fluorescence trace, we take a training/test set approach. Given a fluorescence trace of length T , the first $\lfloor T/2 \rfloor$ timesteps are used in the training set, and the remainder are used for the test set. We solve (1.2) and (3.18) for a range of values of the tuning parameter λ on the training set; in the case of (3.18) we also use a range of threshold values L .

For all tuning parameter values considered, we apply the three measures mentioned earlier to the estimated and true spike trains, and select the tuning parameter values that optimize these measures on the training set. We then solve (1.2) and (3.18) on the test set with the selected values of the tuning parameters, and evaluate test set performance.

As pointed out by Pachitariu and others (2018), estimating the decay rate γ in (1.1) is difficult. Therefore, as in Pachitariu and others (2018), we categorize calcium indicators into three groups based on their decay properties. As in Vogelstein and others (2010), within each calcium indicator rate category, we set $\gamma = 1 - \frac{\Delta}{\phi}$, where Δ is 1 / (frame rate), and ϕ is a time-scale parameter based on the category, defined as

$$\phi = \begin{cases} 0.7, & \text{fast category} \\ 1.25, & \text{medium category} \\ 2, & \text{slow category} \end{cases}.$$

For example, in Figure 3, GCaMP6f is classified as a fast indicator and the data is recorded at 100Hz. Therefore, we take $\gamma \approx 0.986$.

In practice, users typically do not have the benefit of a training set to select the tuning parameter value λ to solve (1.2) or (1.3). Therefore, we recommend using the procedure proposed in de Vries and others (2018), which selects λ based on the firing rate, decay rate γ , and estimated signal-to-noise ratio.

3.1.2. *Results for a single cell.* In Figure 3, we illustrate this procedure for cell 13, GCaMP6f, V1, from Chen and others (2013). Each row corresponds to one of the measures described in Section 3.1.1. The left column displays these measures on the training set, for the solution to (1.2) with different values of λ , and for the post-thresholded solution to (3.18) with different values of λ and L . The right column shows

