Spike inference from calcium imaging using sequential Monte Carlo methods

Joshua Vogelstein[1] and Liam Paninski[2]

[1]joshuav@jhu.edu, Johns Hopkins University
[2]liam@stat.columbia.edu, Columbia University

Reference: Vogelstein JT, Watson BO, Packer AM, Yuste R, Jedynak B, Paninski L. Spike inference from calcium imaging using sequential Monte Carlo methods. Biophysical Journal. *in press.*

Support for JTV was provided by NIDCD DC00109.

LP is supported by an NSF CAREER award and by an Alfred P. Sloan Research Fellowship.

background

- the neural signal of interest is a spike train, i.e., a sequence of binary events, $n_t \in \{0, 1\} \forall t \in (0, T)$
- the observed signals are nonlinear and non-Gaussian functions of the spike trains
- often, the relationship between some external covariates (e.g., a movie) and the resulting spike train is of interest
- one could simultaneously observe an ensemble of neurons using new imaging technologies
- given ensemble spike trains, one would like to learn the connection matrix governing activity
- learning the connection matrix of ensembles of neurons has remained elusive, as neither the experimental technology nor the analytical tools were available. . . until now

definition of terms

States	
F_t	fluorescence
$[\mathrm{Ca}^{2+}]_t$	intracellular calcium concentration
n_t	spike
Parameters	
α	scale
β	offset
σ_F	measurement noise scale
a	"decay" of calcium
A	jump size due to spike
d	baseline of calcium
σ_c	calcium noise scale
λ	probability of spiking
Other	
$S(\cdot)$	Hill Equation: $S(x) = x^m/(x^m + k_d)$
$\varepsilon_{\cdot,t}$	standard normal random variable
Δ	time step size
$\mathcal{B}(n_t; \lambda)$	Bernoulli random variable, $n_t = 1$ w.p. λ and 0 o.w.
T	total number of steps

a simple model

$$\begin{aligned}
F_t &= \alpha S([\mathrm{Ca}^{2+}]_t) + \beta + (S([\mathrm{Ca}^{2+}]_t) + \sigma_F)\varepsilon_{F,t} \\
[\mathrm{Ca}^{2+}]_t &= a[\mathrm{Ca}^{2+}]_{t-1} + A n_t + d + \sigma_c \sqrt{\Delta}\varepsilon_{c,t} \\
n_t &\sim \mathcal{B}(n_t; \lambda\Delta)
\end{aligned}$$

