Statistical challenges in neural data analysis

Liam Paninski

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

February 14, 2010

Support: Sloan Fellowship, NIH collaborative research in computational neuroscience, NSF CAREER, McKnight Scholar award.

The neural code

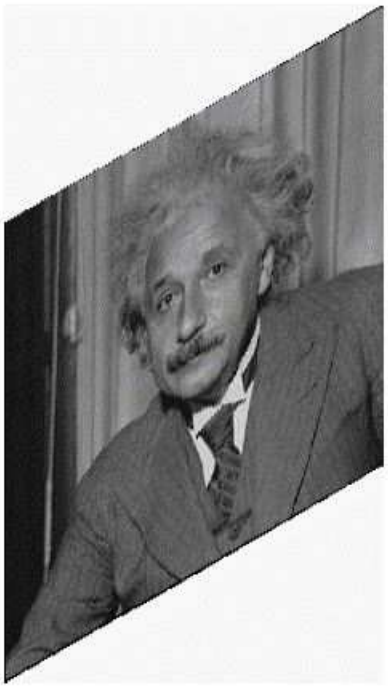

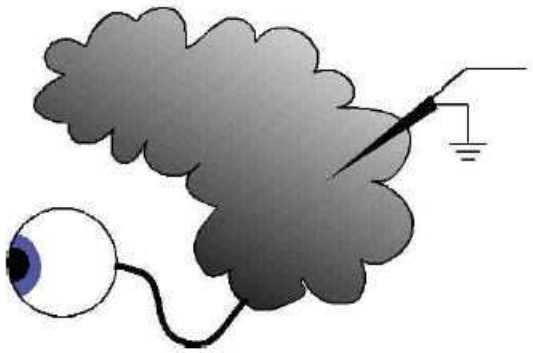

Spike Responses

Basic goal: infer input-output relationship between

- External observables x (sensory stimuli, motor responses...)
- Neural variables y (spike trains, population activity...)

Several levels of neural data analysis

- "Subcellular" level: measurements of intracellular voltage or ionic concentrations (intracellular "patch" electrodes, two-photon imaging, molecular tagging)
- "Circuit" level: electrical activity of single neurons or small groups of isolated neurons (multi-electrode recordings, calcium-sensitive microscopy)
- "Systems" level: blood flow or other indirect measurements of electrical activity in coarsely-defined brain areas (fMRI, EEG, MEG...)

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

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.

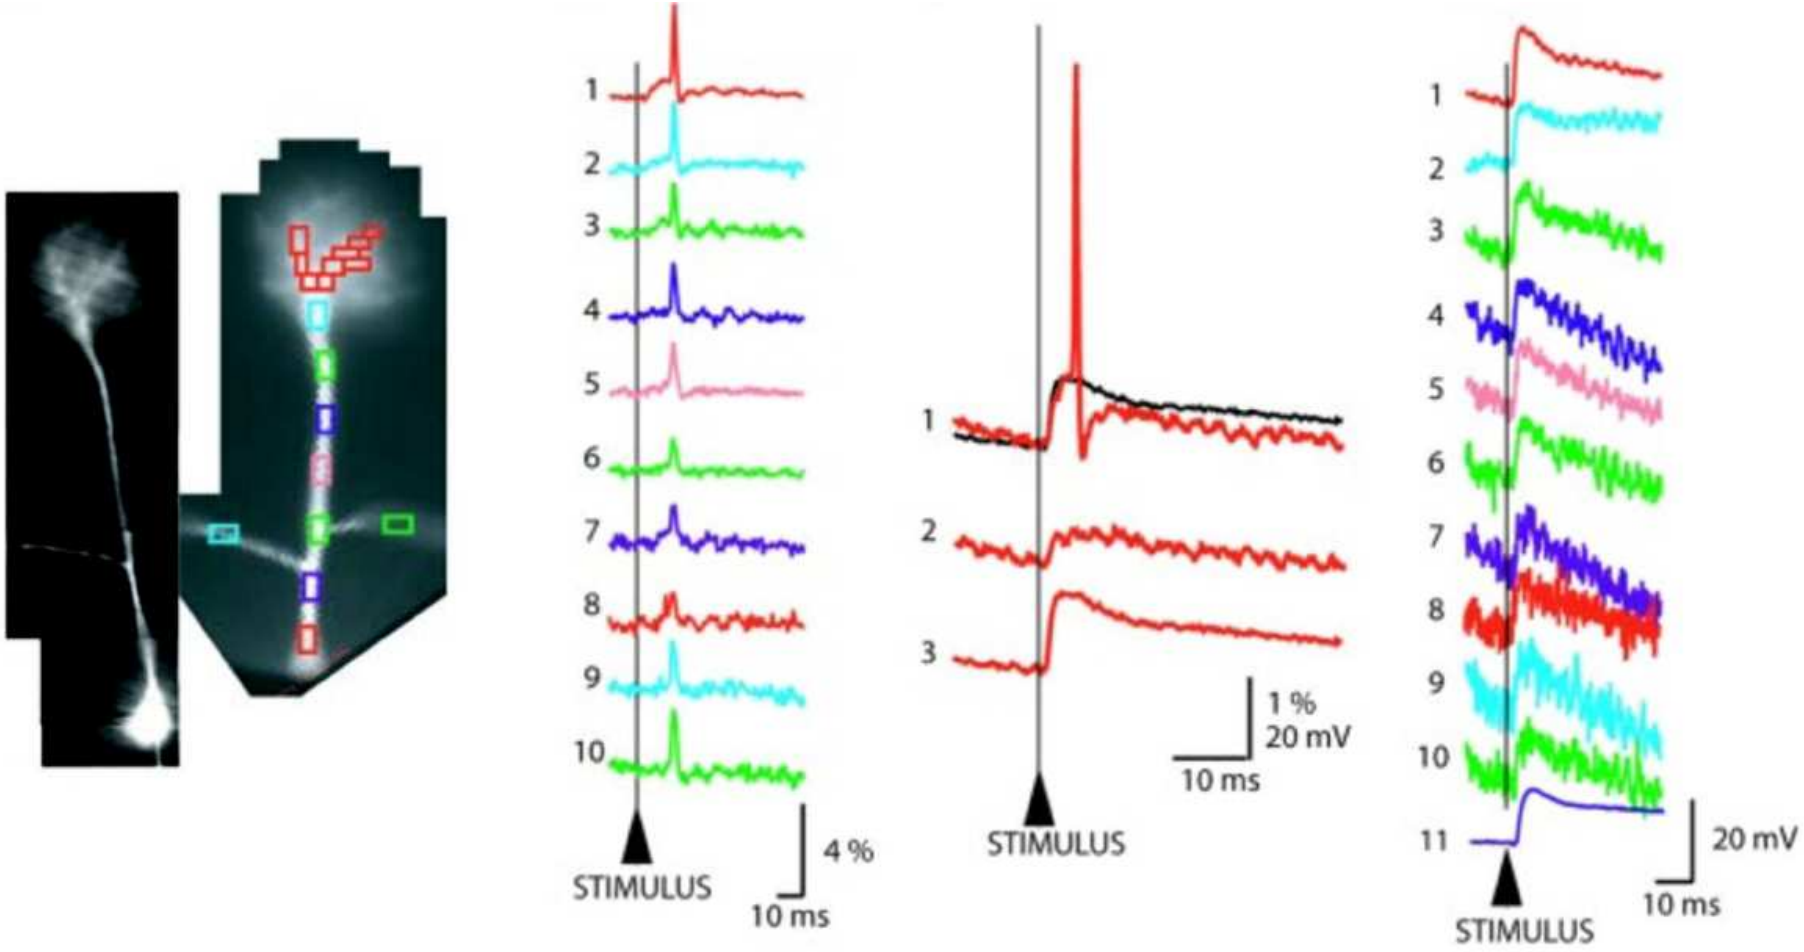

(Djurisic et al., 2004)