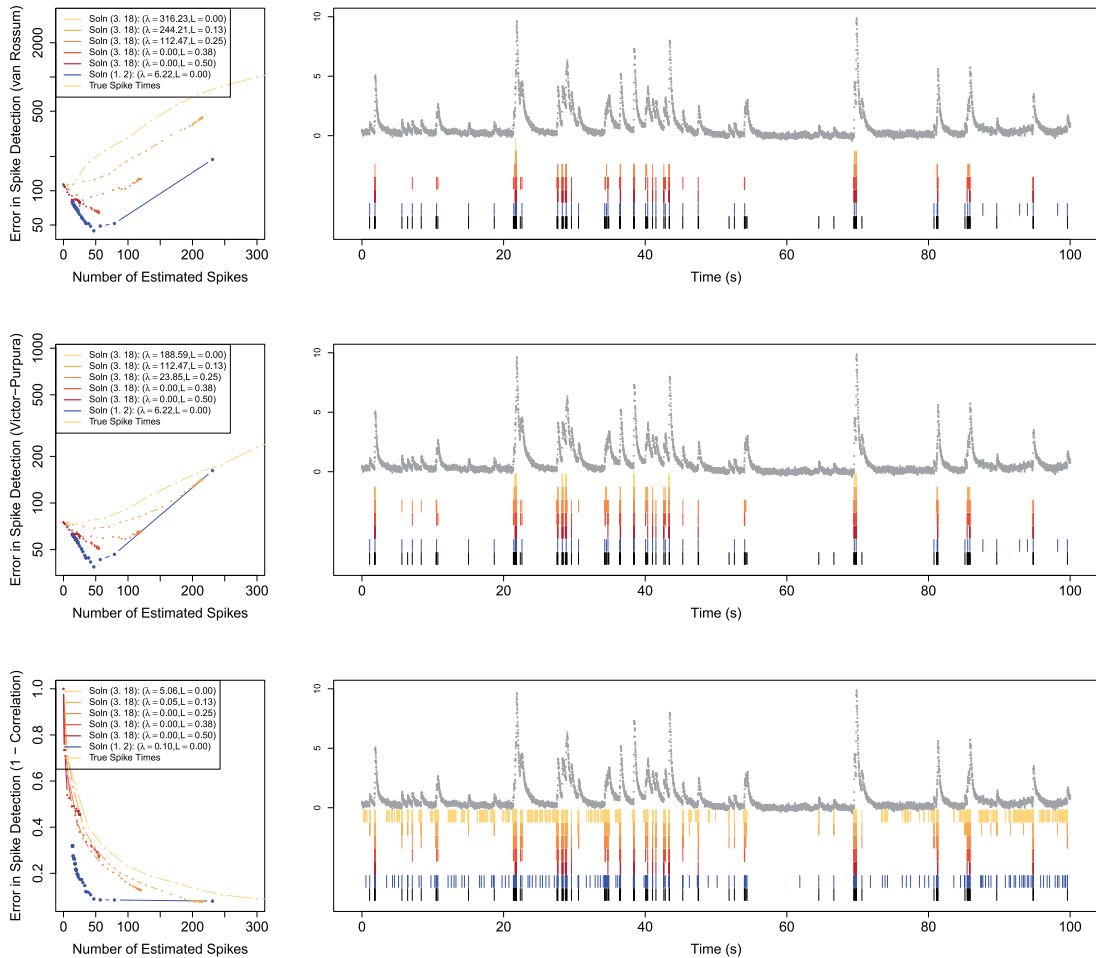


Fig. 3. Illustrative example for cell 13, GCaMP6f, V1, from [Chen and others \(2013\)](#) after preprocessing; see [Theis and others \(2016\)](#). Different spike measures are displayed in each row. *Left*: Performances of the post-thresholded solution to (3.18) and the solution to (1.2). *Right*: The cell’s fluorescence trace is displayed in grey. The estimated spikes on the test set from the “best” choice of the tuning parameter λ , as determined by either van Rossum, Victor-Purpura, or a correlation-based measure on the training set, are displayed under the fluorescence trace. The true spike times, as determined using electrophysiological recording, are shown in black. The colors in the left-hand panels correspond to the colors in the right-hand panel.

the fluorescence trace along with the estimated spikes, on the test set, using tuning parameters selected on the training set.

There are a number of important observations to draw from Figure 3. As measured by van Rossum and Victor-Purpura, the estimated spikes from (1.2) are much more accurate than those estimated (and post-thresholded) using the convex relaxation (3.18). This agrees with our visual inspection of the right hand panel: the estimated spikes from problem (1.2) more closely match the number and timings of the true spikes than those estimated from problem (3.18).

In contrast, if performance is measured by correlation, then the estimated spikes obtained from (3.18) result in slightly better performance than the estimated spikes from (1.2). However, in the training set

there are 75 true spikes, whereas (3.18) outperforms (1.2) when approximately 200 spikes are estimated. Therefore, selecting the tuning parameter for (3.18) based on correlation leads to a *substantial overestimate* of the number of spikes, and therefore, poor overall accuracy in the number and timing of the spikes. This pattern has been observed in other ℓ_1 regularization problems (Zou, 2006; Maidstone and others, 2018), and persists across cells in the `spikefinder` data (results not shown).

To summarize, when van Rossum and Victor-Purpura distance are used to evaluate performance, our proposal (1.2) substantially outperforms the approach in (3.18). When performance is evaluated using correlation, the performance of (3.18) is slightly better than that of (1.2); however, this better performance is achieved when far too many spikes are estimated, indicating that correlation is a poor choice for quantifying the accuracy of spike detection.

3.1.3. Results for all datasets in the `spikefinder` challenge. In this section, we examine the performance of the solutions to (1.2) and (3.18) on all datasets collected as part of the `spikefinder` challenge. For the ten datasets included in this challenge, [Table S1 of the supplementary material](#) available at *Biostatistics* online tabulates the calcium indicator; circuit; publishing authors; average, minimum, and maximum fluorescence trace length; the number of cells measured; and the time-scale classification. In total, there are 174 traces, each of which contains fewer than 100,000 timesteps. We analyze these 174 cells as described in Section 3.1.1.

