Challenges and opportunities in statistical neuroscience

Liam Paninski

Department of Statistics and Center for Theoretical Neuroscience
Columbia University
http://www.stat.columbia.edu/∼liam
liam@stat.columbia.edu
December 9, 2011

Support: NIH/NSF CRCNS, Sloan, NSF CAREER, McKnight Scholar, DARPA.

The coming statistical neuroscience decade

Some notable recent developments:

- machine learning / statistics methods for extracting information from high-dimensional data in a computationally-tractable, systematic fashion
- computing (Moore's law, massive parallel computing, GPUs)
- optical methods for recording and stimulating many genetically-targeted neurons simultaneously
- high-density multielectrode recordings (Litke's 512-electrode retinal readout system; Shepard's 65,536-electrode active array)

Three challenges

1. Reconstructing the full spatiotemporal voltage on a dendritic tree given noisy, intermittently-sampled subcellular measurements
2. Decoding behaviorally-relevant information from multiple spike trains
3. Inferring connectivity from large populations of noisily-observed spike trains

1. Basic goal: understanding dendrites

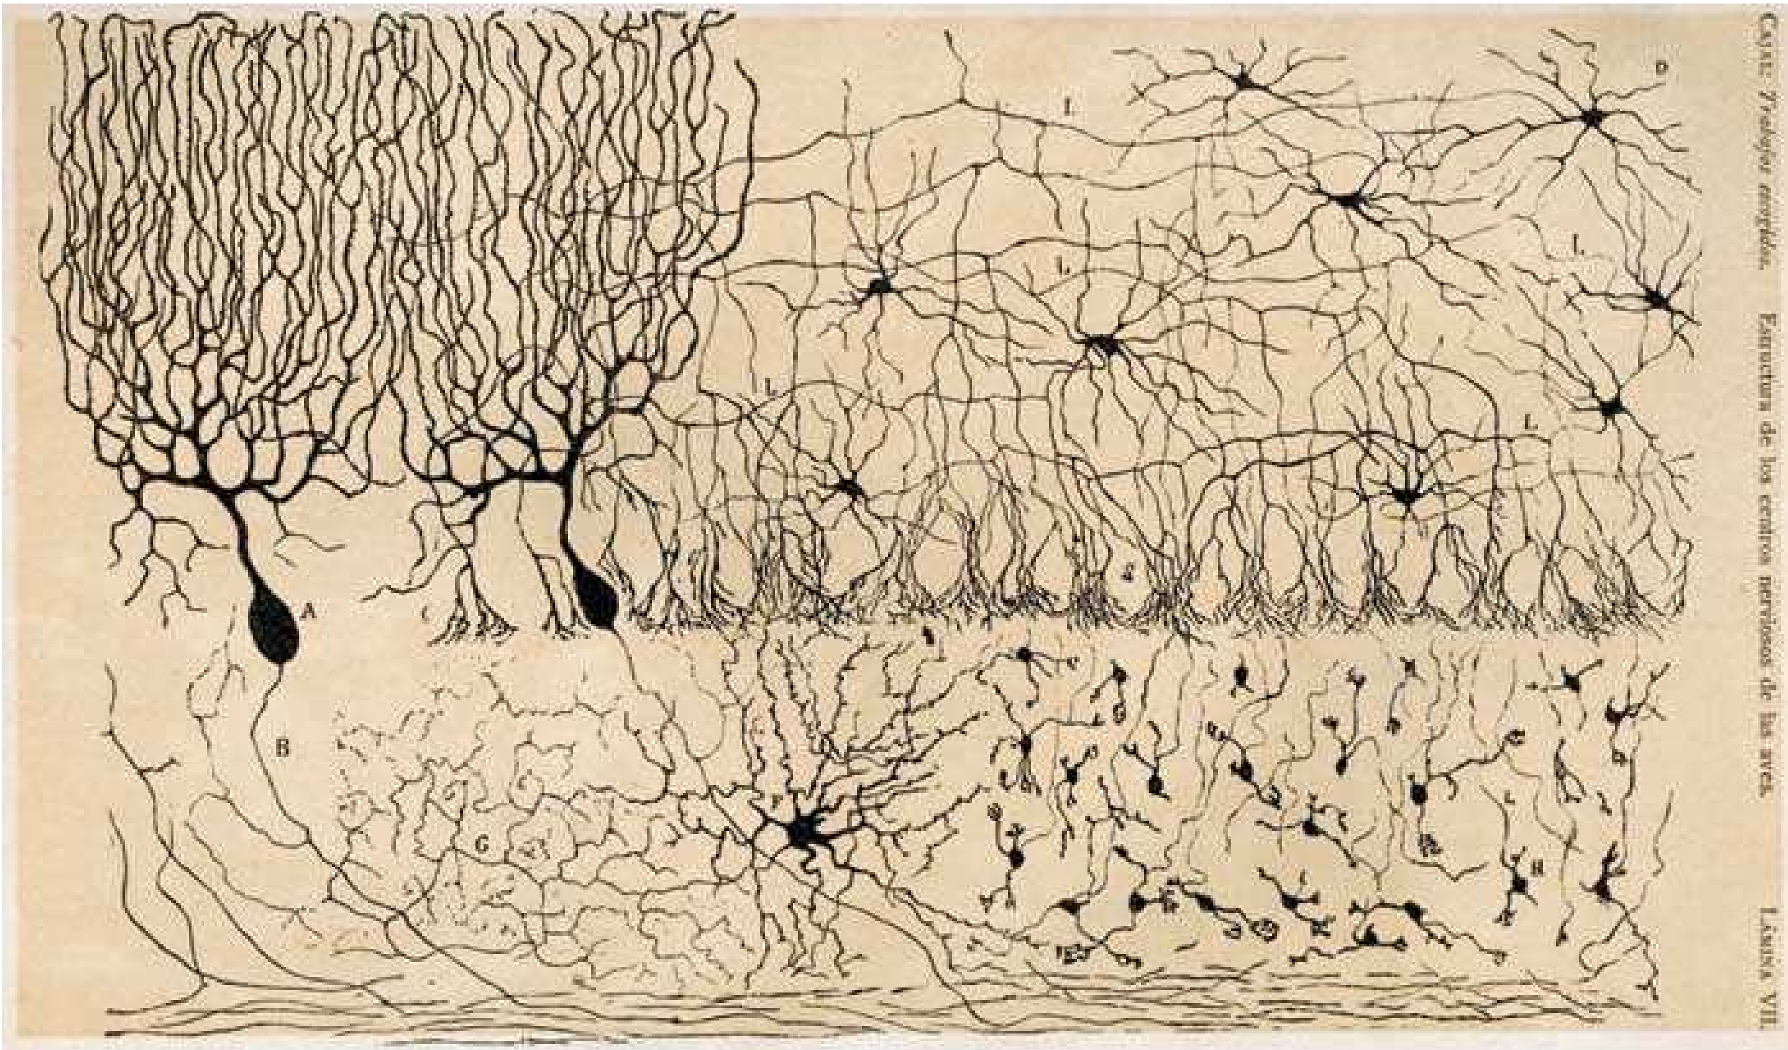

Ramon y Cajal, 1888.

The filtering problem

Spatiotemporal imaging data opens an exciting window on the computations performed by single neurons, but we have to deal with noise and intermittent observations.