a simple schematic

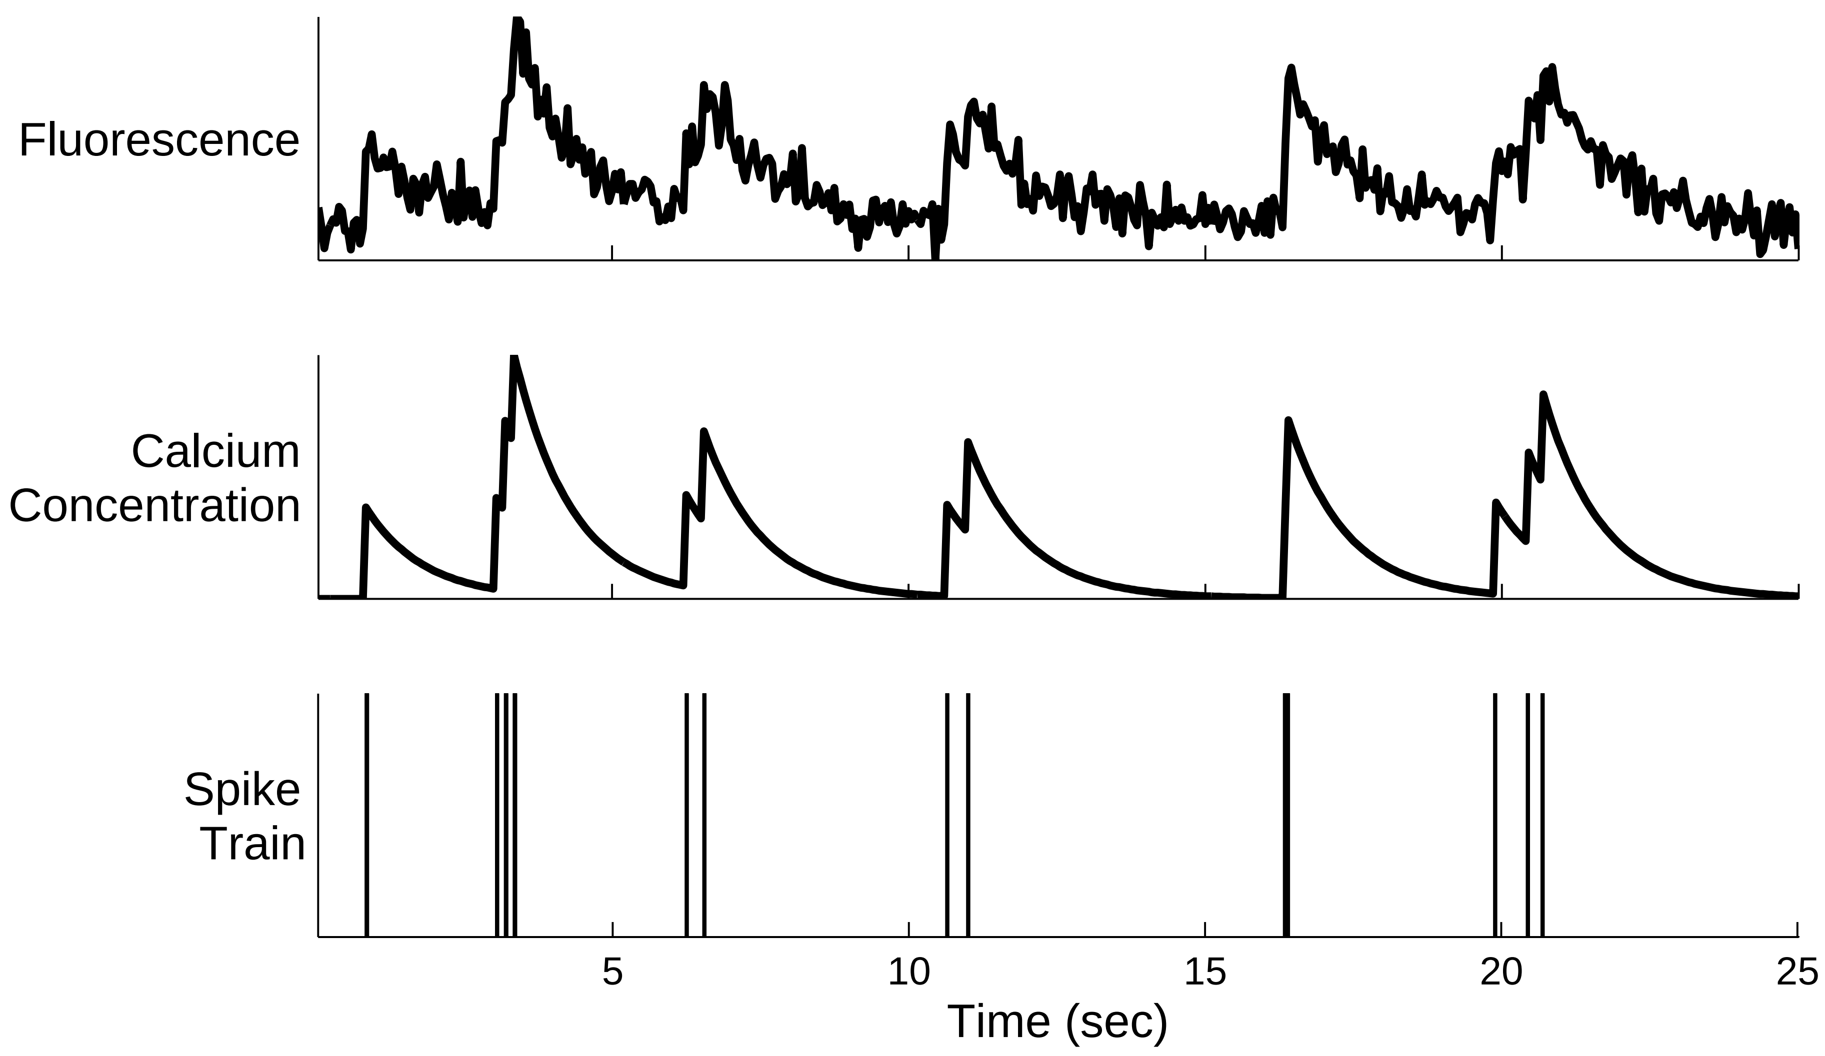

goals and approach

- Given the above model, we would like to find $P(n_t|F_{0:T},\theta) \quad \forall t \in (0,T)$ and

$$\widetilde{\theta} = \underset{\theta}{\operatorname{argmax}} \int P(n_{0:T}|F_{0:T},\theta') \ln P(n_{0:T},F_{0:T}|\theta) dn_{0:T}$$

  where $\theta = \{\alpha, \beta, \sigma_F, a, d, A, \sigma_c, \lambda\}$

- we use a forward-backward smoother to estimate the E-step of an EM algorithm, and gradient ascent to maximize all the parameters in the M-step.

- we develop an optimal one observation ahead sampler, $P_\theta(\{n, [\mathrm{Ca}^{2+}]\}_t | \{n, [\mathrm{Ca}^{2+}]\}_{t-1}, F_t)$ to sample efficiently.