Figure 4a compares the test set performance, with respect to the van Rossum, Victor-Purpura, and correlation measures, for each of the 174 cells. As measured by the van Rossum and Victor-Purpura distance, the solution to (1.2) outperforms the solution to (3.18). However, under the correlation measure, the solution to (3.18) achieves higher correlations than the solution to (1.2). These results are consistent with those on a single cell presented in Section 3.1.2, where it was shown that van Rossum and Victor-Purpura accurately estimate spike times, whereas correlation yields a cruder measure of spike rate and encourages overestimation of the number of spikes.

3.2. The solution to (1.2) outperforms the solution to (1.3)

As mentioned earlier, in this article, we have developed not only an algorithm for solving (1.3) that is much faster than the algorithm proposed in Jewell and Witten (2018), but also an algorithm for solving (1.2), which cannot be solved using techniques from Jewell and Witten (2018). By incorporating the fact that a firing neuron causes an increase, but never a decrease, in the calcium concentration, the estimated spikes from problem (1.2) are closer to the ground truth spikes than the estimated spikes from (1.3). In practice, the solutions to (1.2) and (1.3) are typically quite similar; however, the solution to (1.2) benefits from greater interpretability. See [Appendix S8 of the supplementary material](#) available at *Biostatistics* online for an example.

3.3. Comparison of the estimated spike magnitudes from (1.2) to the true number of spikes.

The data from the `spikefinder` challenge was resampled to 100 Hz before we downloaded it. At this sampling frequency, since one timestep is just 1/100th of a second, there are very few timesteps with more than one true spike. Nonetheless, for instances where there is more than one spike in a single timestep, we wish to ask the question: Do larger values of the estimated spike magnitudes, $\hat{c}_t - \gamma\hat{c}_{t-1}$, correspond to more true spikes (as measured by electrophysiology) in the t th timestep?

Figure 4b investigates whether there is a relationship between the estimated spike magnitude $\hat{c}_t - \gamma\hat{c}_{t-1}$ and the number of spikes measured by electrophysiology at the t th timestep. Because the estimated spike magnitude of $\hat{c}_t - \gamma\hat{c}_{t-1}$ is not directly comparable across cells, for each cell we transform the magnitudes into percentiles. We then compare the percentile of $\hat{c}_t - \gamma\hat{c}_{t-1}$ to the true number of spikes within a 0.1