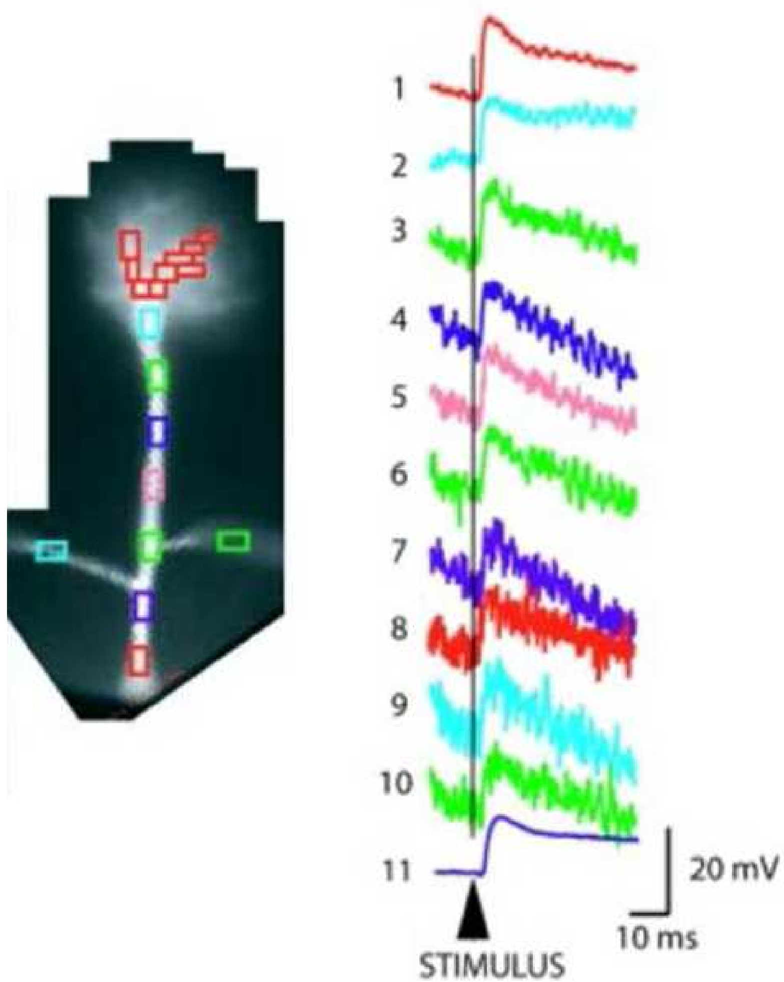

Basic paradigm: compartmental models

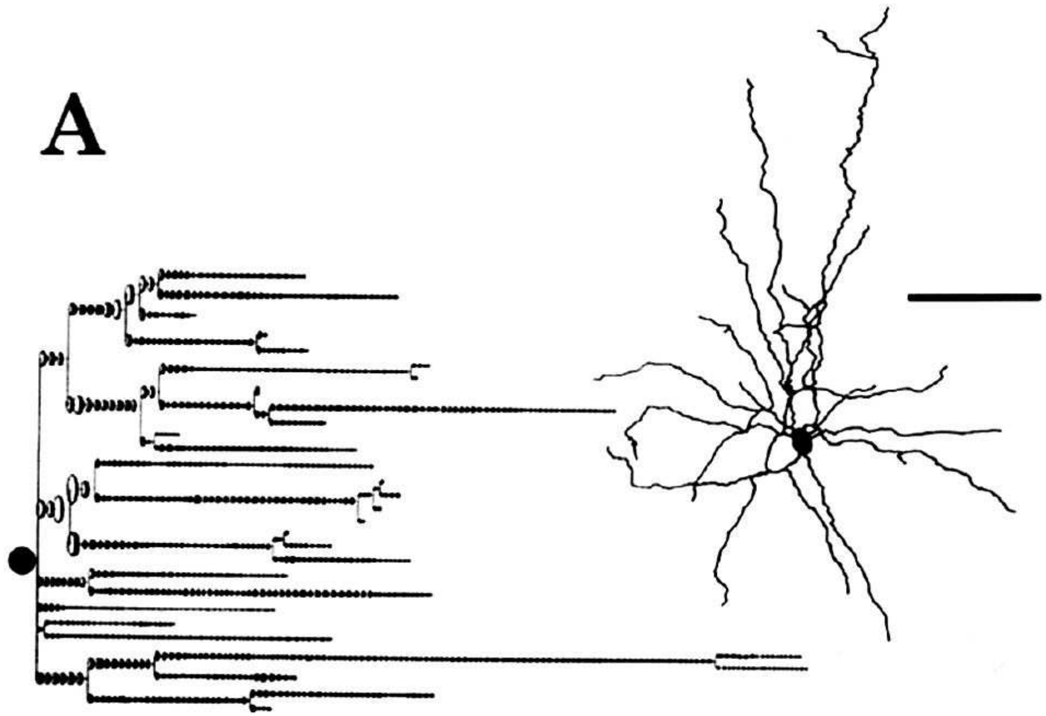

- write neuronal dynamics in terms of equivalent nonlinear, time-varying RC circuits
- leads to a coupled system of stochastic differential equations

Inference of spatiotemporal neuronal state given noisy observations

State-space approach: q_t = state of neuron at time t.

We want $p(q_t|Y_{1:t}) \propto p(q_t, Y_{1:t})$. Markov assumption:

$$p(Q,Y) = p(Q)p(Y|Q) = p(q_1)\left(\prod_{t=2}^{T} p(q_t|q_{t-1})\right)\left(\prod_{t=1}^{T} p(y_t|q_t)\right)$$

To compute $p(q_t, Y_{1:t})$, just recurse

$$p(q_t, Y_{1:t}) = p(y_t|q_t)\int_{q_{t-1}} p(q_t|q_{t-1})p(q_{t-1}, Y_{1:t-1})dq_{t-1}.$$

Linear-Gaussian case: requires $O(\dim(q)^3 T)$ time; in principle, just matrix algebra (Kalman filter). Approximate solutions in more general case via sequential Monte Carlo (Huys and Paninski, 2009).

Major challenge: $\dim(q)$ can be $\approx 10^4$ or greater.

Low-rank approximations

Key fact: current experimental methods provide just a few low-SNR observations per time step.

Basic idea: if dynamics are approximately linear and time-invariant, we can approximate Kalman covariance $C_t = cov(q_t|Y_{1:t})$ as a low-rank perturbation of the prior covariance $C_0 + U_t D_t U_t^T$, with $C_0 = \lim_{t\to\infty} cov(q_t)$.

In many cases we can solve linear equations involving C_0 in $O(\dim(q))$ time; in this case we use the fact that the dendrite is a tree, and fast methods are available to solve the cable equation on a tree.

The necessary recursions — i.e., updating U_t, D_t and the Kalman mean $E(q_t|Y_{1:t})$ — involve linear manipulations of C_0, and can be handled in $O(\dim(q))$ time (Paninski, 2010).

Example: inferring voltage from subsampled observations

(Loading low-rank-speckle.mp4)