Basic paradigm: compartmental models

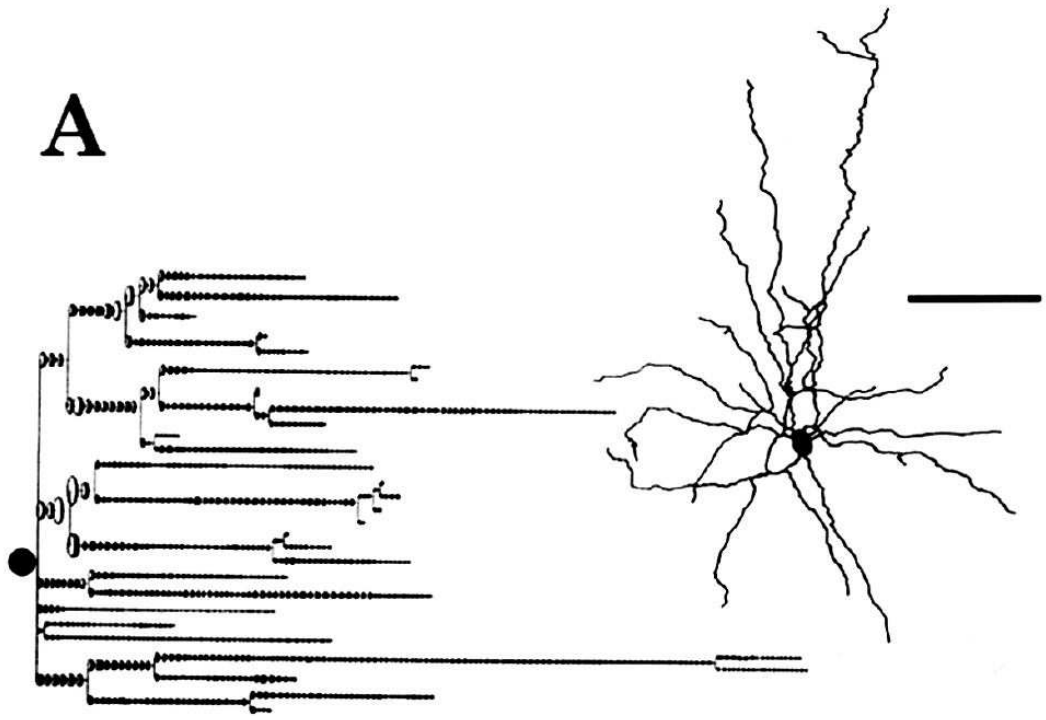

- write neuronal dynamics in terms of equivalent nonlinear, time-varying RC circuits
- leads to a coupled system of stochastic differential equations

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 perturbation of the marginal covariance $C_0 + U_t D_t U_t^T$, with $C_0 = \lim_{t\to\infty} cov(q_t)$.

C_0 is the solution to a Lyapunov equation. It turns out that we can solve linear equations involving C_0 in $O(\dim(q))$ time via Gaussian belief propagation, using the fact that the dendrite is 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, using

$$\begin{aligned} C_t &= [(AC_{t-1}A^T + Q)^{-1} + B_t]^{-1} \\ C_0 + U_t D_t U_t^T &= \left([A(C_0 + U_{t-1}D_{t-1}U_{t-1}^T)A^T + Q]^{-1} + B_t\right)^{-1}, \end{aligned}$$

and can be done in $O(\dim(q))$ time (Paninski, 2009).

Example: inferring voltage from subsampled observations

(Loading low-rank-speckle.mp4)

Applications

- Optimal experimental design: which parts of the neuron should we image? (Submodular optimization; Krause and Guestrin, '07)
- Estimation of biophysical parameters (e.g., membrane channel densities, axial resistance, etc.): reduces to a simple nonnegative regression problem once $V(x,t)$ is known (Huys et al., 2006)
- Detecting location and weights of synaptic input

Application: synaptic locations/weights

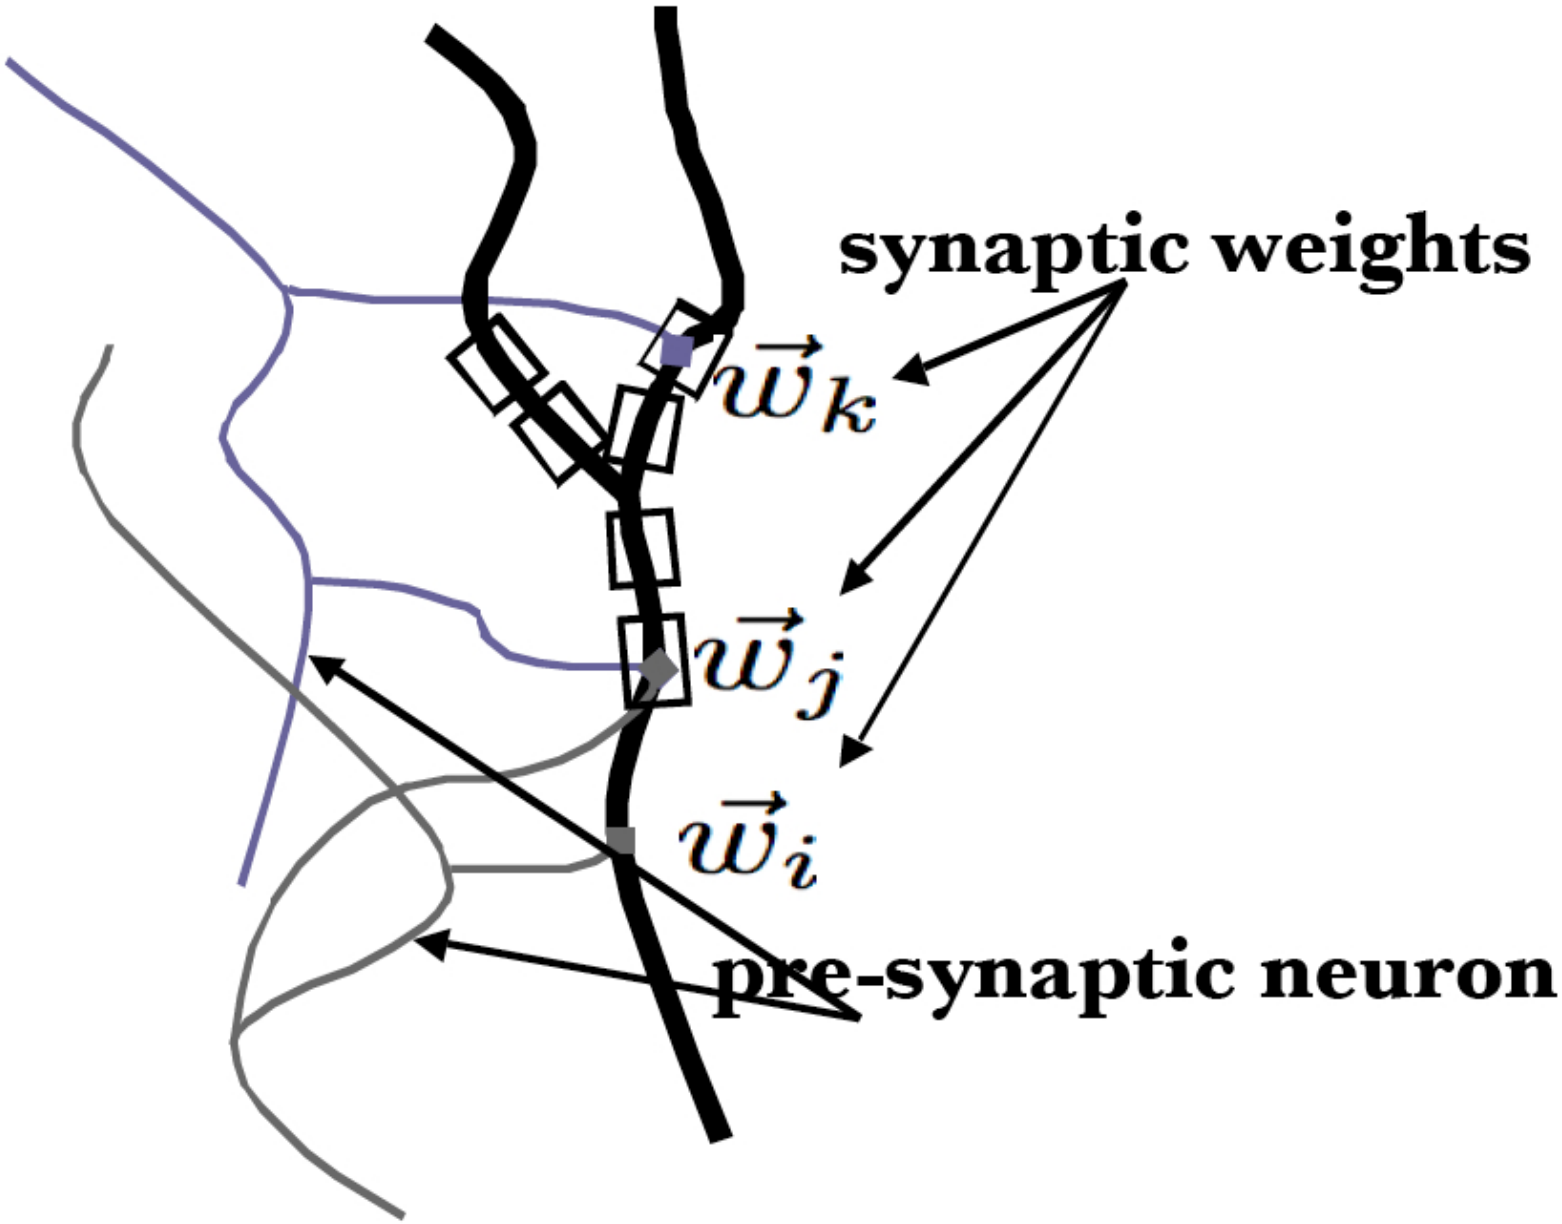

Application: synaptic locations/weights

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).