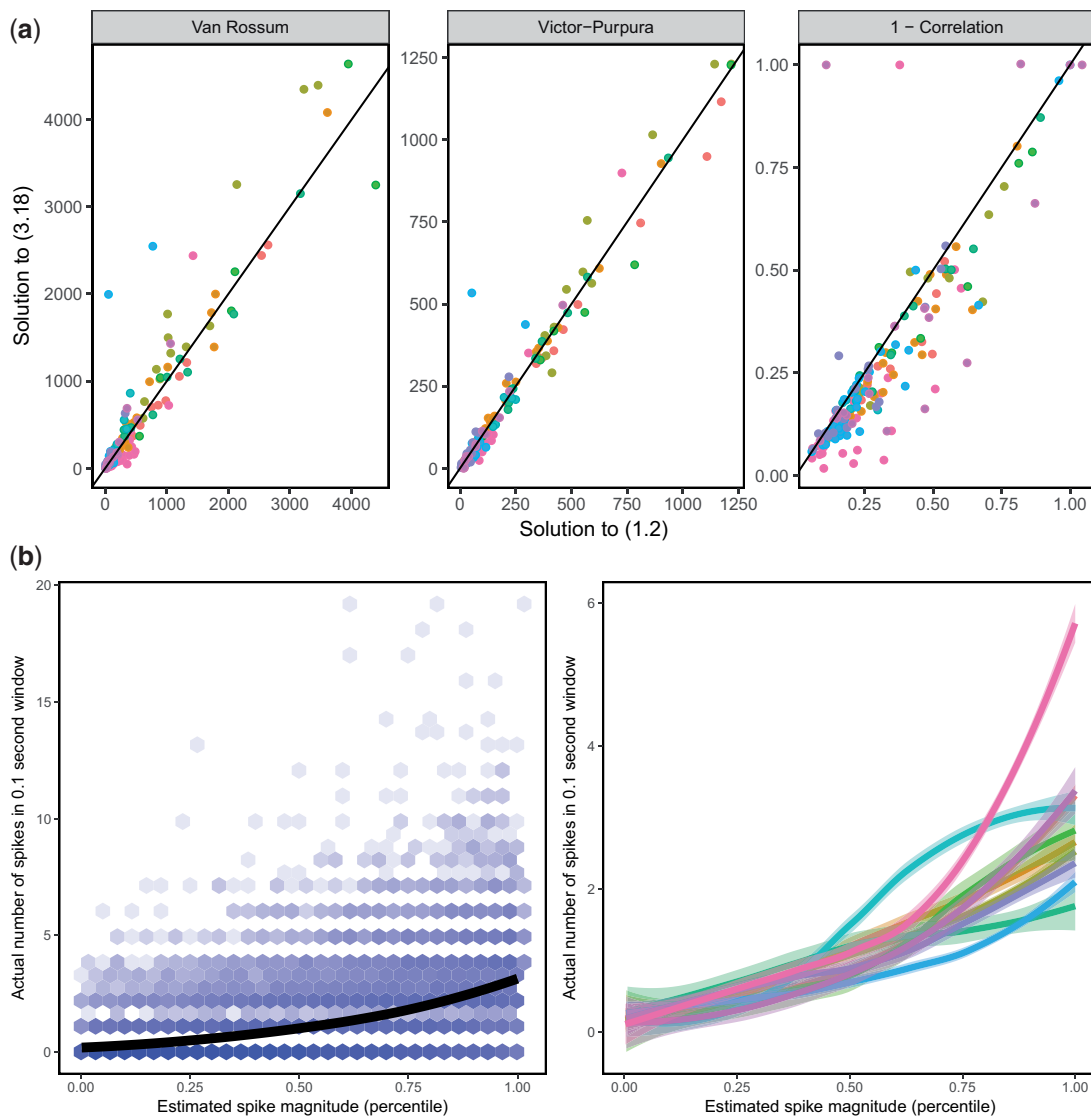


Fig. 4. Results for all datasets in the spikefinder challenge. (a) Optimal van Rossum, Victor-Purpura, and correlation measures for our proposal, (1.2), and a competing proposal, (3.18). Small values of the van Rossum and Victor-Purpura measures suggest accurate estimation of the timing and number of spikes, whereas a large value of the correlation measure suggests accurate estimation of the spike rate, though perhaps an overestimate of the number of spikes. Each dot represents the performance of (1.2) and (3.18) on a single cell, for each of the 174 cells. Cells are colored based on the dataset from which they were obtained (see Table S1 of the supplementary material available at *Biostatistics* online). (b) Large increases in the estimated spike magnitude, $\hat{c}_t - \gamma\hat{c}_{t-1}$, are associated with more true spikes, as measured by electrophysiology, at the t th timestep. For each cell, we transform the spike magnitudes into percentiles, and then compare the percentile of $\hat{c}_t - \gamma\hat{c}_{t-1}$ to the true number of spikes within a 0.1 second window of t . *Left*: For each cell in each of the ten datasets, we display each timestep for which a spike is estimated to occur; however, to avoid overplotting, hexagonal bins are used to represent points covered by the hexagon; darker colors indicate more points. The black curve represents the loess fit across all of the points. *Right*: Loess curves along with 95% confidence intervals for each dataset. Cells are colored based on the dataset from which they were obtained (see Table S1 of the supplementary material available at *Biostatistics* online). Details are provided in Section 3.3.

second window of t . Figure 4b displays the percentiles and the number of spikes across all 174 traces on a test set; tuning parameters were chosen to optimize the van Rossum distance on a training set. The left panel displays a loess curve fit to all ten datasets, and the right panel shows the loess curves along with 95% confidence intervals for each dataset. As expected, a larger value of $\hat{c}_t - \gamma \hat{c}_{t-1}$ is associated with more spikes in the ground truth data.