Application: synaptic locations/weights

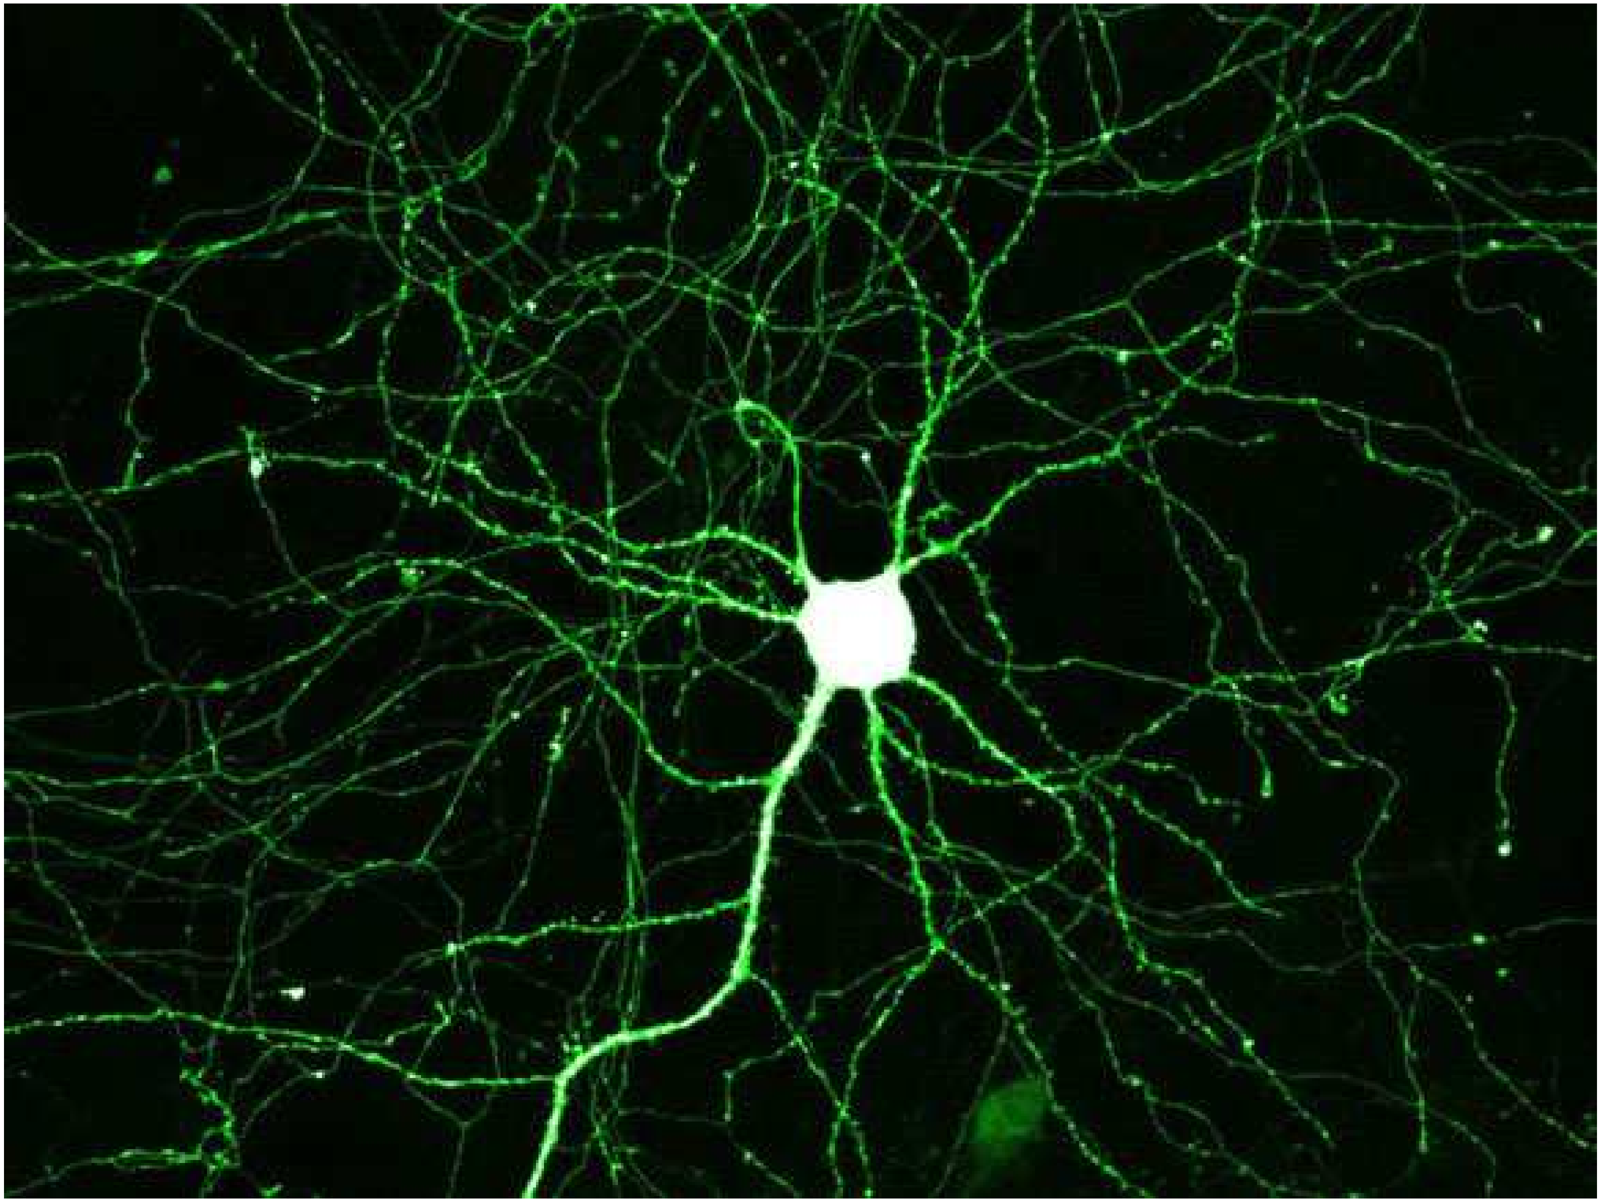

Application: synaptic locations/weights

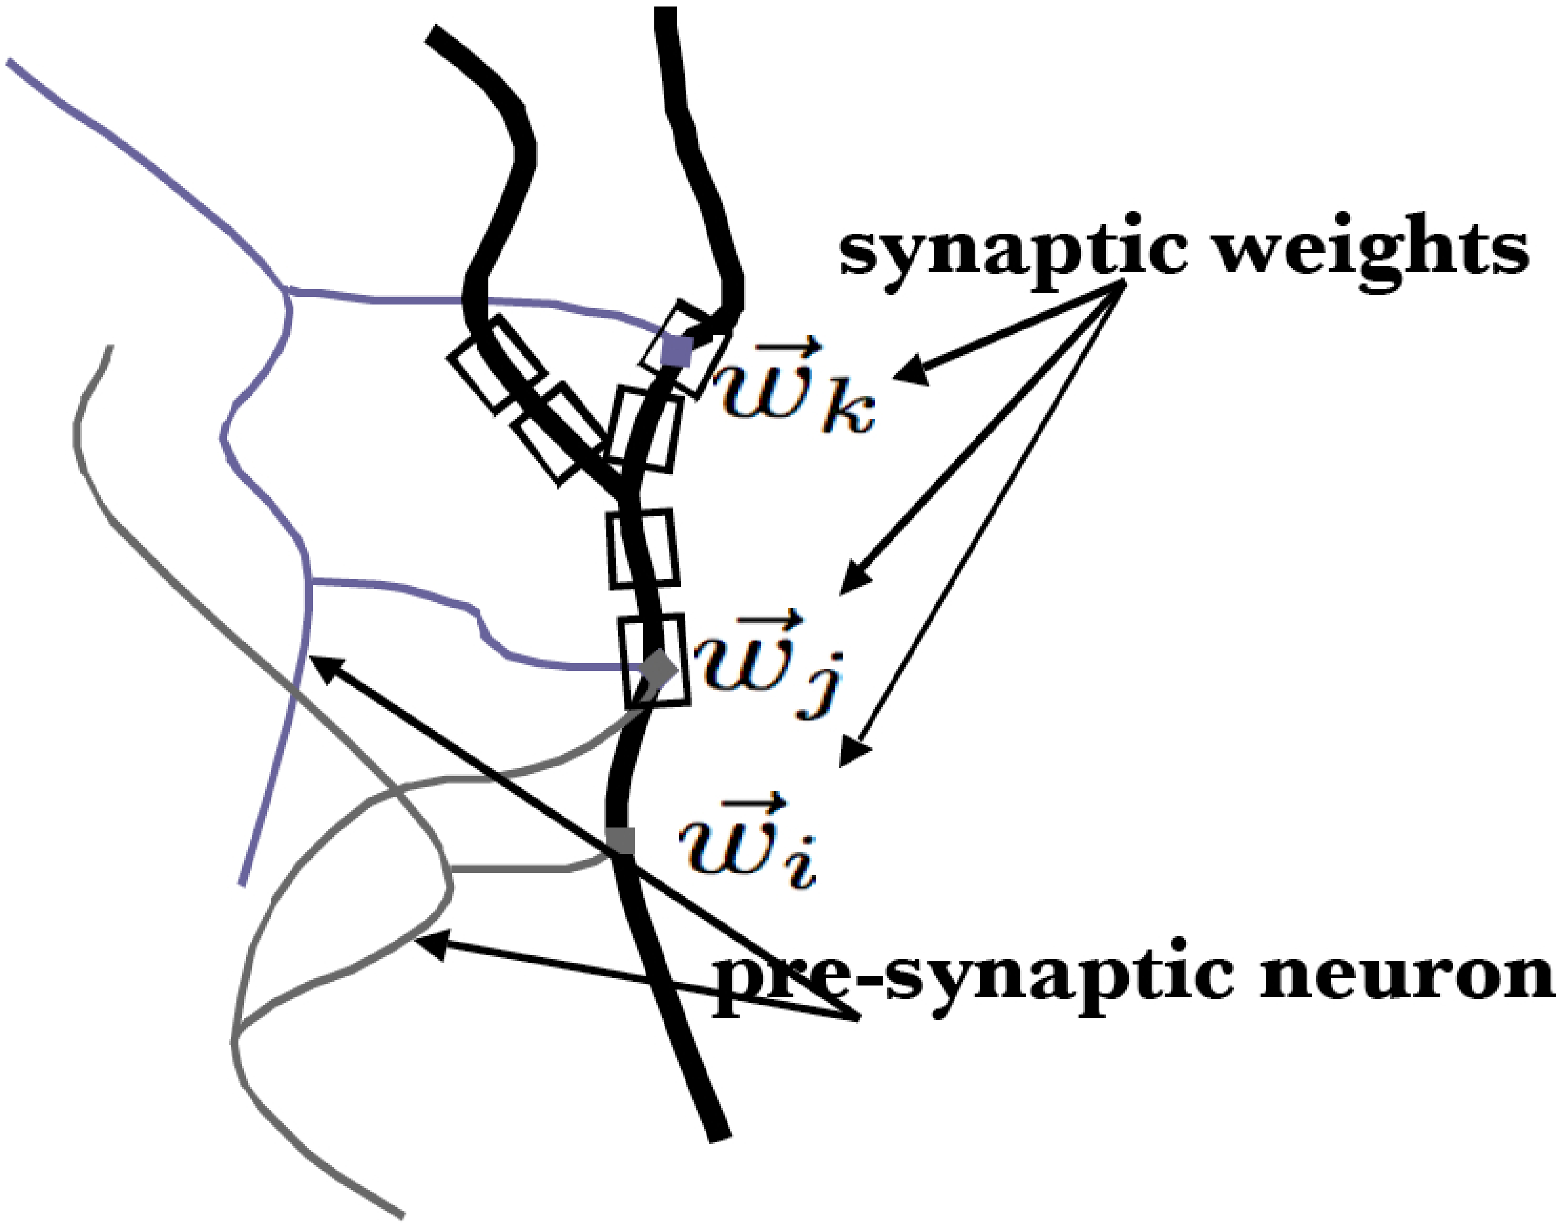

Application: synaptic locations/weights

Including known terms:

$$d\vec{V}/dt = A\vec{V}(t) + W\vec{U}(t) + \vec{\epsilon}(t);$$

$U_j(t)$ = known input terms.

Example: $U(t)$ are known presynaptic spike times, and we want to detect which compartments are connected (i.e., infer the weight matrix W).

Loglikelihood is quadratic; L_1-penalized loglikelihood can be optimized efficiently with homotopy approach. Total computation time is $O(NTk)$: N = # compartments, T = # timesteps, k = # nonzero weights.