Detecting synapses

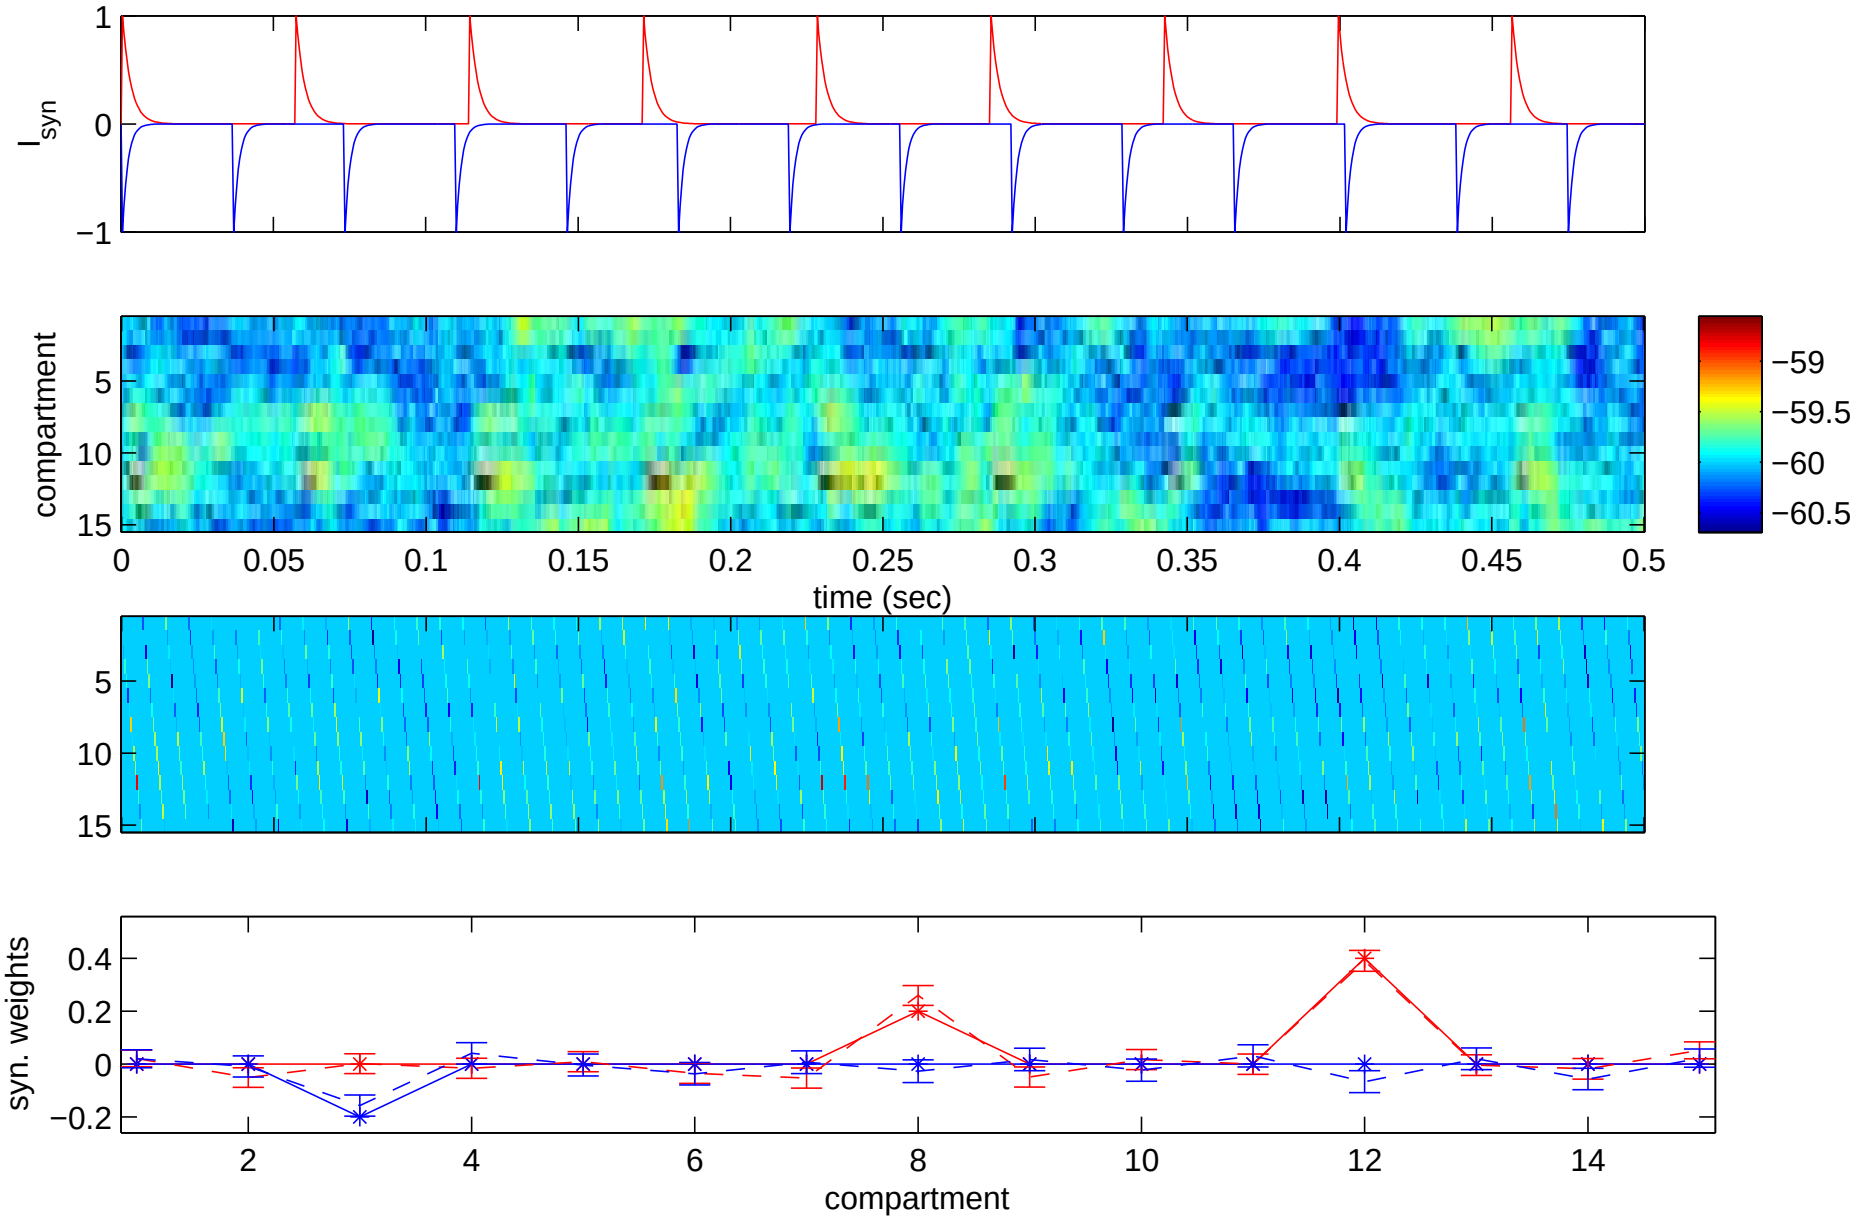

(Paninski and Ferreira, 2008; Paninski et al., 2010)