4. DISCUSSION

Determining the times at which a neuron fires from a calcium imaging dataset is a challenging and important problem. In this article, we build upon the nonconvex approach for spike deconvolution proposed in Jewell and Witten (2018). Though Jewell and Witten (2018) proposed a tractable algorithm for solving the nonconvex problem, it is prohibitively slow to run on large populations of neurons for which long recordings are available. The algorithm proposed in this paper solves the optimization problem of Jewell and Witten (2018) for fluorescence traces of 100 000 timesteps in less than a second. Moreover, Algorithm S8.1 of the supplementary material available at *Biostatistics* online overcomes a limitation of Jewell and Witten (2018) by avoiding “negative” spikes; that is, a decrease in the calcium concentration due to a spike. We show that these algorithms have excellent performance, relative to existing convex relaxations, as quantified by the van Rossum and Victor-Purpura measures, on datasets collected as part of the `spikefinder` challenge (<http://spikefinder.codeneuro.org/>). Moreover, Algorithm S8.1 of the supplementary material available at *Biostatistics* online was recently used to decode data from nearly 60 000 neurons in the Allen Institute for Brain Science’s “platform paper” for the Allen Brain Observatory (de Vries and others, 2018).

In this article, we assume that the calcium concentration decays exponentially according to a first-order auto-regressive model. Although this is typically a good approximation, there are datasets for which—due to different experimental conditions—spike times estimated from (1.2) and (1.3) are systematically biased due to model misspecification. In future work, we propose to extend the functional pruning framework to more general calcium models.

In this article, we focus on developing *point estimates* of the times at which a neuron spikes. However, it is also of interest to propagate uncertainty from the deconvolution procedure to downstream analyses that rely on the spike times. It remains an open question to define the notion of confidence associated with a set of estimated spikes.

SUPPLEMENTARY MATERIAL

Supplementary material is available at <http://biostatistics.oxfordjournals.org>.

ACKNOWLEDGMENTS

We thank Michael Buice, Peter Ledochowitsch, and Michael Oliver at the Allen Institute for Brain Science and Ilana Witten at Princeton for helpful conversations. *Conflict of Interest:* None declared.

FUNDING

S.J. received funding from the Natural Sciences and Engineering Research Council of Canada. T.D.H. is partially supported by the Natural Sciences and Engineering Research Council of Canada (RGPGR 448167-2013), and by Canadian Institutes of Health Research (EP1-120608 and EP1-120609). This work was partially supported by Engineering and Physical Sciences Research Council Grant EP/N031938/1 to P.F., and National Science Foundation CAREER DMS-1252624, National Institutes of Health grants DP5OD009145, R01DA047869, and R01EB026908, and a Simons Investigator Award in Mathematical Modeling of Living Systems to D.W.