particle filter outperforms optimal linear filter in simulations

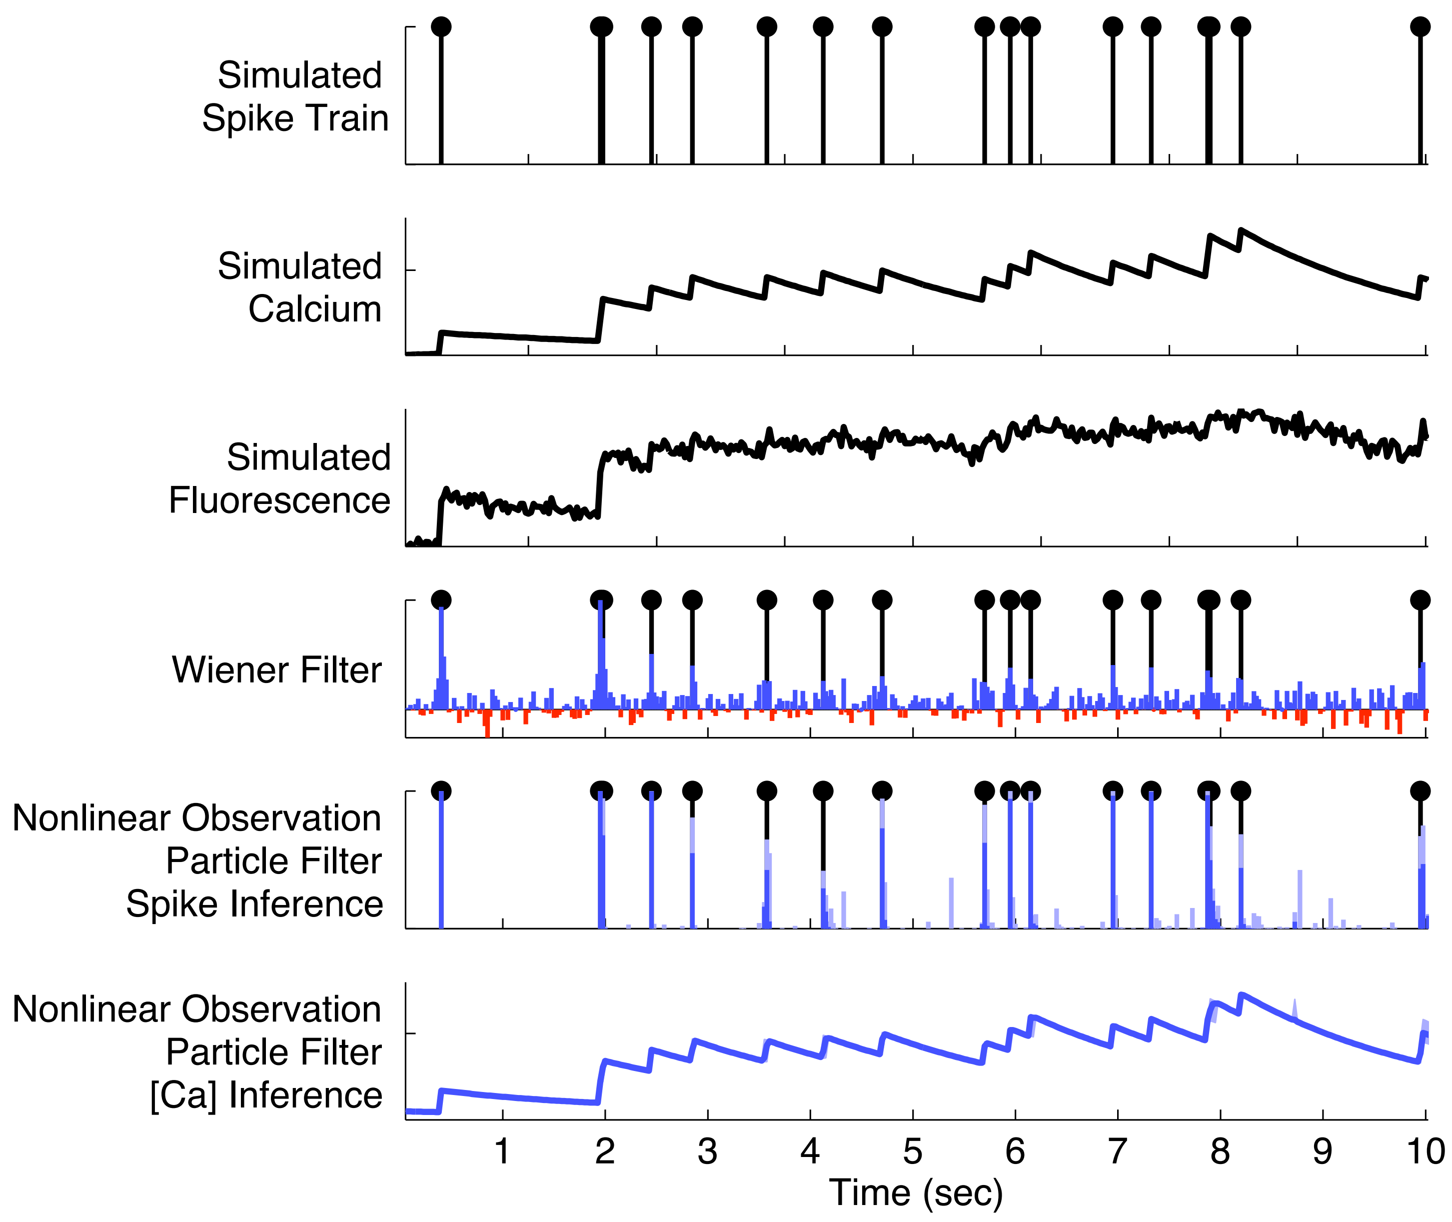

particle filter outperforms optimal linear filter in real data with ground truth

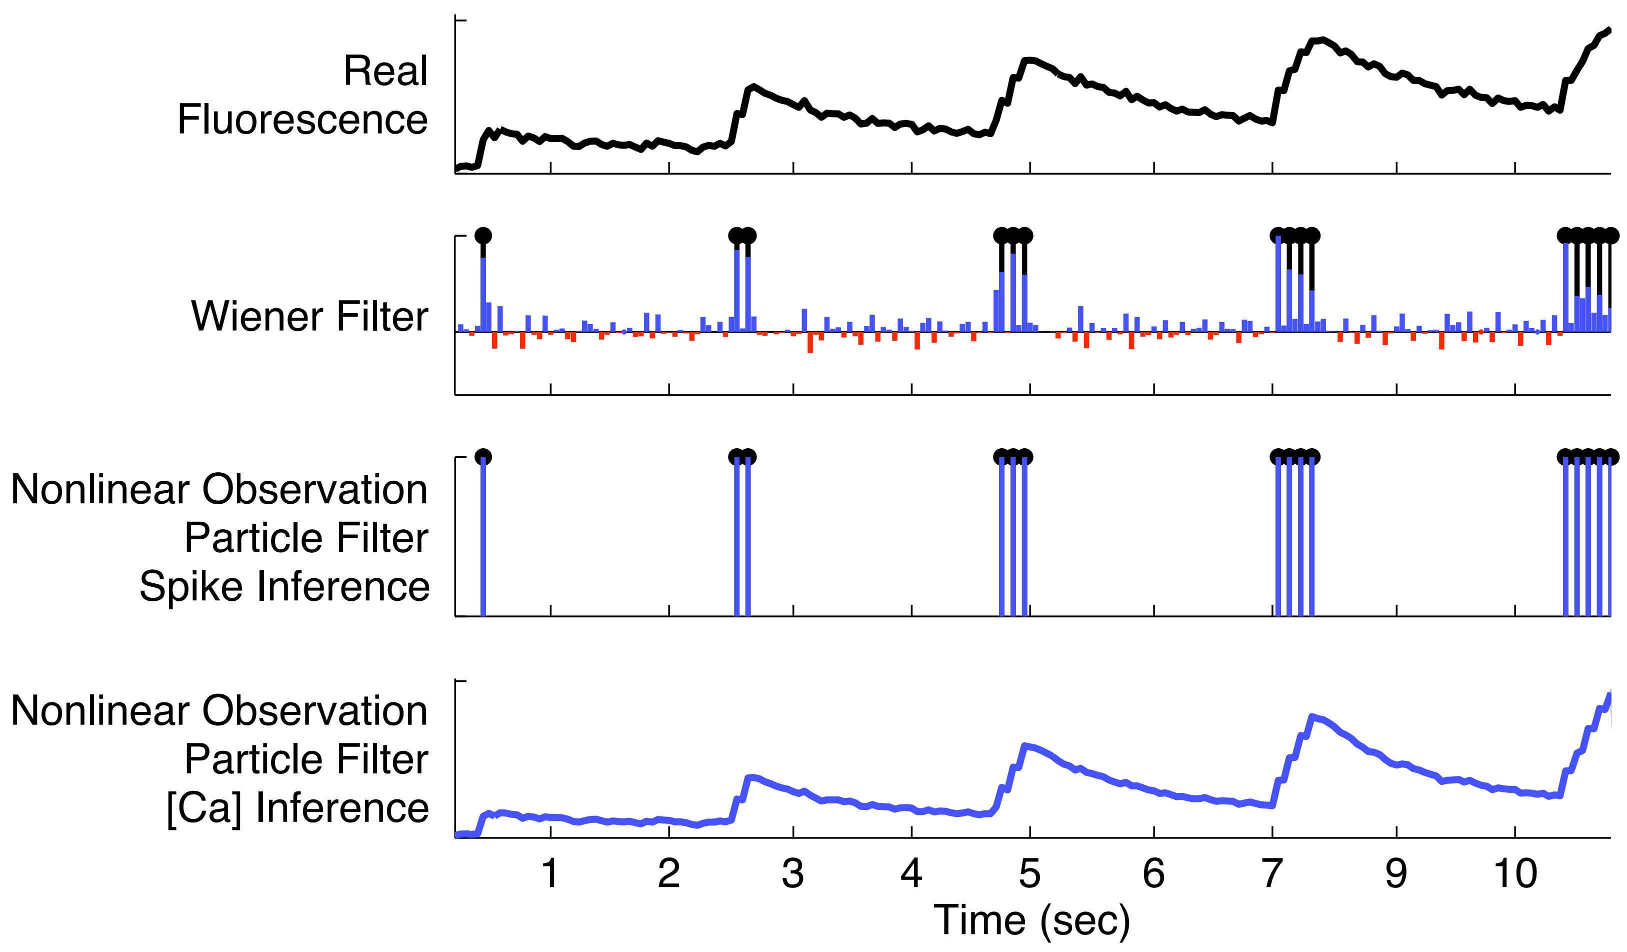

a less simple model: intermittent observations, parametric neural model

$$
\begin{aligned}
F_t &= \alpha S([\mathrm{Ca}^{2+}]_t) + \beta + (S([\mathrm{Ca}^{2+}]_t) + \sigma_F)\varepsilon_{F,t}, \quad t \in \mathcal{T}_o \subseteq (0,T) \\
[\mathrm{Ca}^{2+}]_t &= a[\mathrm{Ca}^{2+}]_{t-1} + An_t + d + \sigma_c\sqrt{\Delta}\varepsilon_{c,t} \\
n_t &\sim \mathcal{B}\big(n_t; f(\boldsymbol{k}'\boldsymbol{x}_t)\big)
\end{aligned}
$$