Part 2: optimal decoding of spike train data

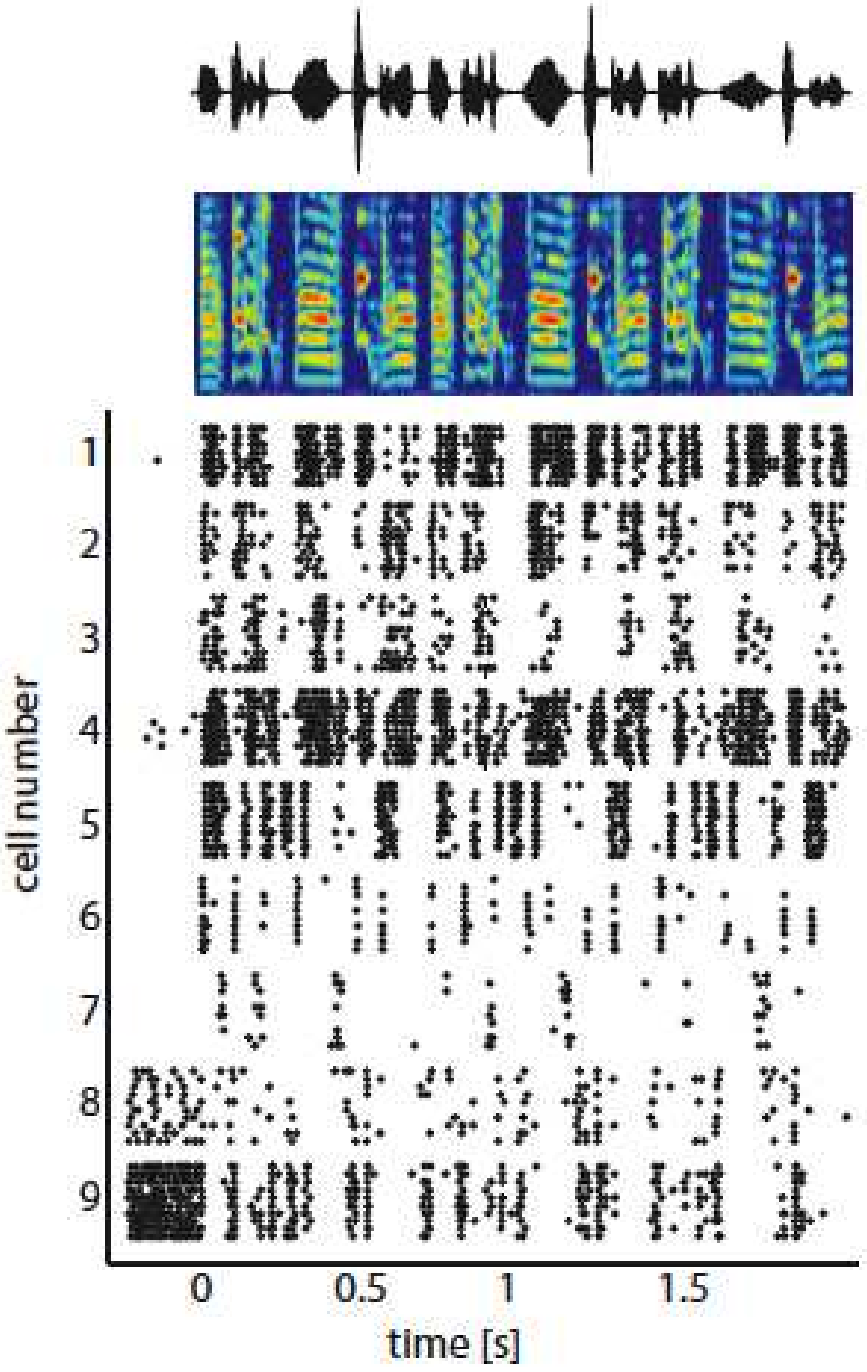

Semiparametric GLM

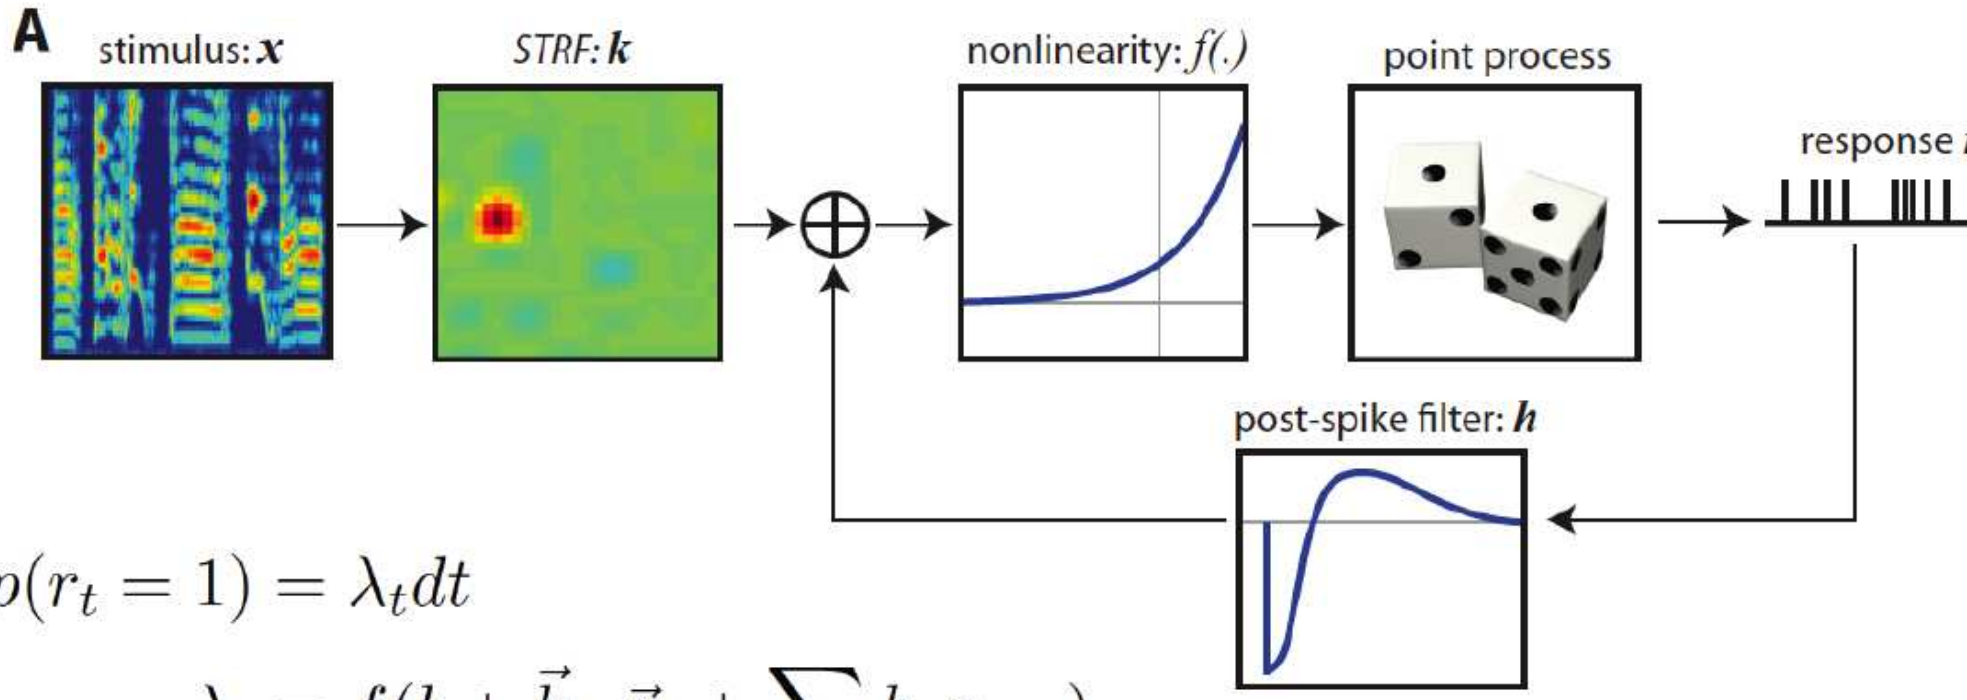

$$p(r_t = 1) = \lambda_t dt$$

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

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.

MAP stimulus decoding

It is reasonable to estimate the song X that led to a response R via the MAP

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

(Note that X is very high-dimensional!) For this model, we have:

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

Two basic observations:

- If $\log p(X)$ is concave, then so is $\log p(X|R)$, since each $\log p(r_t|X, Y_{\dots,t-1})$ is.
- Hessian H of $\log p(R|X)$ w.r.t. X is banded: each $p(r_t|X, R_{\dots,t-1})$ 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.

— similar speedups available in constrained cases (Paninski et al., 2010), and in MCMC setting (e.g., fast hybrid Monte Carlo methods (Ahmadian et al., 2010)).