Application: synaptic locations/weights

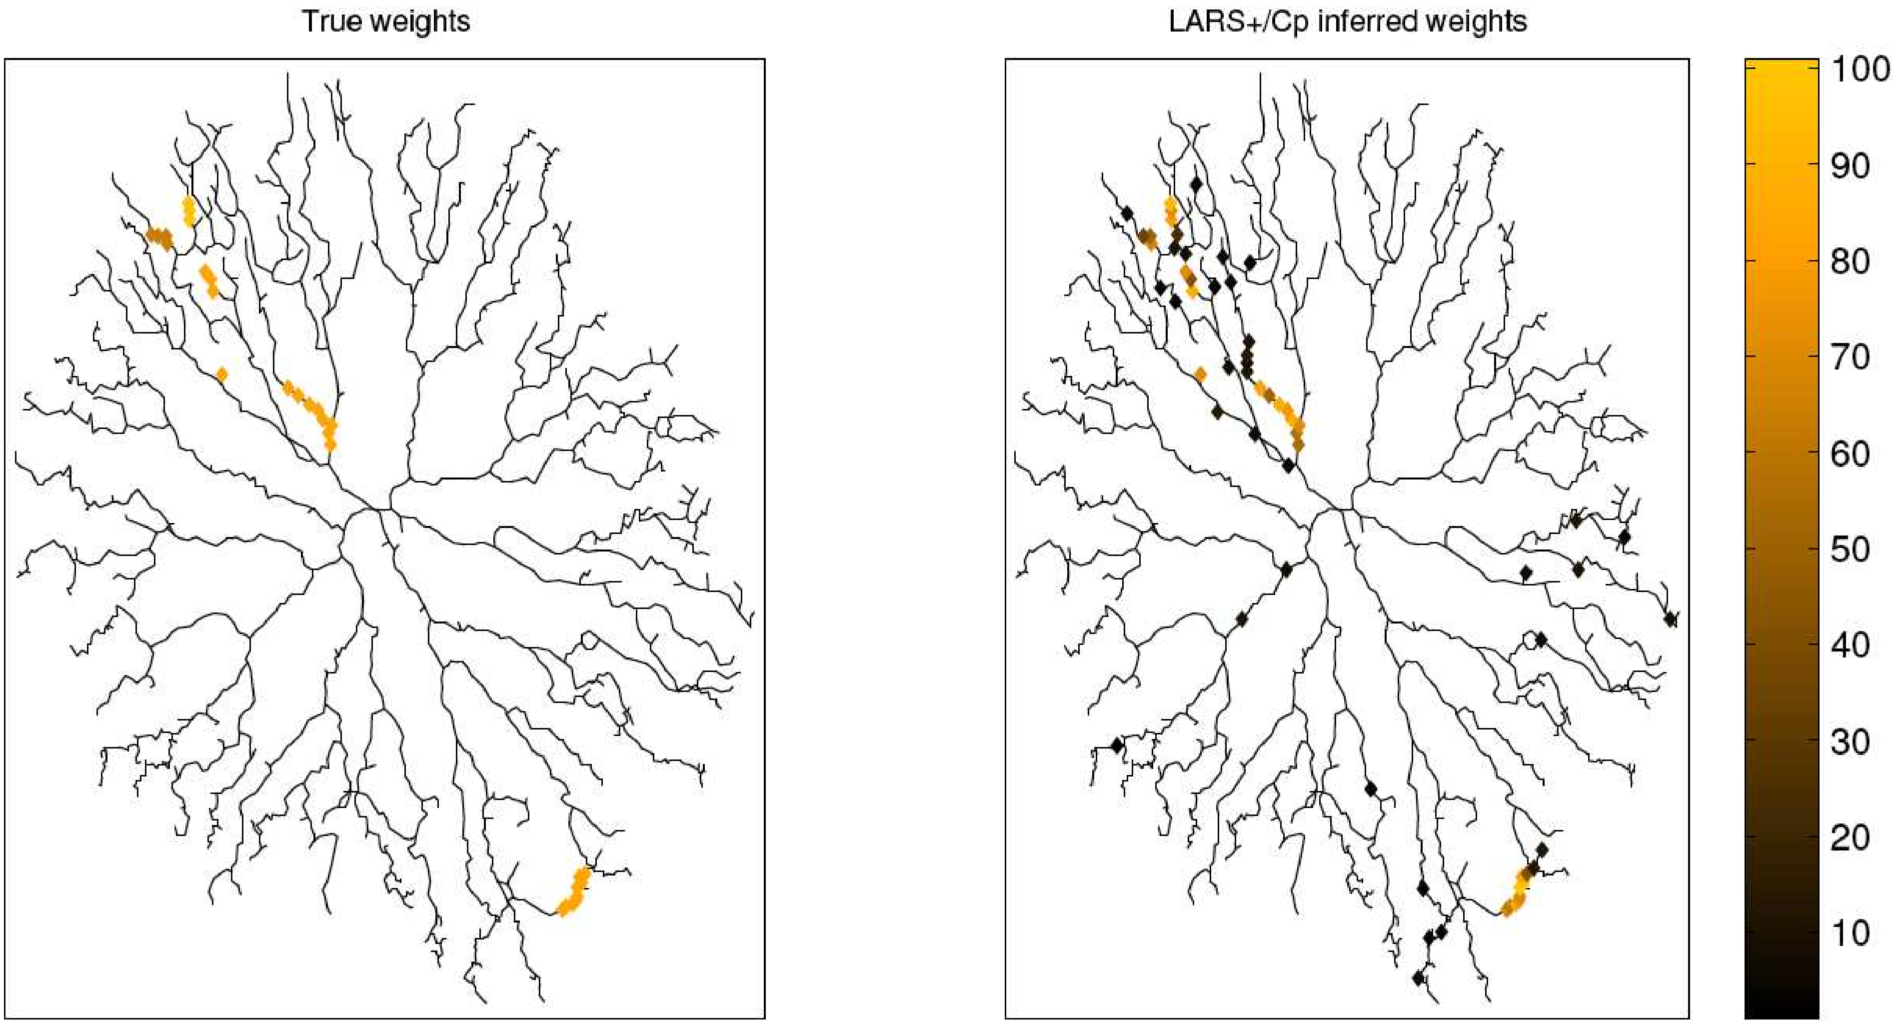

Part 2: optimal decoding of spike train data

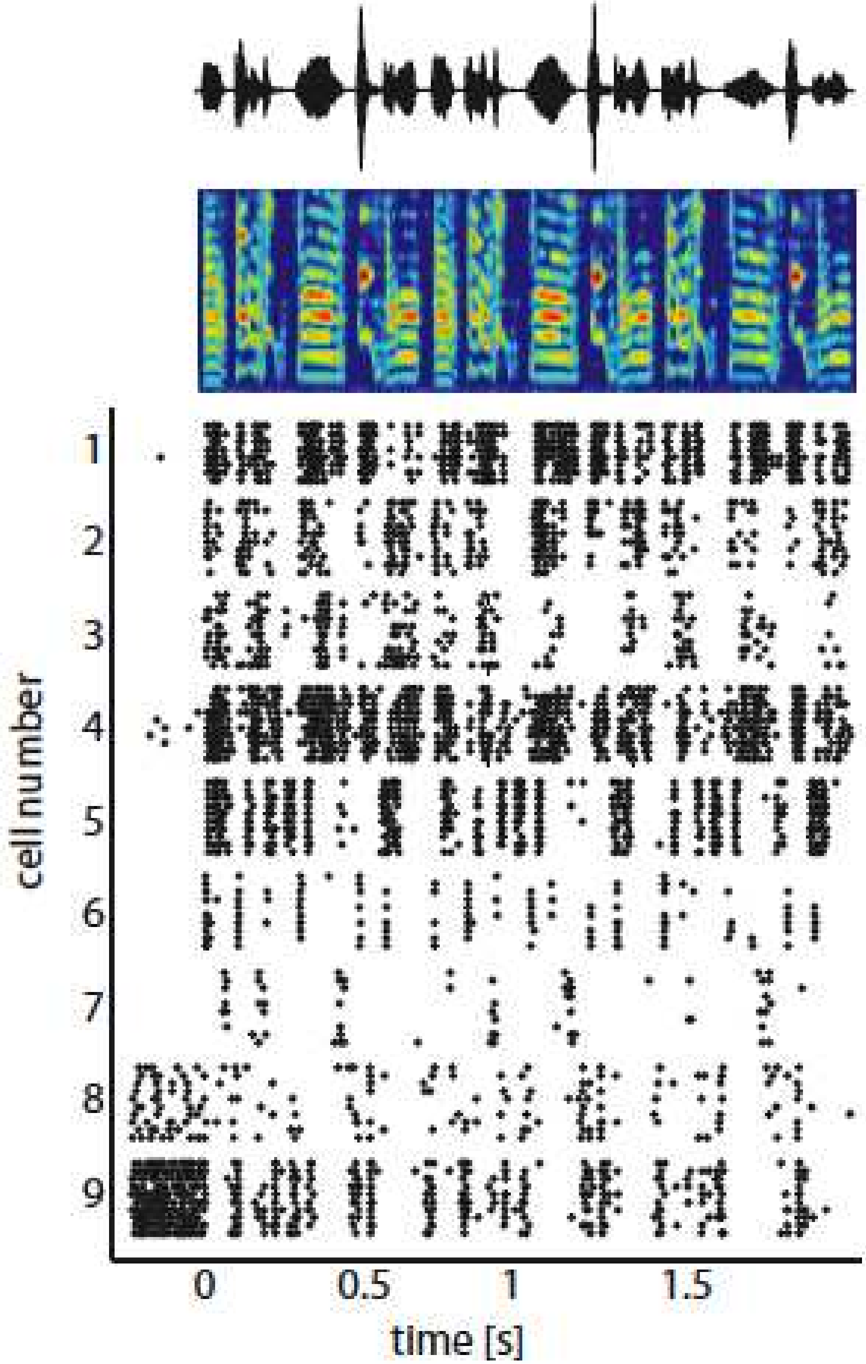

Encoding model

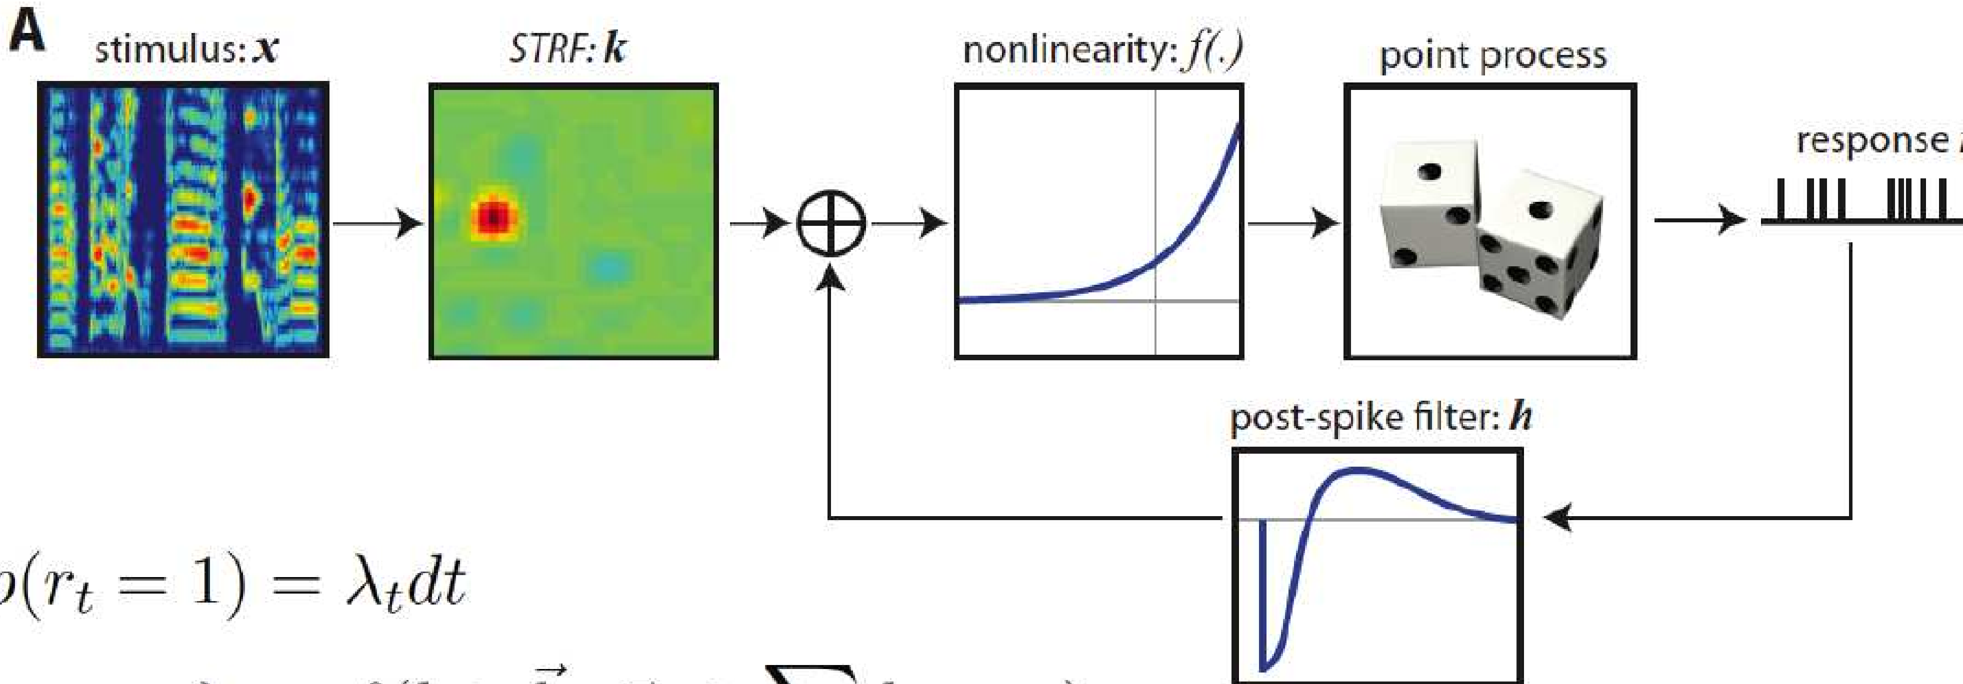

$$p(r_t = 1) = \lambda_t dt$$

$$\lambda_t = f(b + \vec{k} \cdot \vec{x}_t + \sum_j h_j r_{t-j})$$

Parameters $(\vec{k}, h)$ estimated by L_1-penalized maximum likelihood (concave); f estimated by log-spline (Calabrese, Woolley et al. 2009). Currently the best predictive model of these spike trains.