- observations occur at a subset of time steps
- $f(\cdot)$ is a link function that is both convex and log-concave
- $\boldsymbol{k}$ is a linear filter
- $\boldsymbol{x}_t$ is the time-varying input to the neuron, including external covariates and spike histories

a less simple schematic

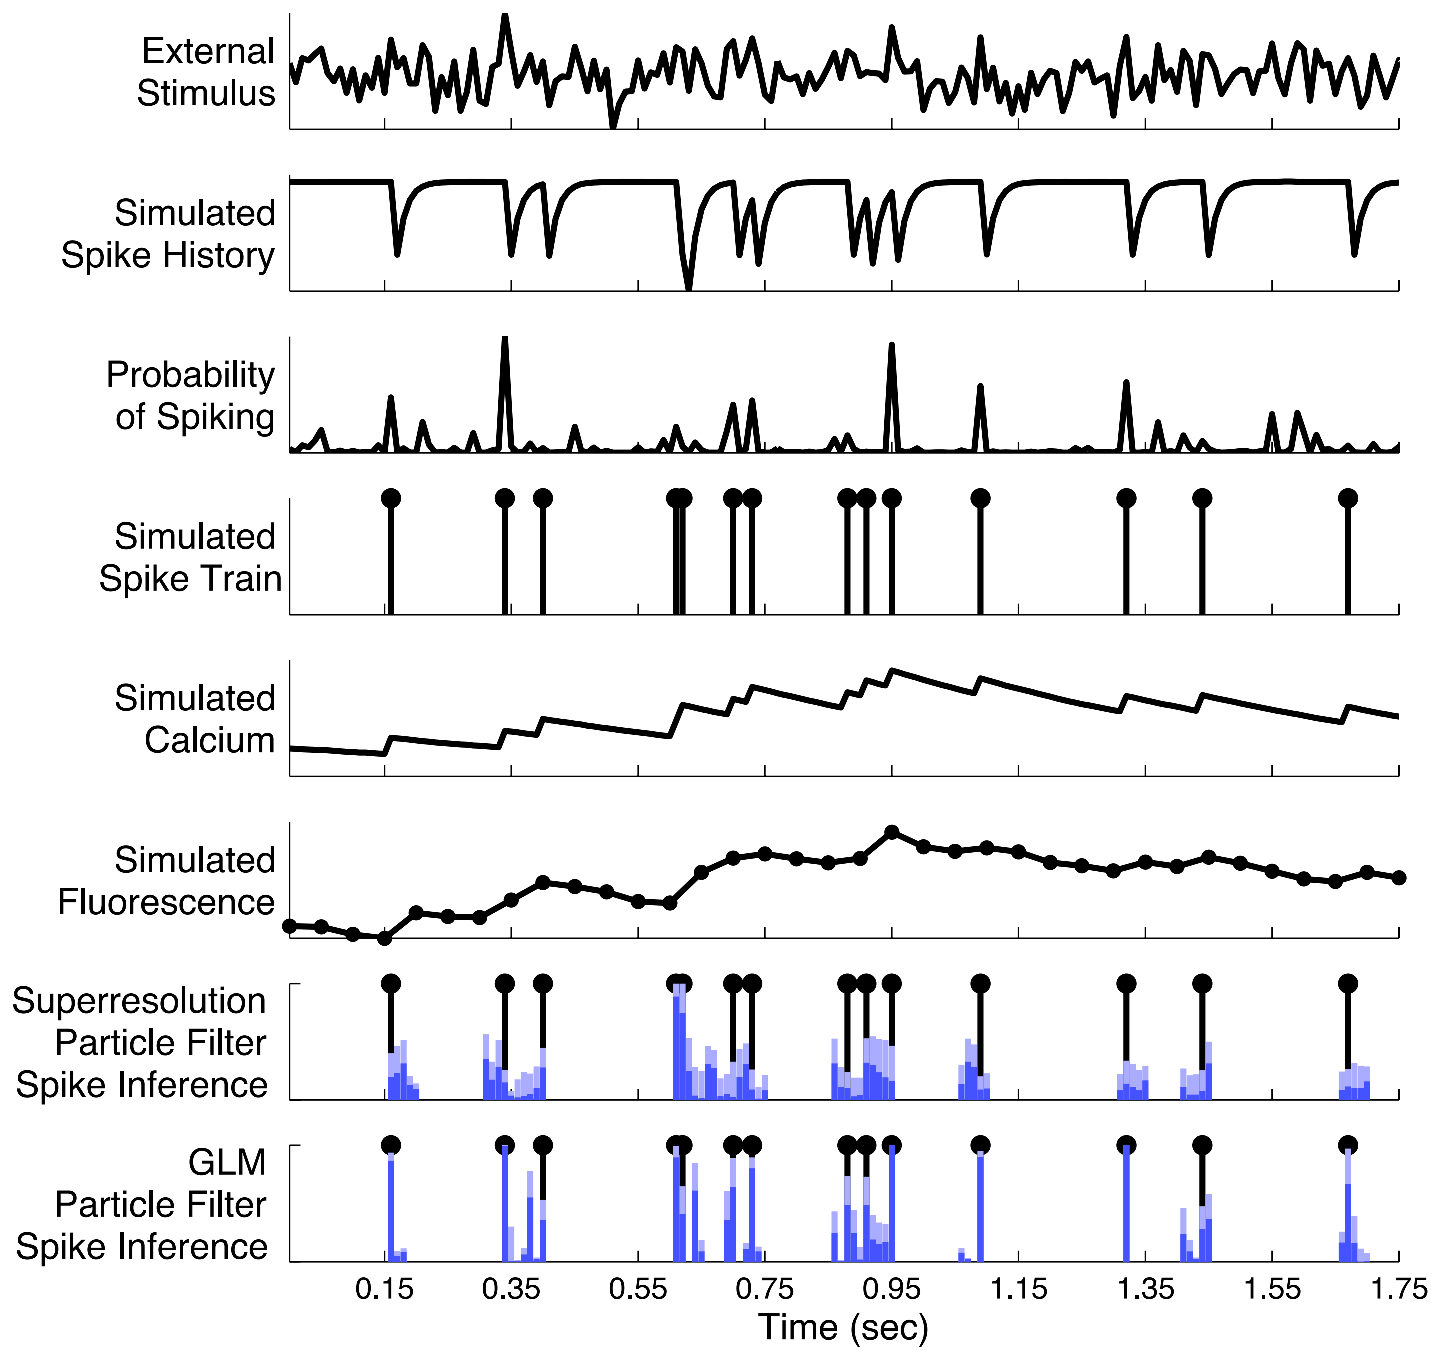

superresolution: array of results upon subsampling in temporal domain

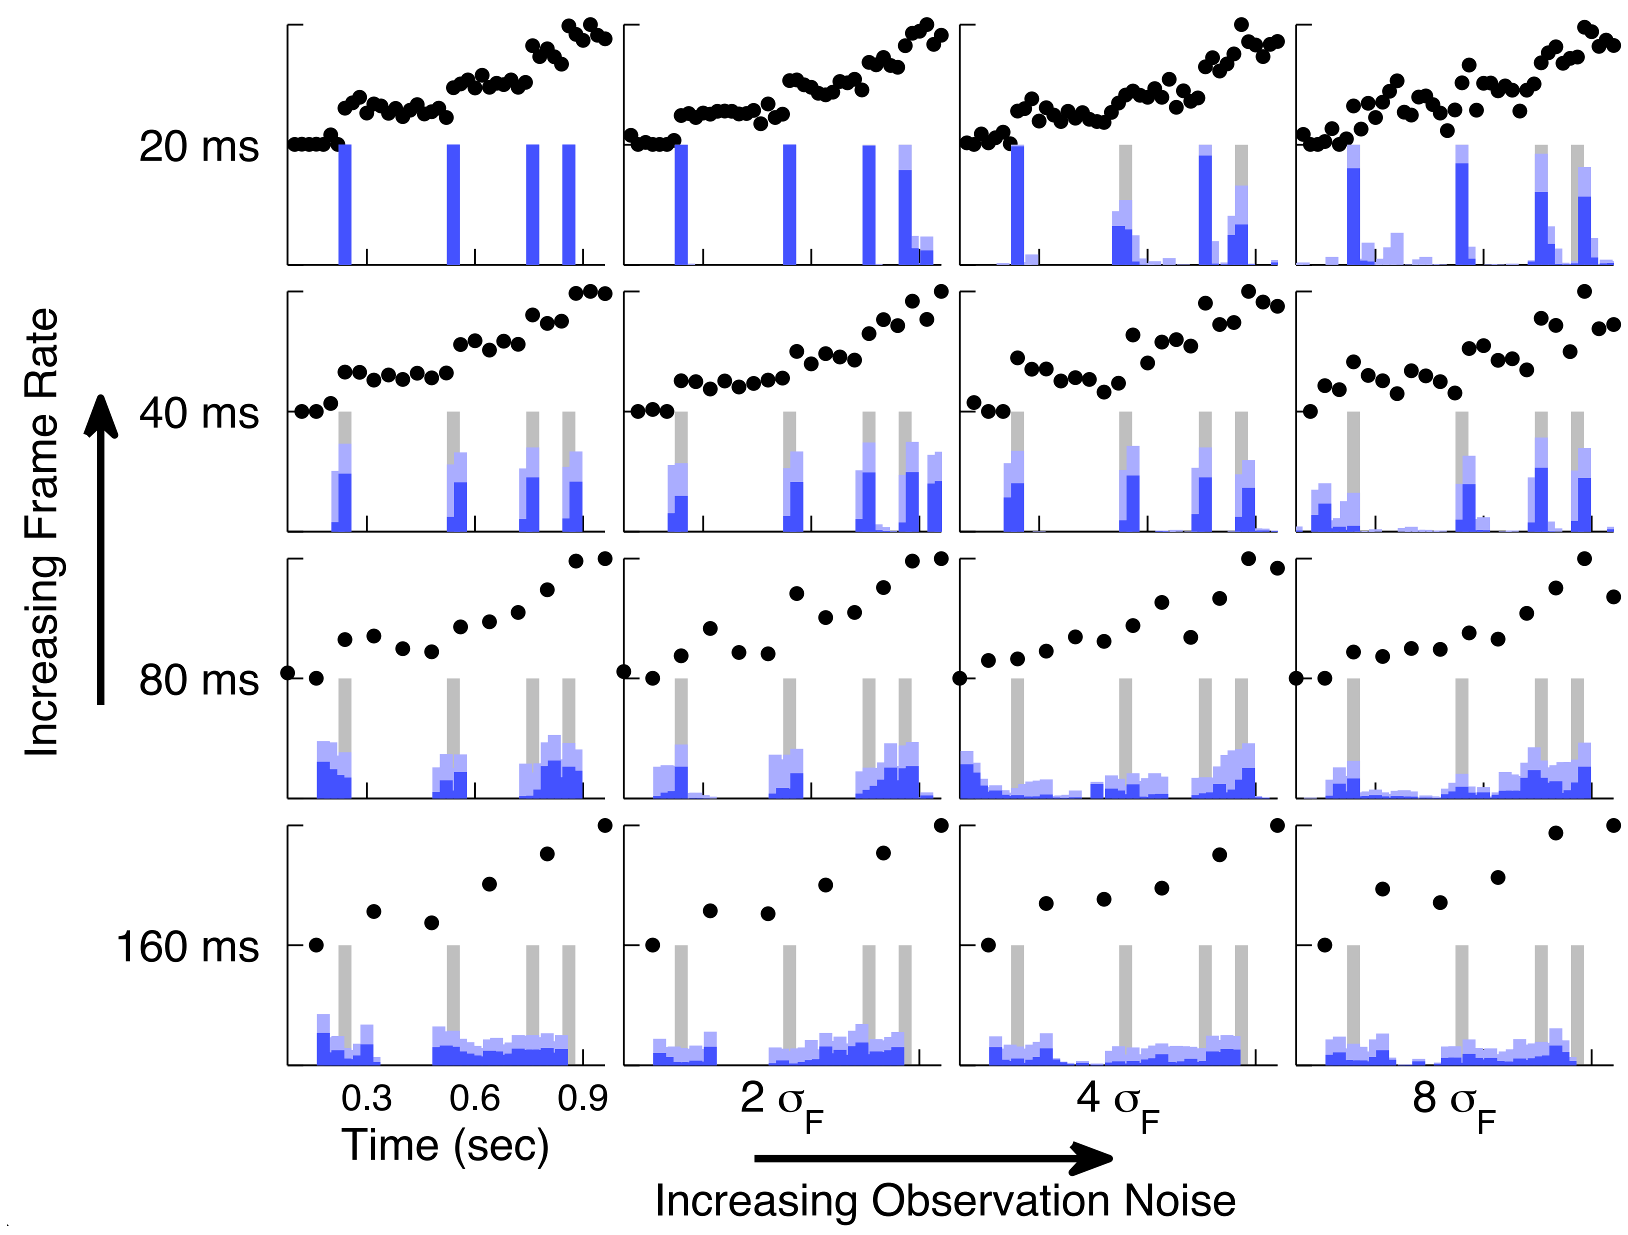

main result on real data with ground truth

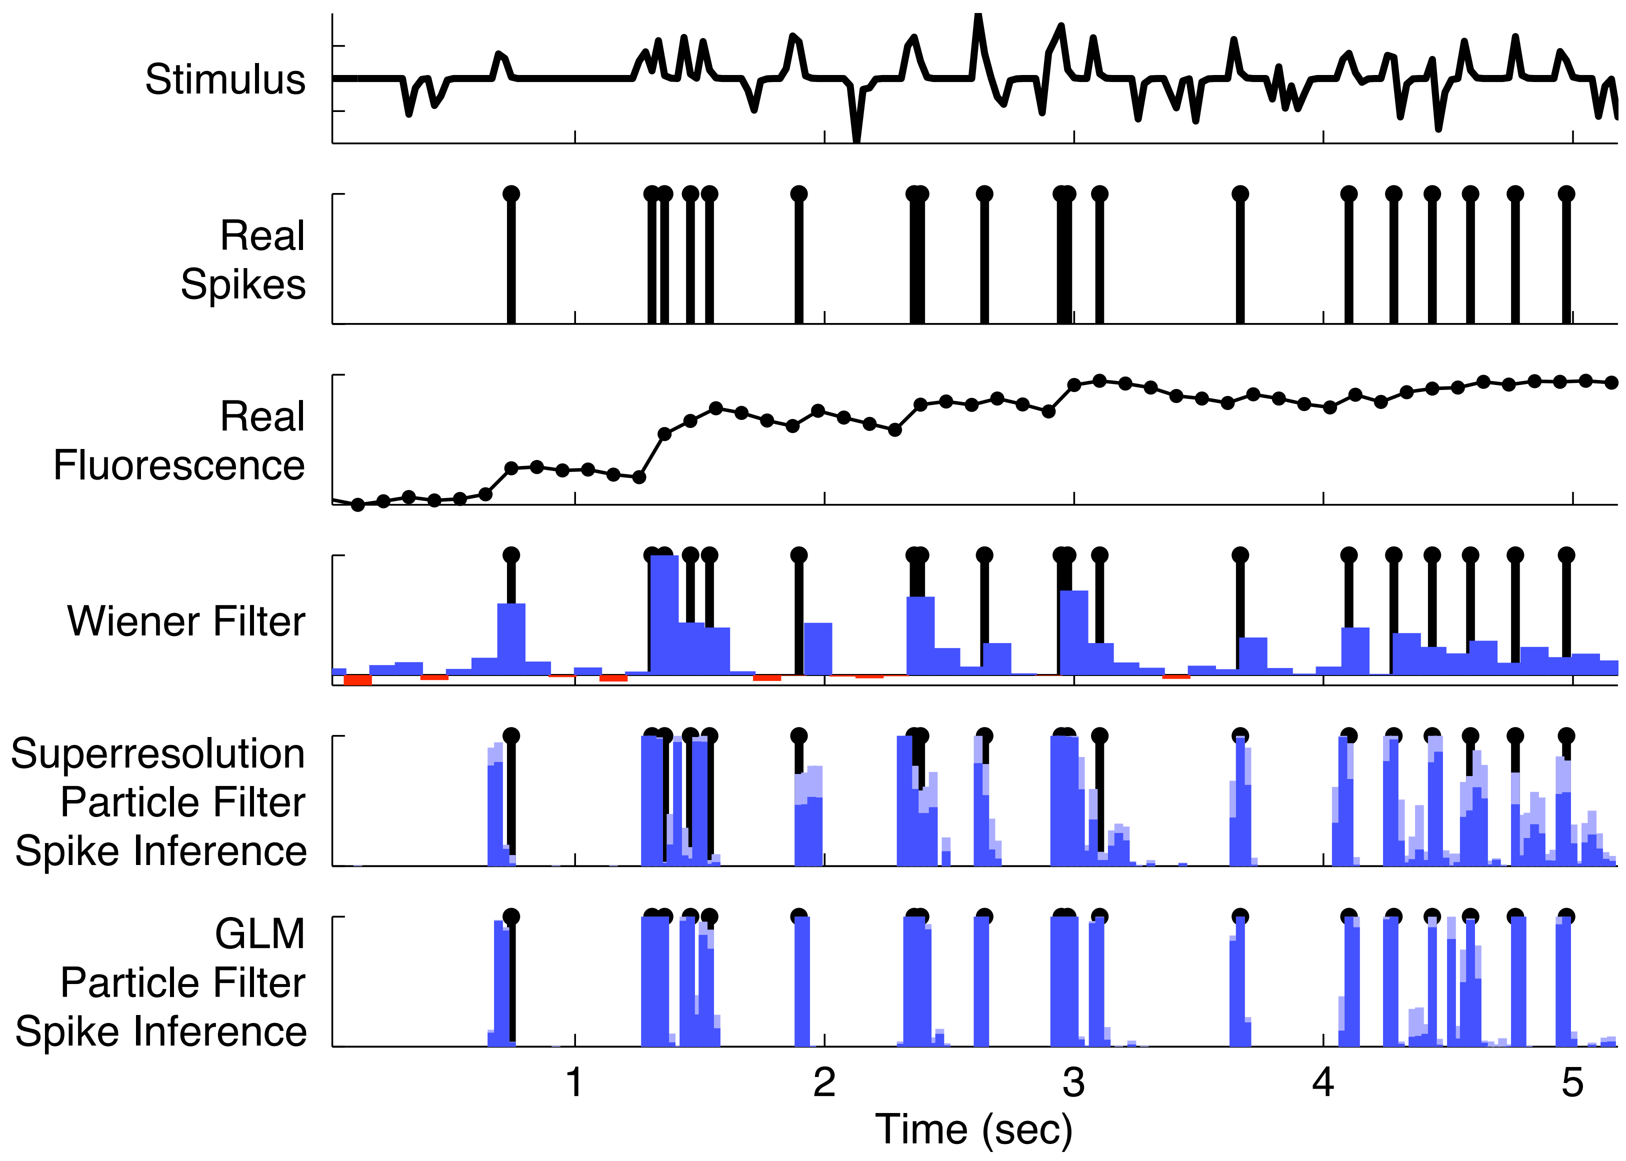

an even less simple model: an ensemble of N neurons

$$F_{i,t} = \alpha_i S([\mathrm{Ca}^{2+}]_{i,t}) + \beta_i + (S([\mathrm{Ca}^{2+}]_{i,t}) + \sigma_{i,F})\varepsilon_{F_i,t}$$