Application: fast optimal decoding

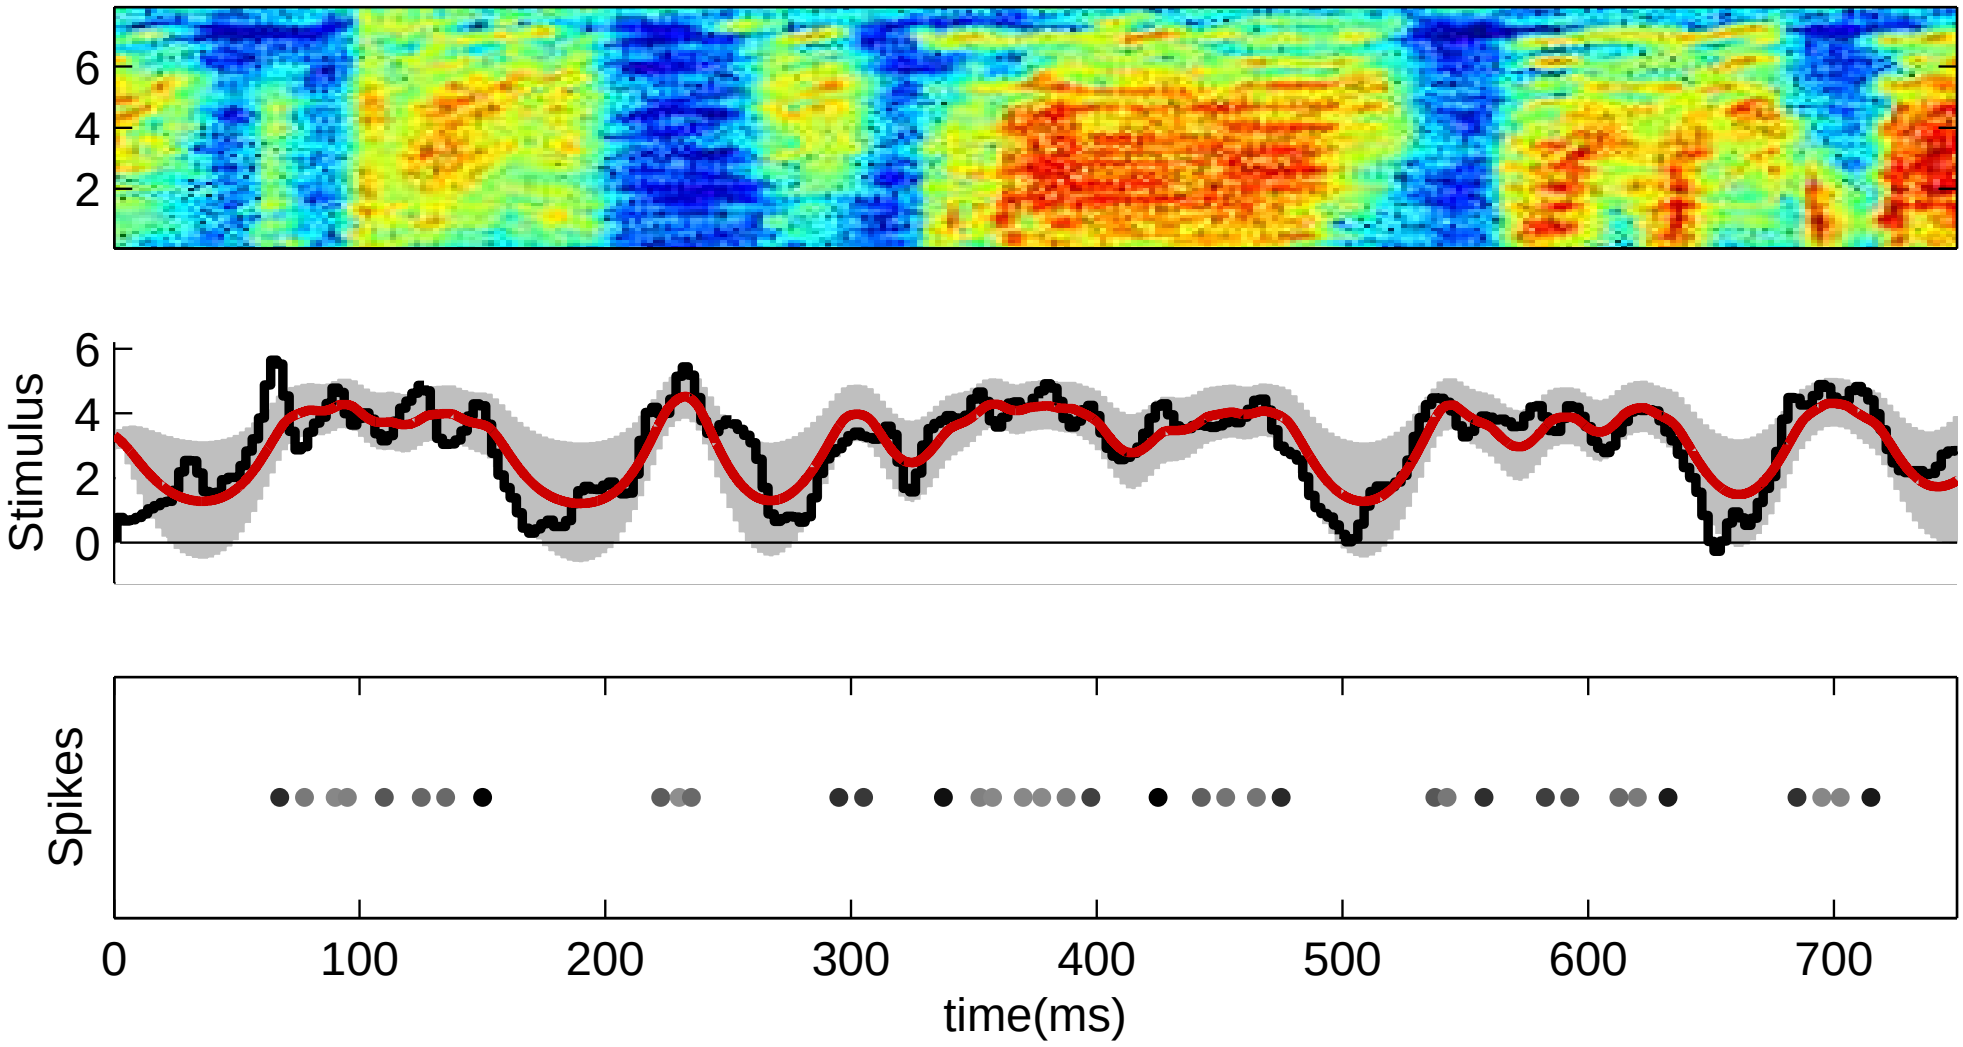

Decoding a full song

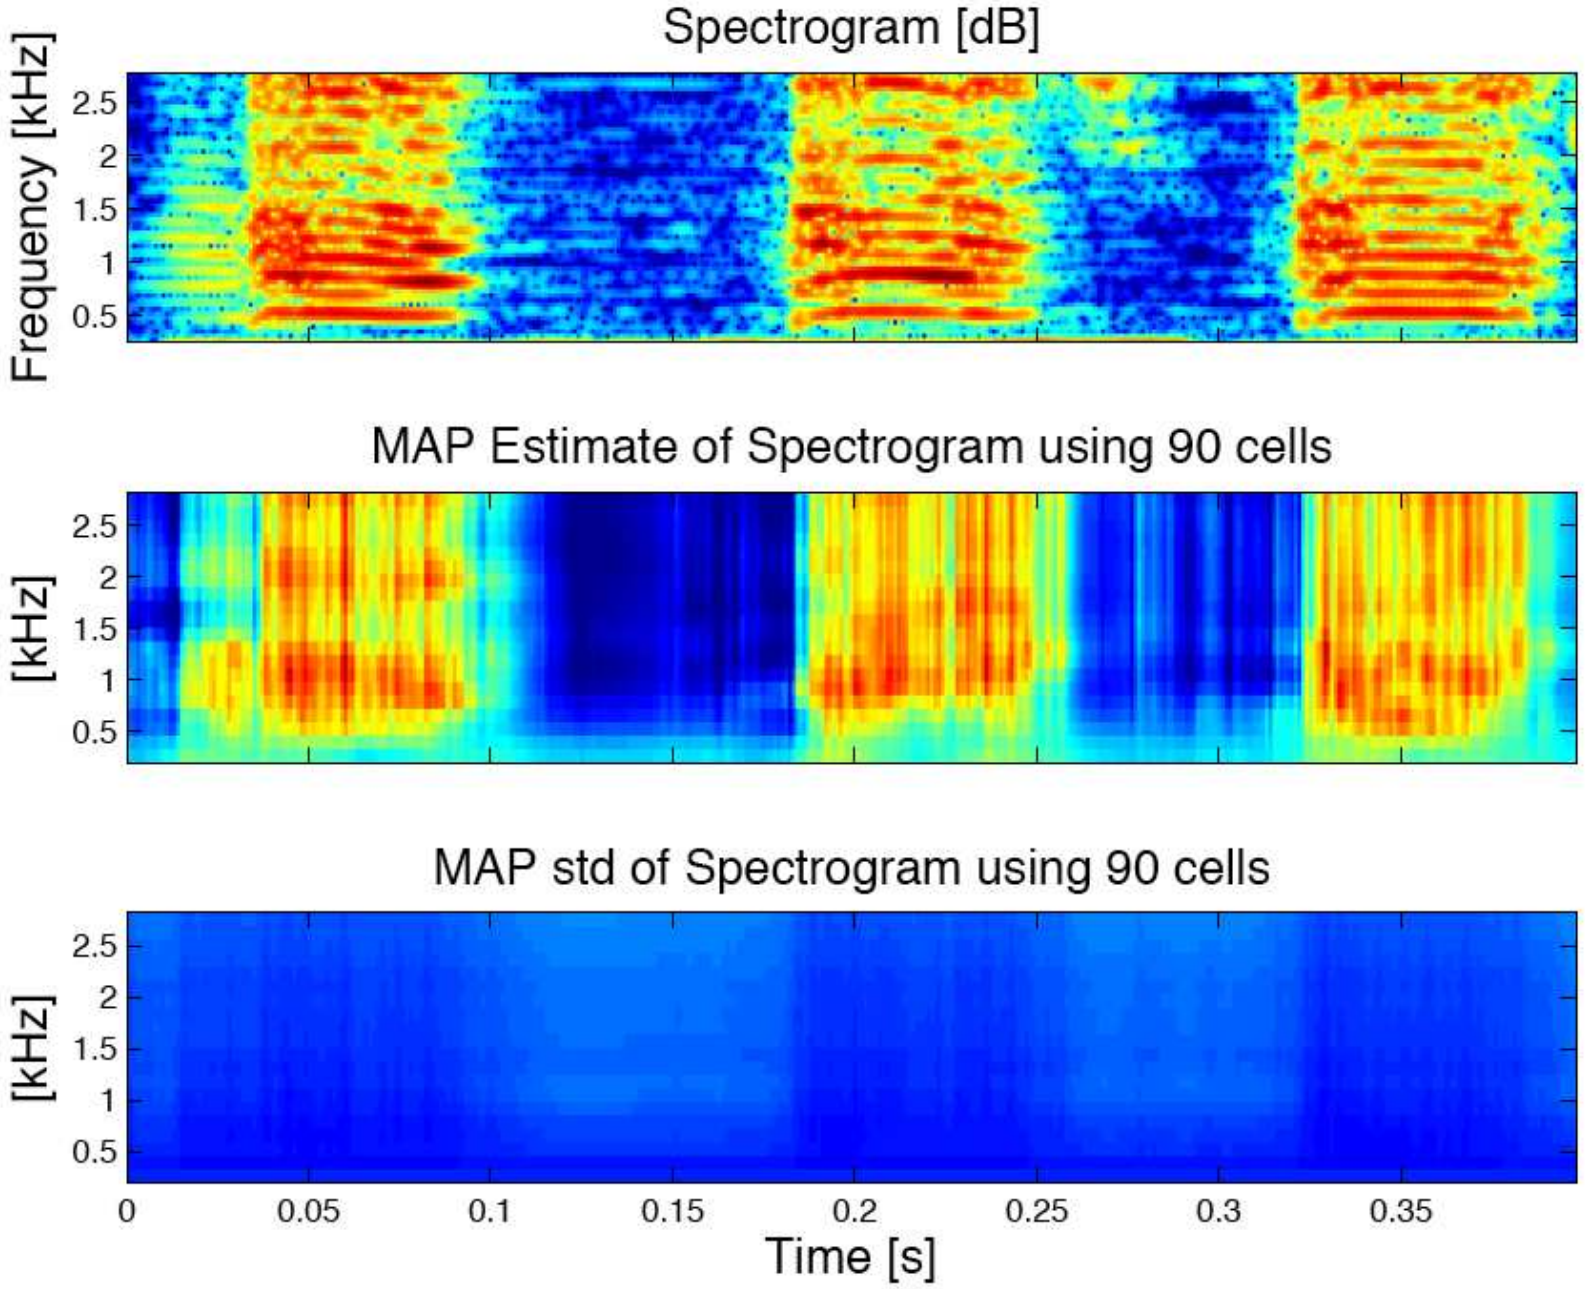

(Ramirez et al 2010)

Application: optimal stimulus design

Idea: we have full control over the stimuli we present. Can we choose stimuli $\vec{x}_t$ to maximize the informativeness of each trial?

— More quantitatively, optimize $I(r_t; \theta | \vec{x}_t)$ with respect to $\vec{x}_t$.

Maximizing $I(r_t; \theta; \vec{x}_t) \implies$ minimizing uncertainty about θ.

In general, very hard to do: high-d integration over θ to compute $I(r_t; \theta | \vec{x}_t)$, high-d optimization to select best $\vec{x}_t$.

GLM setting + low-rank matrix methods make this surprisingly tractable: $O(\dim(\theta)^2)$ computation (Lewi et al., 2009).