Maximum a posteriori stimulus decoding

Bayesian MAP decoder:

$$\hat{X} = \arg\max_X p(X|R).$$

(X = spectrogram; very high-dimensional!) For this model, we have:

$$\begin{aligned}\log p(X|R) &= \log p(X) + \log p(R|X) + const. \\ &= \log p(X) + \sum_t \log p(r_t|X, R_{\ldots,t-1}) + const.\end{aligned}$$

Two basic observations:

- If $\log p(X)$ is concave, then so is $\log p(X|R)$, since each $\log p(r_t|X, Y_{\ldots,t-1})$ is.
- Hessian H of $\log p(R|X)$ w.r.t. X is banded: each $p(r_t|X, R_{\ldots,t-1})$ depends on X locally in time, and therefore contributes a banded term.

Newton's method iteratively solves $HX_{dir} = \nabla$. Solving banded systems requires $O(T)$ time, so computing MAP requires $O(T)$ time if log-prior is concave with a banded Hessian.

Application: fast optimal decoding

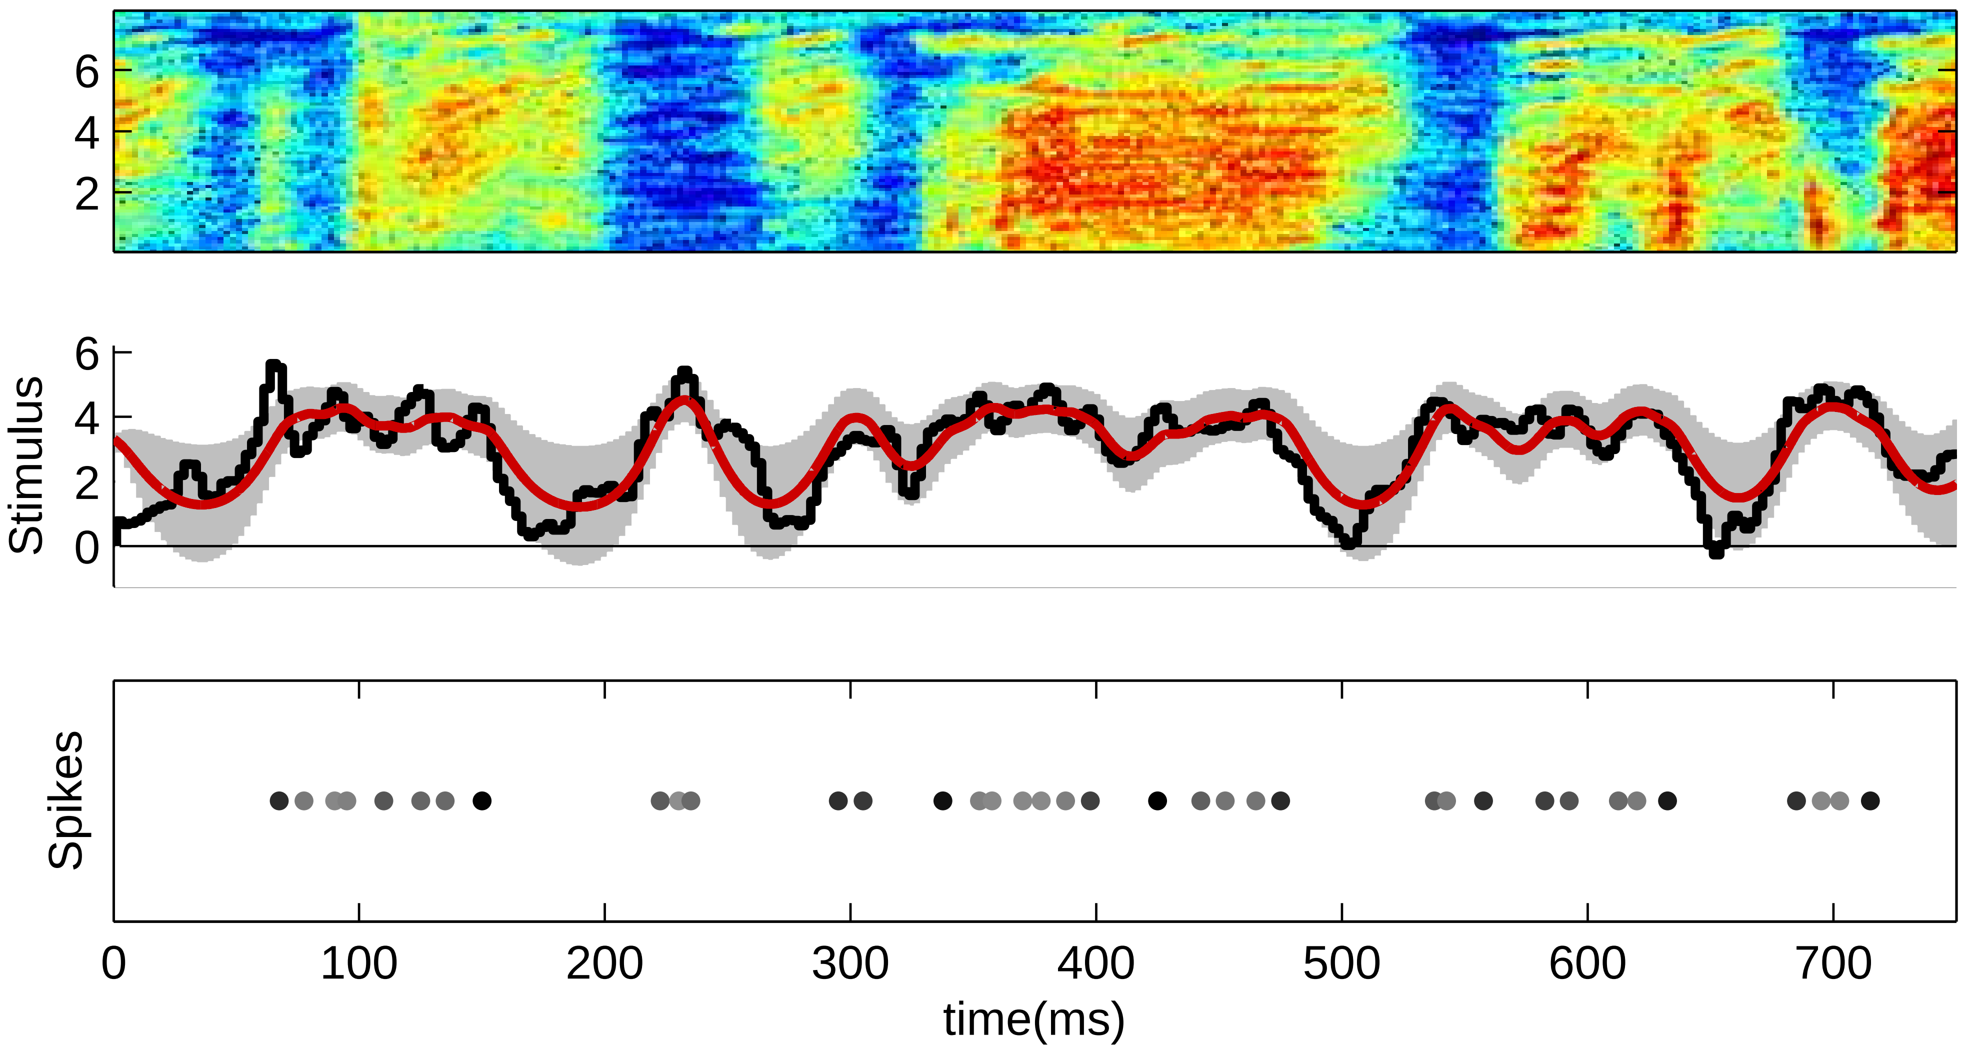

Decoding a full song

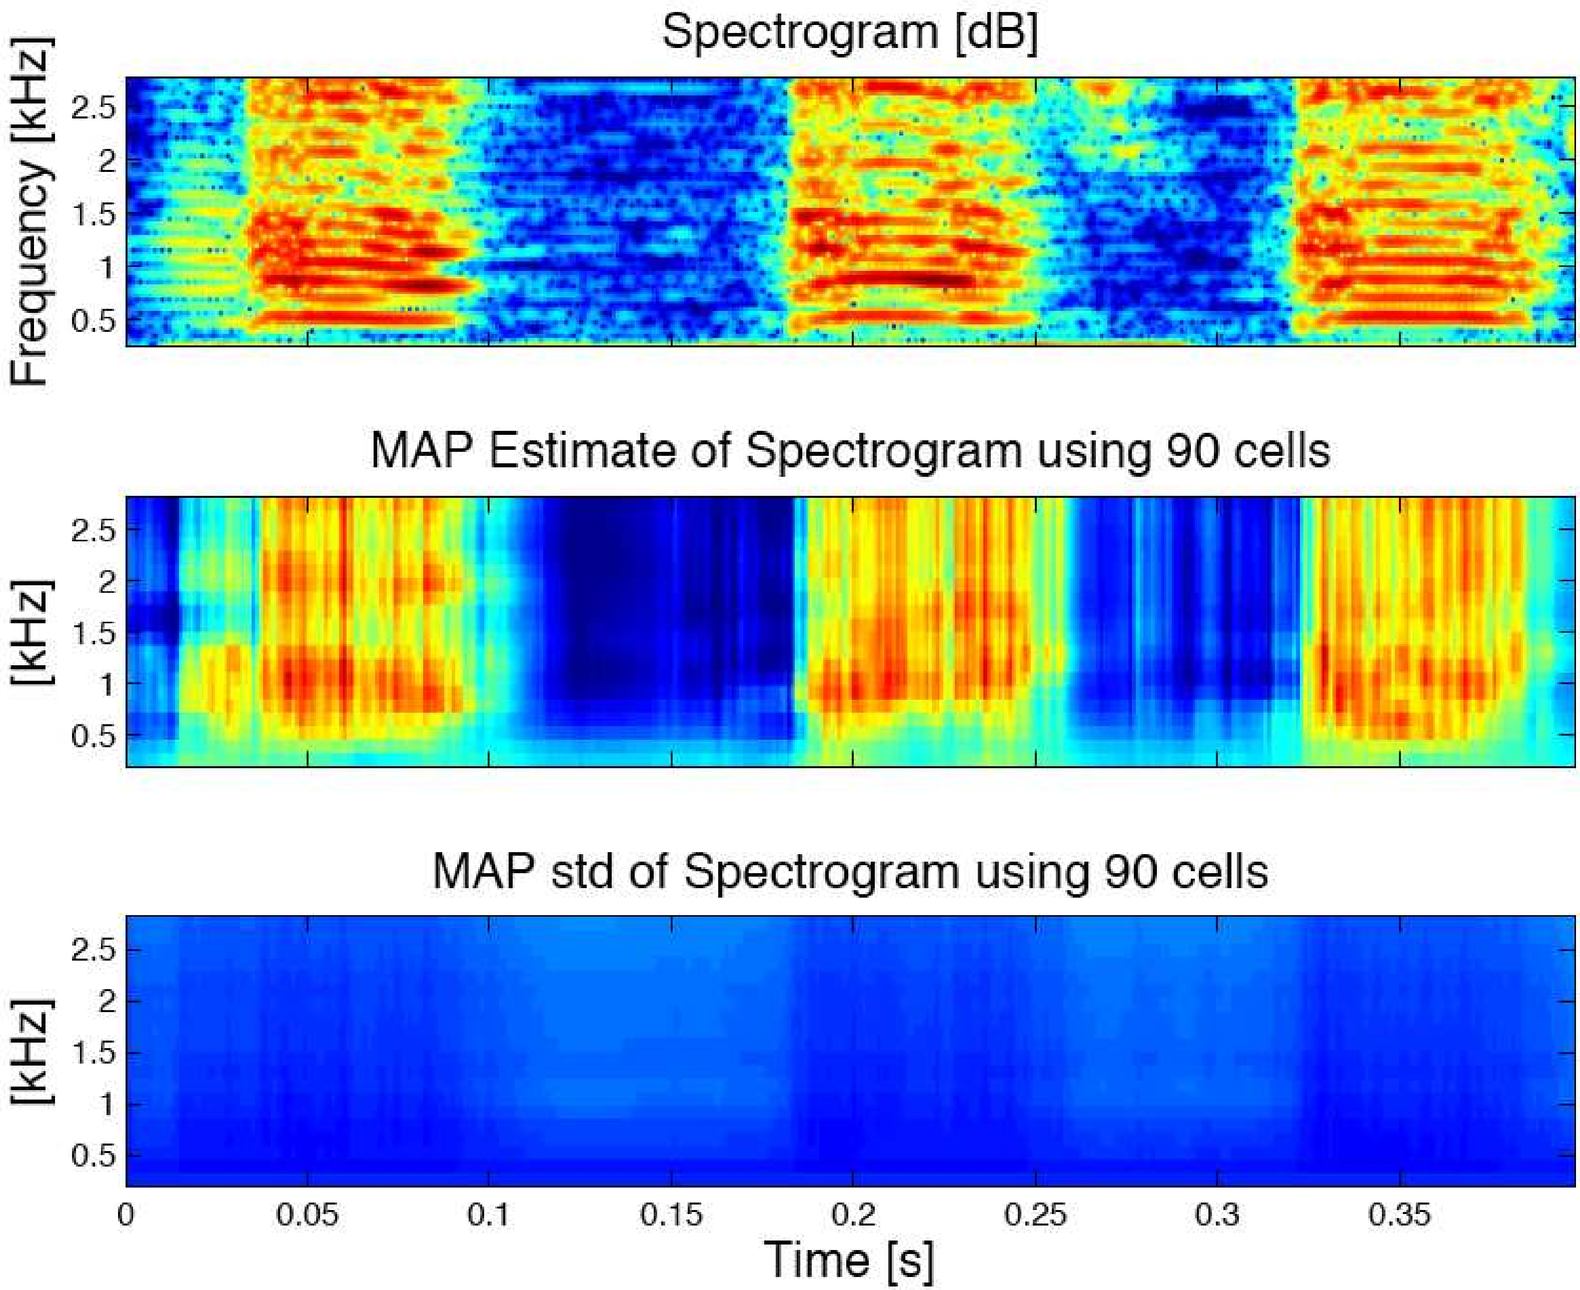

(Ramirez et al., 2010)