REFERENCES

- AHRENS, M. B., ORGER, M. B., ROBSON, D. N., LI, J. M. AND KELLER, P. J. (2013). Whole-brain functional imaging at cellular resolution using light-sheet microscopy. *Nature Methods* **10**, 413–420.
- CHEN, T.-W., WARDILL, T. J., SUN, Y., PULVER, S. R., RENNINGER, S. L., BAOHAN, A., SCHREITER, E. R., KERR, R. A., ORGER, M. B., JAYARAMAN, V. *and others*. (2013). Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature* **499**, 295–300.
- DE VRIES, S. E. J., LECOQ, J., BUICE, M. A., GROBLEWSKI, P. A., OCKER, G. K., OLIVER, M., FENG, D., CAIN, N., LEDOCHOWITSCH, P., MILLMAN, D. *and others*. (2018). A large-scale, standardized physiological survey reveals higher order coding throughout the mouse visual cortex. *bioRxiv:359513*. doi: <https://doi.org/10.1101/359513>.
- DENEUX, T., KASZAS, A., SZALAY, G., KATONA, G., LAKNER, T., GRINVALD, A., RÓZSA, B. AND VANZETTA, I. (2016). Accurate spike estimation from noisy calcium signals for ultrafast three-dimensional imaging of large neuronal populations *in vivo*. *Nature Communications* **7**, 12190.
- DOMBECK, D. A., KHABBAZ, A. N., COLLMAN, F., ADELMAN, T. L. AND TANK, D. W. (2007). Imaging large-scale neural activity with cellular resolution in awake, mobile mice. *Neuron* **56**, 43–57.
- DYER, E. L., DUARTE, M. F., JOHNSON, D. H. AND BARANIUK, R. G. (2010). Recovering spikes from noisy neuronal calcium signals via structured sparse approximation. In: Vigneron, V., Zarzoso, V., Moreau, E., Gribonval, R. and Vincent, E. (editors), *Latent Variable Analysis and Signal Separation*. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 604–611.
- DYER, E. L., STUDER, C., ROBINSON, J. T. AND BARANIUK, R. G. (2013). A robust and efficient method to recover neural events from noisy and corrupted data. In: *Neural Engineering (NER), 2013 6th International IEEE/EMBS Conference on IEEE*. San Diego, CA, USA: IEEE, pp. 593–596.
- FRIEDRICH, J. AND PANINSKI, L. (2016). Fast active set methods for online spike inference from calcium imaging. In: Lee, D.D., Sugiyama, M., Luxburg, U.V., Guyon, I. and Garnett, R. (editors), *Advances In Neural Information Processing Systems*. Barcelona, Spain: Curran Associates, Inc., pp. 1984–1992.
- FRIEDRICH, J., ZHOU, P. AND PANINSKI, L. (2017). Fast online deconvolution of calcium imaging data. *PLoS Computational Biology* **13**, e1005423.
- GERSTNER, W., KISTLER, W. M., NAUD, R. AND PANINSKI, L. (2014). *Neuronal Dynamics: From Single Neurons to Networks and Models of Cognition*. Cambridge: Cambridge University Press.
- GREWE, B. F., LANGER, D., KASPER, H., KAMPA, B. M. AND HELMCHEN, F. (2010). High-speed *in vivo* calcium imaging reveals neuronal network activity with near-millisecond precision. *Nature Methods* **7**, 399–405.
- HAYNES, K., ECKLEY, I. A. AND FEARNHEAD, P. (2017). Computationally efficient changepoint detection for a range of penalties. *Journal of Computational and Graphical Statistics* **26**, 134–143.
- HOCKING, T. D., RIGAILL, G., FEARNHEAD, P. AND BOURQUE, G. (2017). A log-linear time algorithm for constrained changepoint detection. *arXiv preprint arXiv:1703.03352*.
- HOLEKAMP, T. F., TURAGA, D. AND HOLY, T. E. (2008). Fast three-dimensional fluorescence imaging of activity in neural populations by objective-coupled planar illumination microscopy. *Neuron* **57**, 661–672.
- HOUGHTON, C. AND KREUZ, T. (2012). On the efficient calculation of van rossum distances. *Network: Computation in Neural Systems* **23**, 48–58.
- JACKSON, B., SCARGLE, J. D., BARNES, D., ARABHI, S., ALT, A., GIOUMOUSIS, P., GWIN, E., SANGTRAKULCHAROEN, P., TAN, L. AND TSAI, T. T. (2005). An algorithm for optimal partitioning of data on an interval. *IEEE Signal Processing Letters* **12**, 105–108.
- JEWELL, S. AND WITTEN, D. (2018). Exact spike train inference via ℓ_0 optimization. *The Annals of Applied Statistics* **12**, 2457–2482.