Application to songbird data: choosing an optimal stimulus sequence

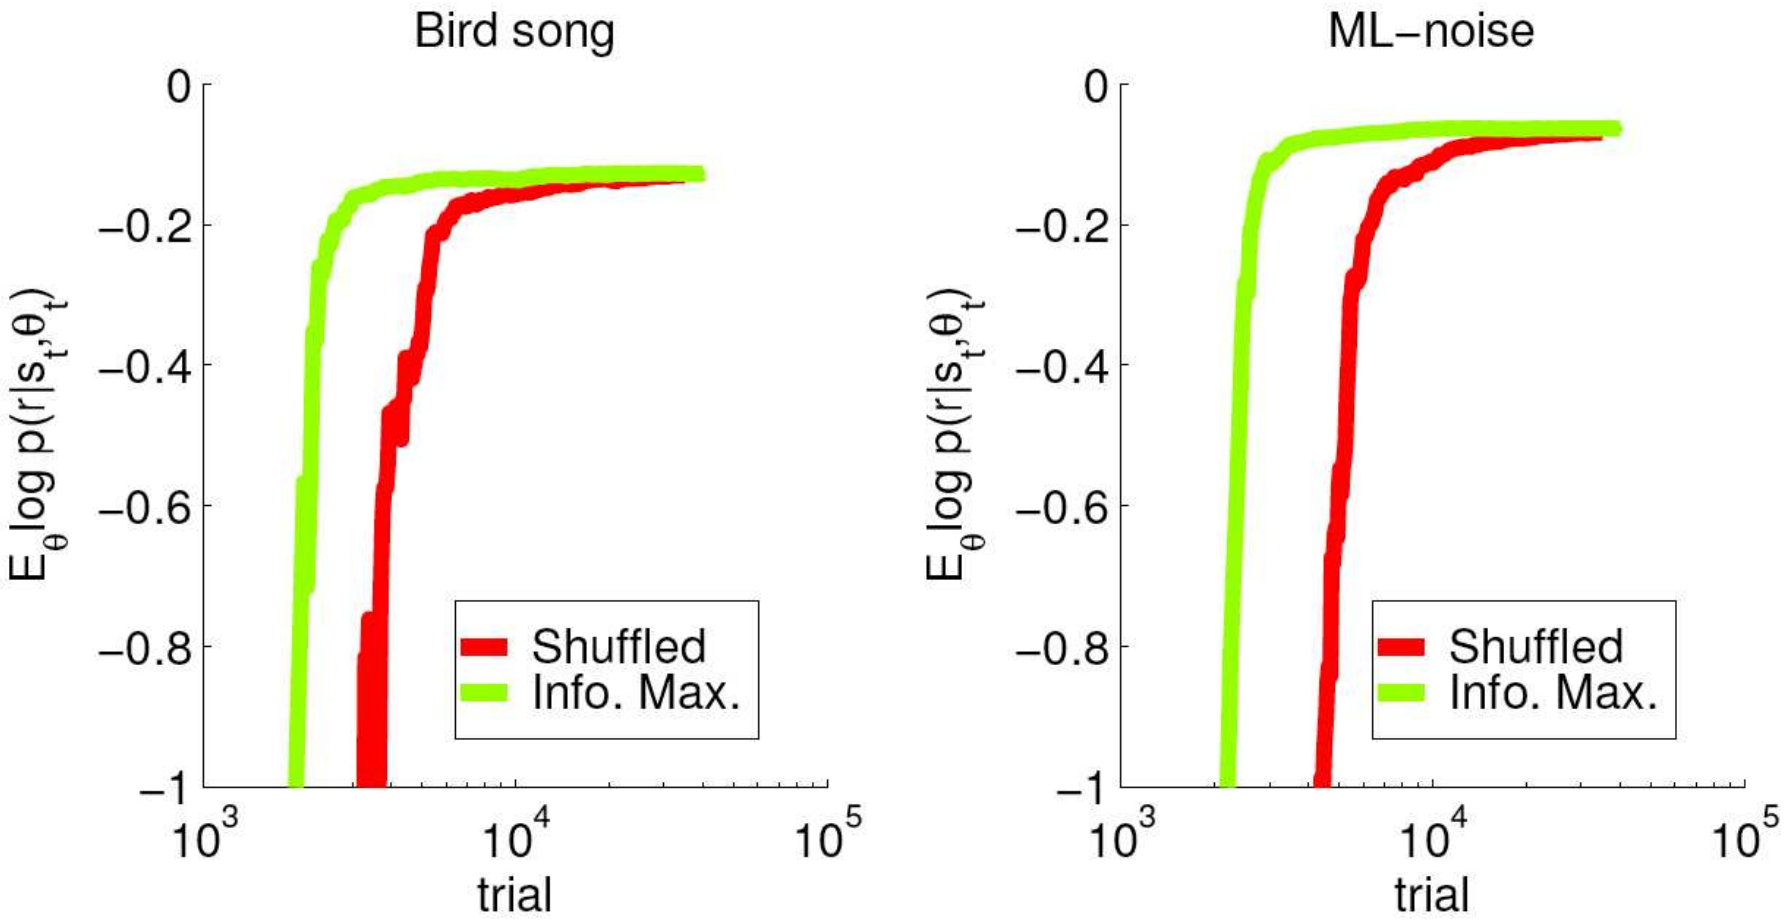

— infomax speeds convergence by a factor of three or more.