Part 3: circuit inference

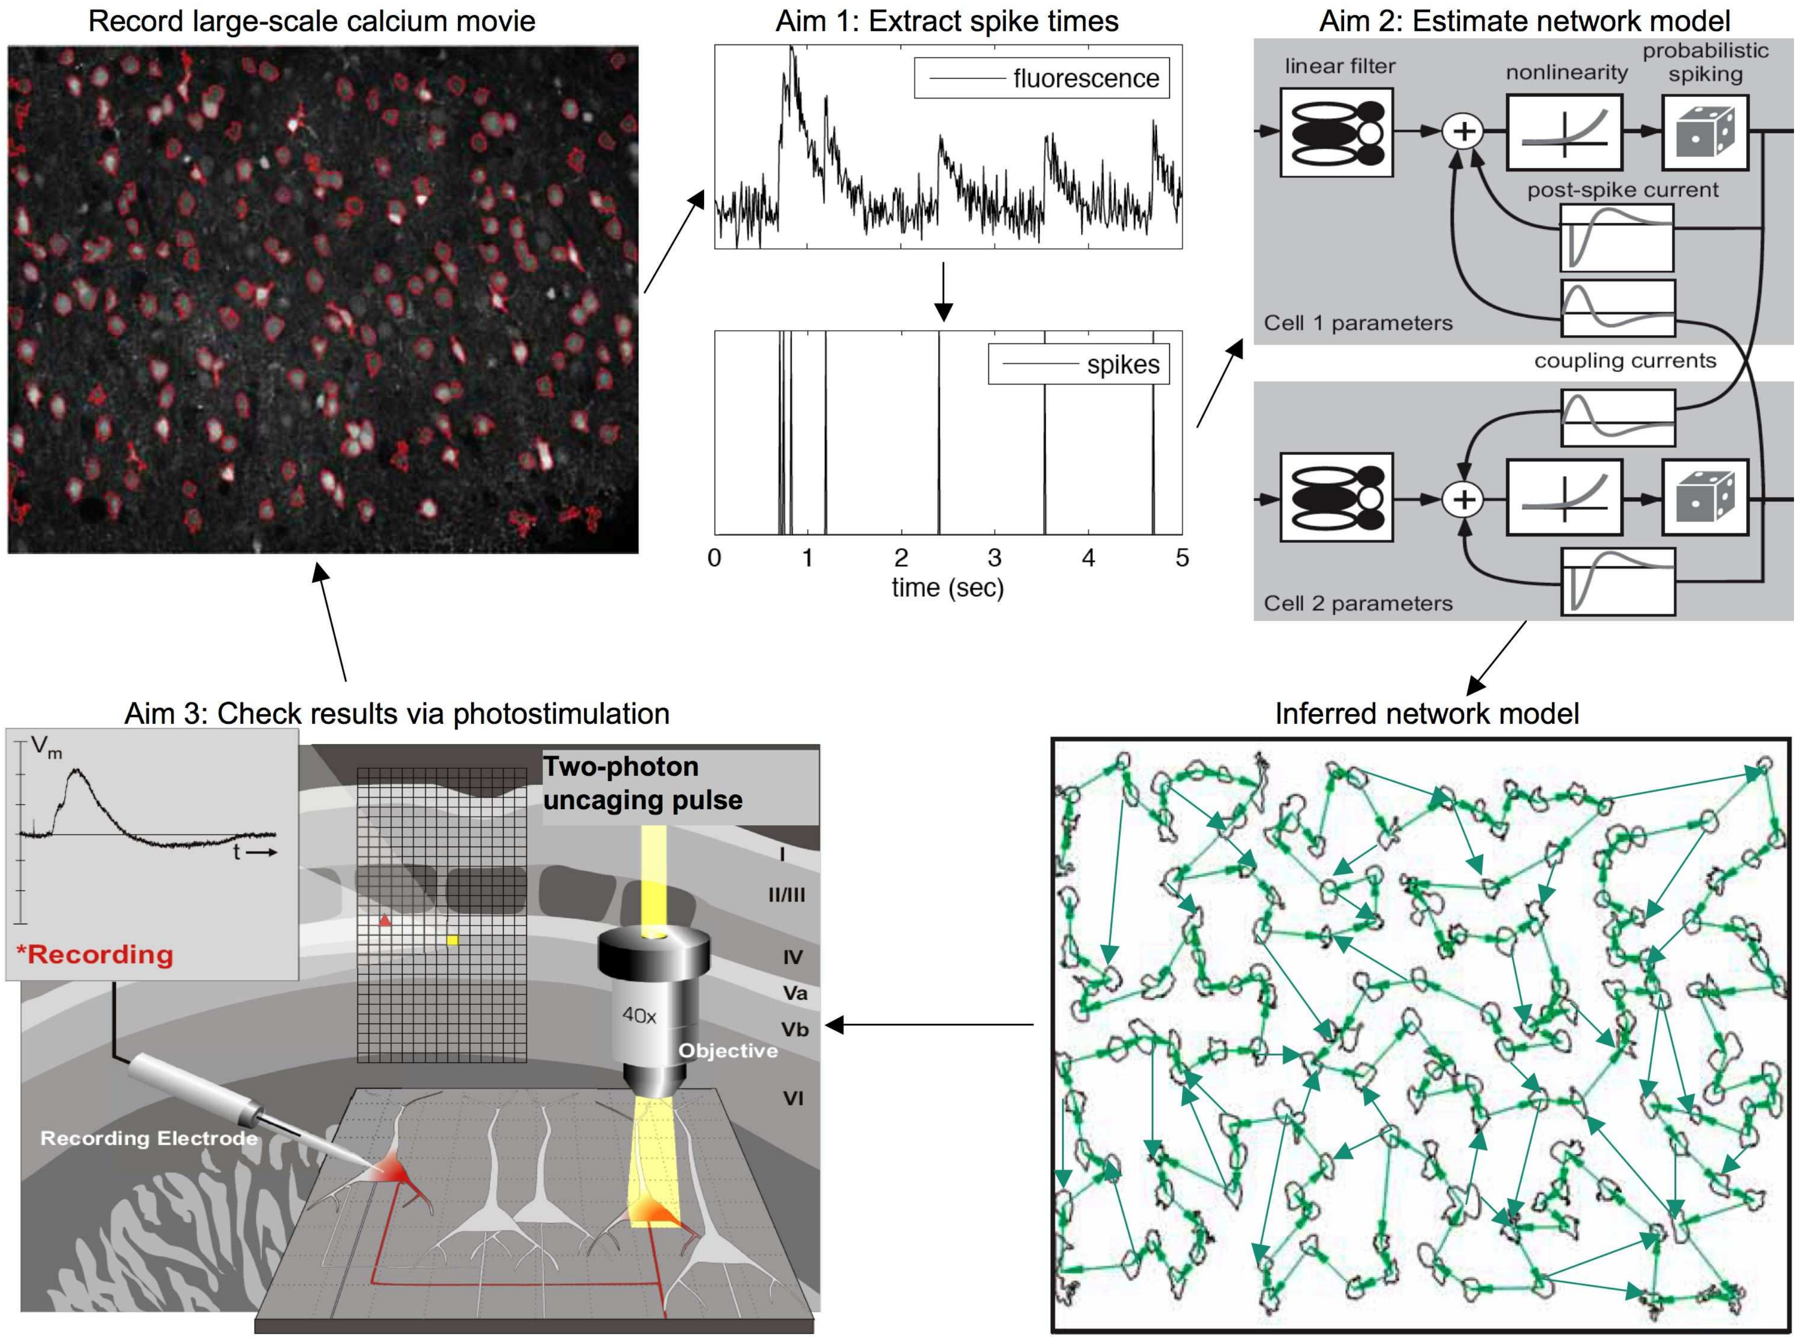

Challenge: slow, noisy calcium data

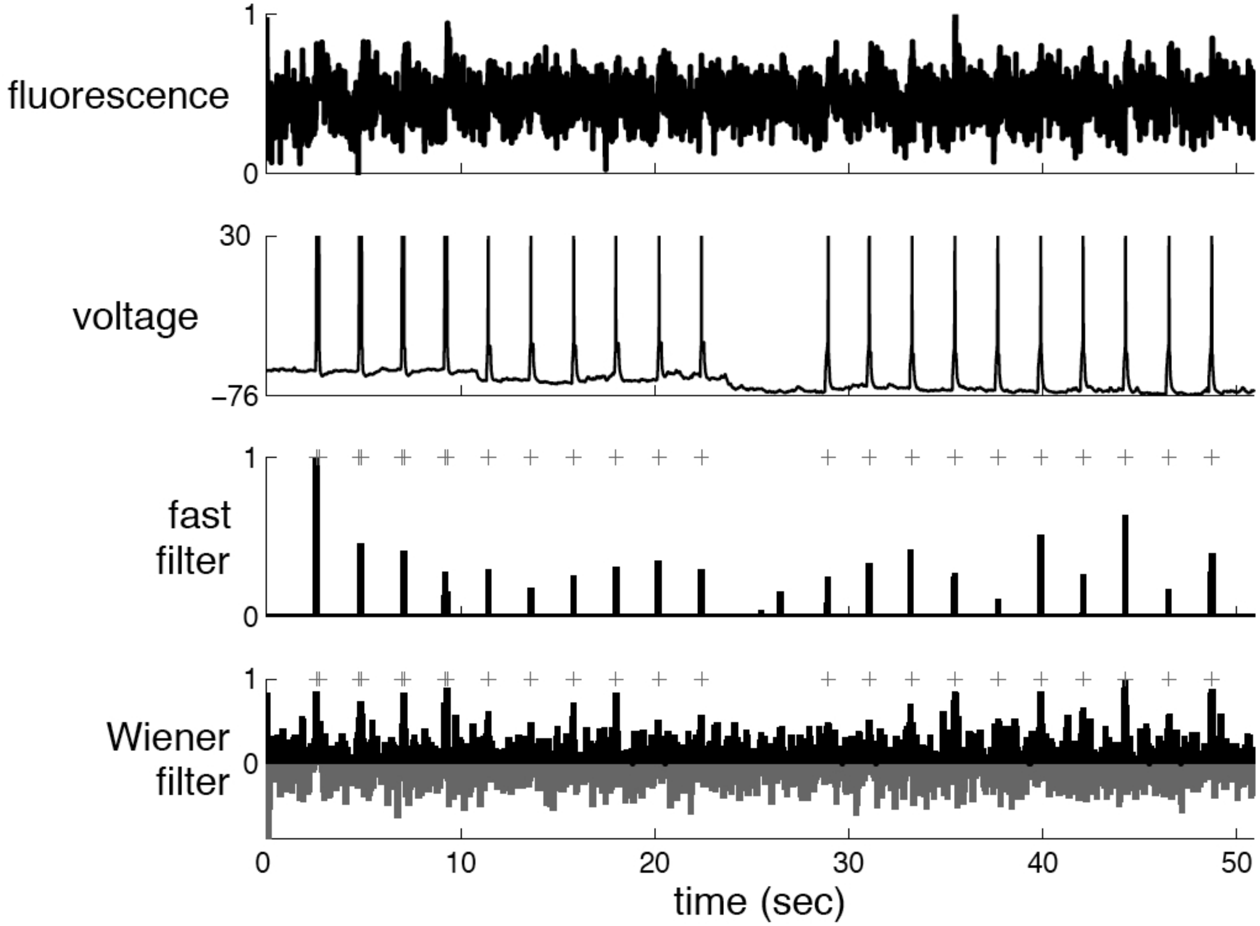

First-order model:

$$C_{t+dt} = C_t - dtC_t/\tau + r_t;\ r_t > 0;\ y_t = C_t + \epsilon_t$$

— $\tau \approx 100$ ms; nonnegative deconvolution problem. Can be solved by $O(T)$ relaxed constrained interior-point optimization (Vogelstein et al., 2010) or sequential Monte Carlo (Vogelstein et al., 2009).