- KILLICK, R., FEARNHEAD, P. AND ECKLEY, I. A. (2012). Optimal detection of changepoints with a linear computational cost. *Journal of the American Statistical Association* **107**, 1590–1598.
- MAIDSTONE, R., FEARNHEAD, P. AND LETCHFORD, A. (2018). Detecting changes in slope with an l_0 penalty. *Journal of Computational and Graphical Statistics*, doi: 10.1080/10618600.2018.1512868.
- MAIDSTONE, R., HOCKING, T., RIGAILL, G. AND FEARNHEAD, P. (2017). On optimal multiple changepoint algorithms for large data. *Statistics and Computing* **27**, 519–533.
- PACHITARIU, M., STRINGER, C. AND HARRIS, K. D. (2018). Robustness of spike deconvolution for calcium imaging of neural spiking. *Journal of Neuroscience* **38**, 7976–7985.
- PNEVMATIKAKIS, E. A., MEREL, J., PAKMAN, A. AND PANINSKI, L. (2013). Bayesian spike inference from calcium imaging data. In: *Signals, Systems and Computers, 2013 Asilomar Conference on IEEE*. pp. 349–353. doi: 10.1109/ACSSC.2013.6810293.
- PREVEDEL, R., YOON, Y.-G., HOFFMANN, M., PAK, N., WETZSTEIN, G., KATO, S., SCHRÖDEL, T., RASKAR, R., ZIMMER, M., BOYDEN, E. S. and others. (2014). Simultaneous whole-animal 3D imaging of neuronal activity using light-field microscopy. *Nature Methods* **11**, 727–730.
- QUIROGA, R. Q. AND PANZERI, S. (2009). Extracting information from neuronal populations: information theory and decoding approaches. *Nature Reviews Neuroscience* **10**, 173.
- REINAGEL, P. AND REID, R. C. (2000). Temporal coding of visual information in the thalamus. *Journal of Neuroscience* **20**, 5392–5400.
- RIGAILL, G. (2015). A pruned dynamic programming algorithm to recover the best segmentations with 1 to k_{\max} change-points. *Journal de la Société Française de Statistique* **156**, 180–205.
- SASAKI, T., TAKAHASHI, N., MATSUKI, N. AND IKEGAYA, Y. (2008). Fast and accurate detection of action potentials from somatic calcium fluctuations. *Journal of Neurophysiology* **100**, 1668–1676.
- THEIS, L., BERENS, P., FROUDARAKIS, E., REIMER, J., ROSÓN, M. R., BADEN, T., EULER, T., TOLIAS, A. S. AND BETHGE, M. (2016). Benchmarking spike rate inference in population calcium imaging. *Neuron* **90**, 471–482.
- VAN ROSSUM, M. C. (2001). A novel spike distance. *Neural Computation* **13**, 751–763.
- VICTOR, J. D. AND PURPURA, K. P. (1996). Nature and precision of temporal coding in visual cortex: a metric-space analysis. *Journal of Neurophysiology* **76**, 1310–1326.
- VICTOR, J. D. AND PURPURA, K. P. (1997). Metric-space analysis of spike trains: theory, algorithms and application. *Network: Computation in Neural Systems* **8**, 127–164.
- VLADIMIROV, N., MU, Y., KAWASHIMA, T., BENNETT, D. V., YANG, C.-T., LOOGER, L. L., KELLER, P. J., FREEMAN, J. AND AHRENS, M. B. (2014). Light-sheet functional imaging in fictively behaving zebrafish. *Nature Methods* **11**, 883.
- VOGELSTEIN, J. T., PACKER, A. M., MACHADO, T. A., SIPPY, T., BABADI, B., YUSTE, R. AND PANINSKI, L. (2010). Fast nonnegative deconvolution for spike train inference from population calcium imaging. *Journal of Neurophysiology* **104**, 3691–3704.
- VOGELSTEIN, J. T., WATSON, B. O., PACKER, A. M., YUSTE, R., JEDYNAK, B. AND PANINSKI, L. (2009). Spike inference from calcium imaging using sequential monte carlo methods. *Biophysical Journal* **97**, 636–655.
- YAKSI, E. AND FRIEDRICH, R. W. (2006). Reconstruction of firing rate changes across neuronal populations by temporally deconvolved ca^{2+} imaging. *Nature Methods* **3**, 377–383.
- ZOU, H. (2006). The adaptive lasso and its oracle properties. *Journal of the American Statistical Association* **101**, 1418–1429.

[Received March 28, 2018; revised December 3, 2018; accepted for publication December 10, 2018]