Part 3: circuit inference

Record large-scale calcium movie

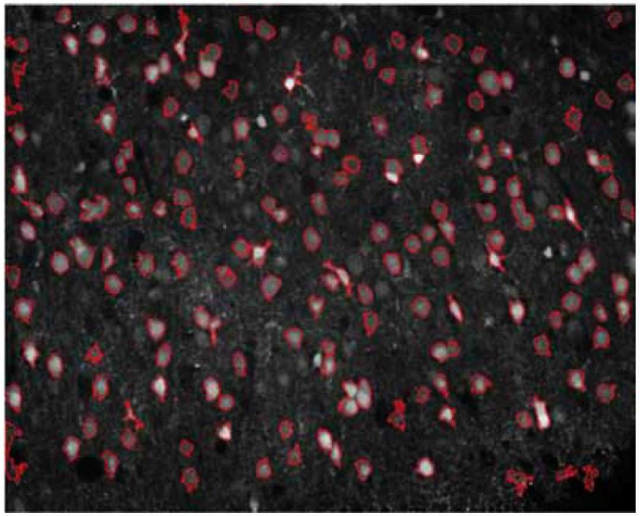

Aim 1: Extract spike times

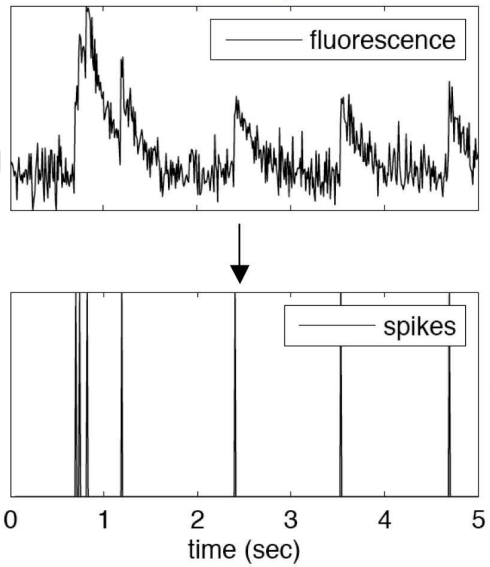

Aim 2: Estimate network model

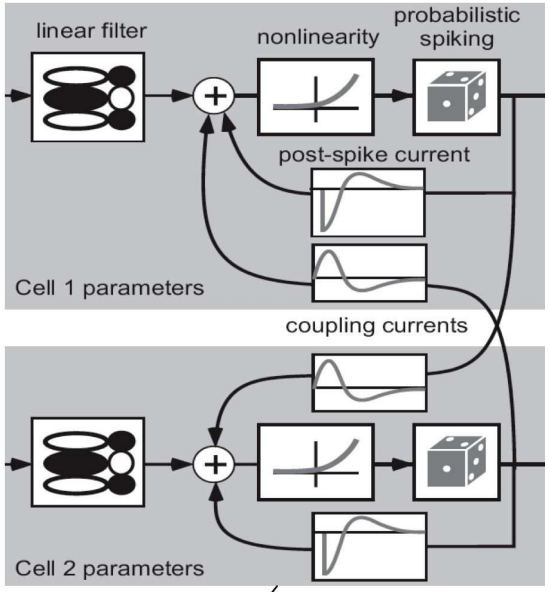

Aim 3: Check results via photostimulation

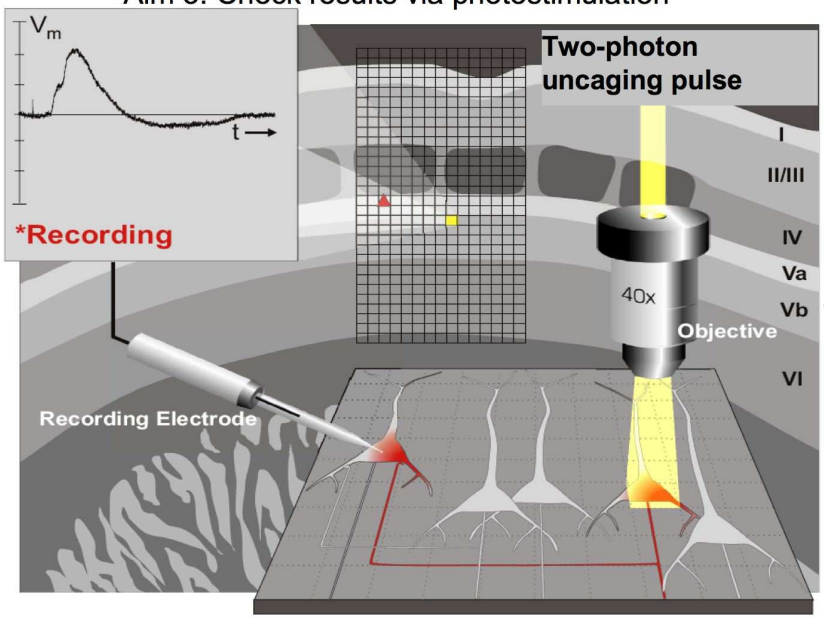

Inferred network model

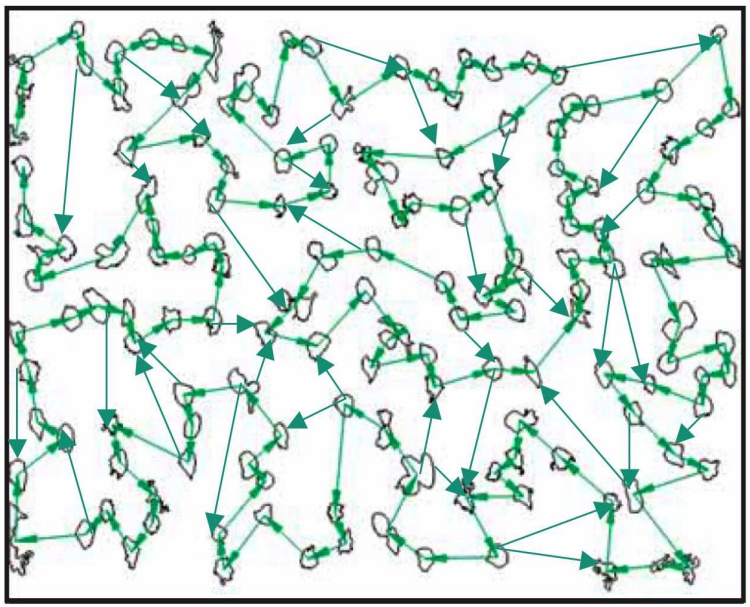

Challenge: slow, noisy calcium data

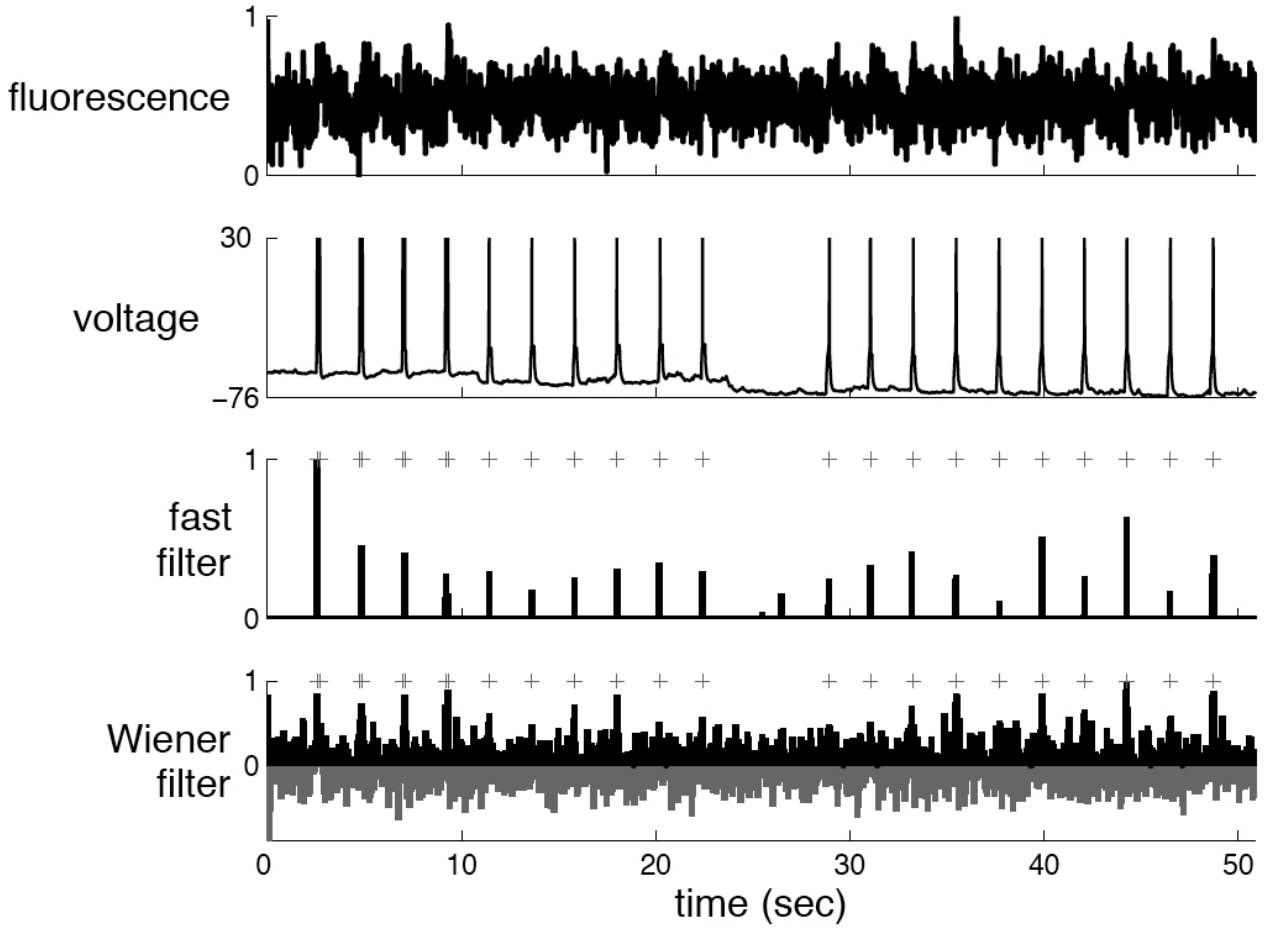

First-order model:

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

— $\tau \approx 100$ ms; nonnegative deconvolution problem. Can be solved by $O(T)$ relaxed constrained interior-point optimization (Vogelstein et al., 2008) or sequential Monte Carlo (Vogelstein et al., 2009).

Monte Carlo EM approach

Given the spike times in the network, L_1-penalized likelihood optimization is easy. But we only have noisy calcium observations Y; true spike times are hidden variables. Thus an EM approach is natural.

- E step: sample spike train responses R from $p(R|Y, \theta)$