Simulated circuit inference

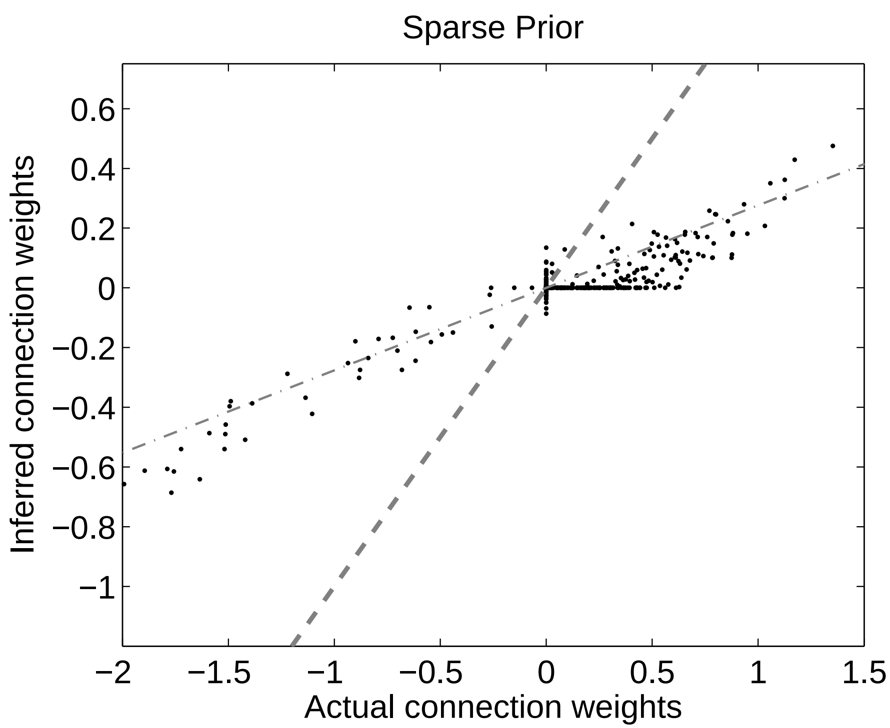

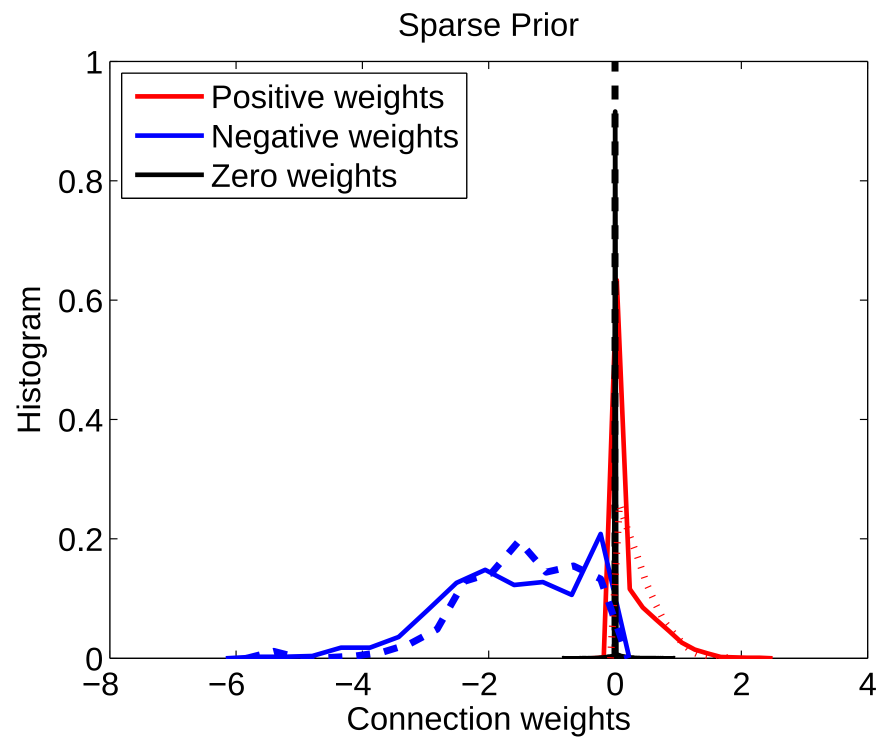

— Connections are inferred with the correct sign in conductance-based integrate-and-fire networks with biologically plausible connectivity matrices (Mishchenko et al., 2010).

— With appropriate approximations, online circuit inference in networks of ≈ 100 neurons is tractable on a desktop. Methods parallelize easily for larger networks.

Optimal control of spike timing

Optimal experimental design and neural prosthetics applications require us to perturb the network at will. How can we make a neuron fire exactly when we want it to?

Assume bounded inputs; otherwise problem is trivial.

Start with a simple model:

$$\lambda_t = f(\vec{k} * I_t + h_t).$$

Now we can just optimize the likelihood of the desired spike train, as a function of the input I_t, with I_t bounded.

Concave objective function over convex set of possible inputs I_t + Hessian is banded $\implies O(T)$ optimization (Ahmadian et al., 2010).

Optical conductance-based control of spiking

$$V_{t+dt} = V_t + dt\left(-gV_t + g_t^i(V^i - V_t) + g_t^e(V^e - V_t)\right) + \sqrt{dt}\sigma\epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0,1)$$

$$g_{t+dt}^i = g_t^i + dt\left(-\frac{g_t^i}{\tau_i} + a_{ii}L_t^i + a_{ie}L_t^e\right);\ g_{t+dt}^e = g_t^e + dt\left(-\frac{g_t^e}{\tau_i} + a_{ee}L_t^e + a_{ei}L_t^i\right)$$

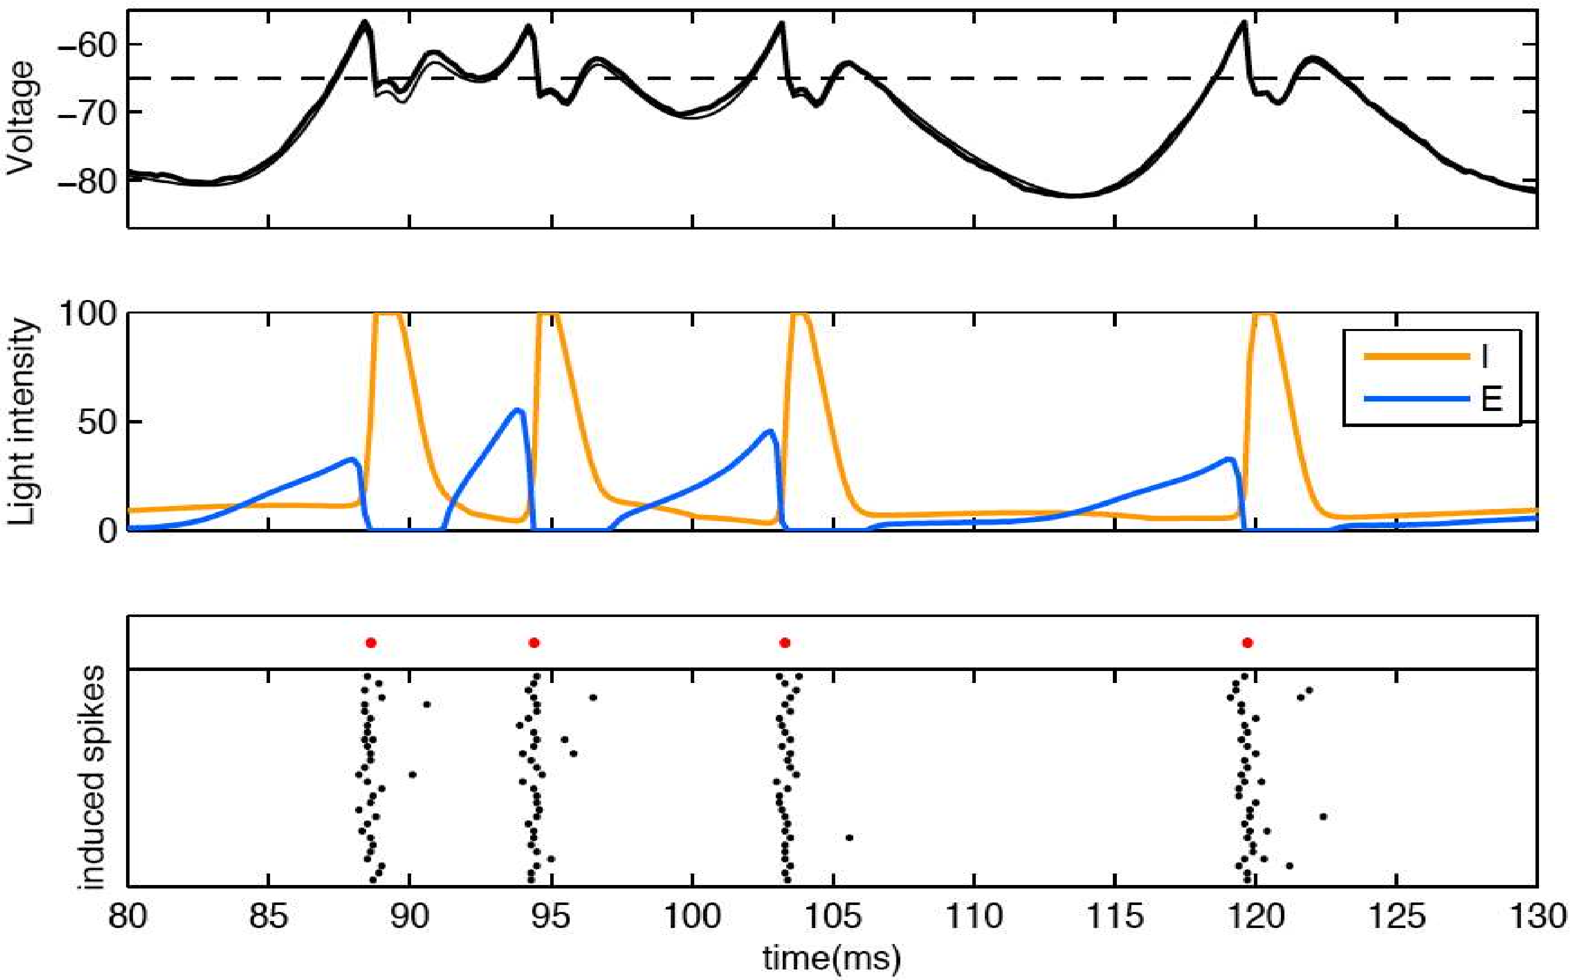

Conclusions

- GLM and state-space approaches provide flexible, powerful methods for answering key questions in neuroscience
- Close relationships between encoding, decoding, and experimental design (Paninski et al., 2007)
- Log-concavity, banded matrix methods make computations very tractable
- Experimental methods progressing rapidly; many new challenges and opportunities for breakthroughs based on statistical ideas

References

Ahmadian, Y., Packer, A., Yuste, R., and Paninski, L. (2010). Fast optimal control of spike trains. *Under review.*

Djurisic, M., Antic, S., Chen, W. R., and Zecevic, D. (2004). Voltage imaging from dendrites of mitral cells: EPSP attenuation and spike trigger zones. *J. Neurosci.*, 24(30):6703–6714.

Huys, Q. and Paninski, L. (2009). Model-based smoothing of, and parameter estimation from, noisy biophysical recordings. *PLOS Computational Biology*, 5:e1000379.

Knopfel, T., Diez-Garcia, J., and Akemann, W. (2006). Optical probing of neuronal circuit dynamics: genetically encoded versus classical fluorescent sensors. *Trends in Neurosciences*, 29:160–166.

Mishchenko, Y., Vogelstein, J., and Paninski, L. (2010). A Bayesian approach for inferring neuronal connectivity from calcium fluorescent imaging data. *Annals of Applied Statistics*, In press.

Paninski, L. (2010). Fast Kalman filtering on quasilinear dendritic trees. *Journal of Computational Neuroscience*, 28:211–28.

Paninski, L., Pillow, J., and Lewi, J. (2007). Statistical models for neural encoding, decoding, and optimal stimulus design. In Cisek, P., Drew, T., and Kalaska, J., editors, *Computational Neuroscience: Progress in Brain Research*. Elsevier.

Vogelstein, J., Packer, A., Machado, T., Sippy, T., Babadi, B., Yuste, R., and Paninski, L. (2010). Fast non-negative deconvolution for spike train inference from population calcium imaging. *J. Neurophys.*, In press.

Vogelstein, J., Watson, B., Packer, A., Jedynak, B., Yuste, R., and Paninski, L. (2009). Model-based optimal inference of spike times and calcium dynamics given noisy and intermittent calcium-fluorescence imaging. *Biophysical Journal*, 97:636–655.