$$[\mathrm{Ca}^{2+}]_{i,t} = a_i[\mathrm{Ca}^{2+}]_{i,t-1} + A_i n_{i,t} + d_i + \sigma_{c_i}\sqrt{\Delta}\varepsilon_{c_i,t}$$

$$n_{i,t} \sim \mathcal{B}\big(n_{i,t}; f(\boldsymbol{k}_i'\widetilde{\boldsymbol{x}}_t)\big)$$

• note that $\widetilde{\boldsymbol{x}}_t$ is augmented relative to $\boldsymbol{x}_t$, as it also includes the impact of other neurons (and $\boldsymbol{k}_i$ reflects this change as well)

why is this hard

- this is a high-dimensional inference problem (the number of hidden states scales linearly with N)
- we are interested in inferring the connection matrix, where the # of elements in this matrix scales quadratically with N
- we use SMC to generate a good proposal distribution in the context of a blockwise gibbs-metropolis sampler
- we condition samples on both previous spiking history of neuron i, and future spiking of all other neurons
- all computations are recursive, so our algorithm is linear in T
- probability of acceptance tends to 1 as the number of particles increases and/or coupling terms are weak
- other people have been thinking along similar lines (e.g., Neal et al. 2003; Andrieu et al., submitted)

pseudocode for a SMC Metropolis

for each neuron **do**
  generate $N-1$ particles, sampling according to $P(\cdot|\{n,[\mathrm{Ca}^{2+}]\}_{i,t-1}^{(l)},\{n,[\mathrm{Ca}^{2+}]\}_{\setminus i,t-1},F_t)$
  add current path to population of particles to obtain a restricted space and compute appropriately normalized transition probabilities
  use standard forward-backward sampling algorithm to sample z from this augmented space
  compute $q(z)$, probability of sampling z using forward-backward recursion
  compute probability of acceptance $r=[q(y)(p(z)]/[q(z)p(y)]$ where y is the current path, and $p(z)$ is the posterior
end for

summary

- we use particle filters to infer spike trains from nonlinear and non-Gaussian observations of neural activity
- we can incorporate a parametric model governing spiking activity to refine our inferences
- using this model, we can obtain superresolution
- all the parameters may be estimated using a very short sequence of observations (and does not ever require obtaining ground truth)
- in weakly-coupled ensembles of neurons, we propose a novel scheme to infer the connection matrix, in which the sampler takes advantage of the spike trains from all the neurons