- M step: given sampled spike trains, perform L_1-penalized likelihood optimization to update parameters θ.

E step is hard part here. Use the fact that neurons interact fairly weakly; thus we need to sample from a collection of weakly-interacting Markov chains, via Metropolis-within-blockwise-Gibbs forward-backward methods (Neal et al., 2003).

Simulated circuit inference

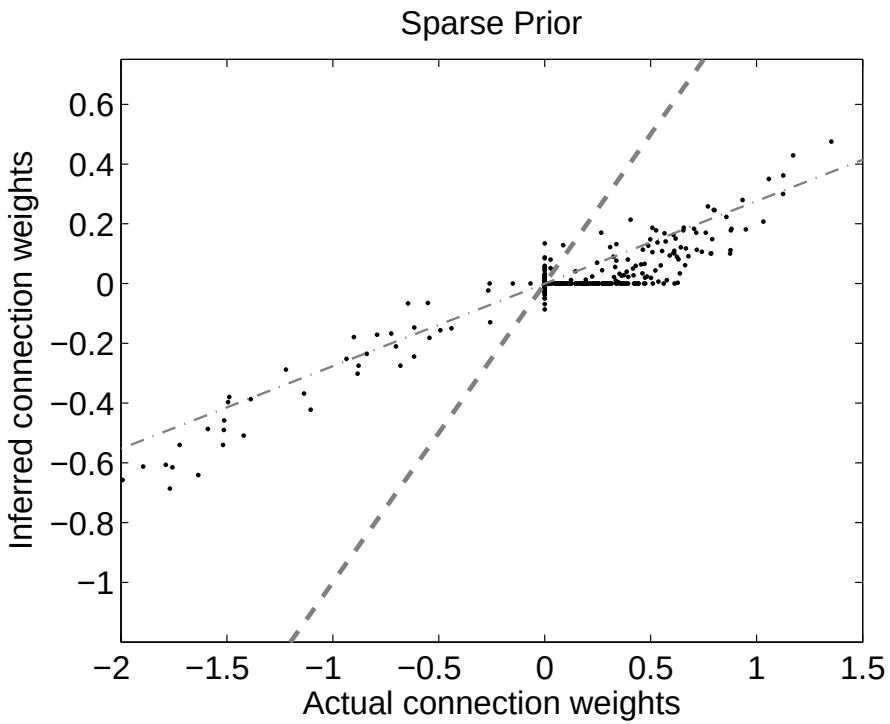

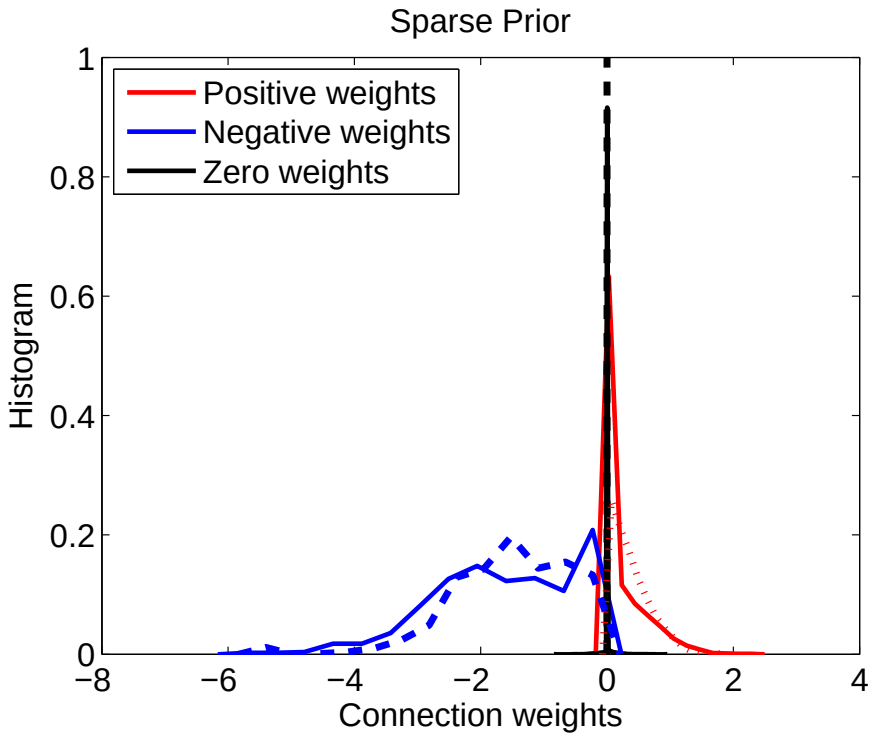

— Connections are inferred with the correct sign in conductance-based integrate-and-fire networks with biologically plausible connectivity matrices (Mishchencko et al., 2009).

Good news: connections are inferred with the correct sign. But process is slow; current work focusing on improved sampling methods (exploiting hybrid forward-backward blockwise-Gibbs approach).

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:

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

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.

Optimal electrical control of spike timing

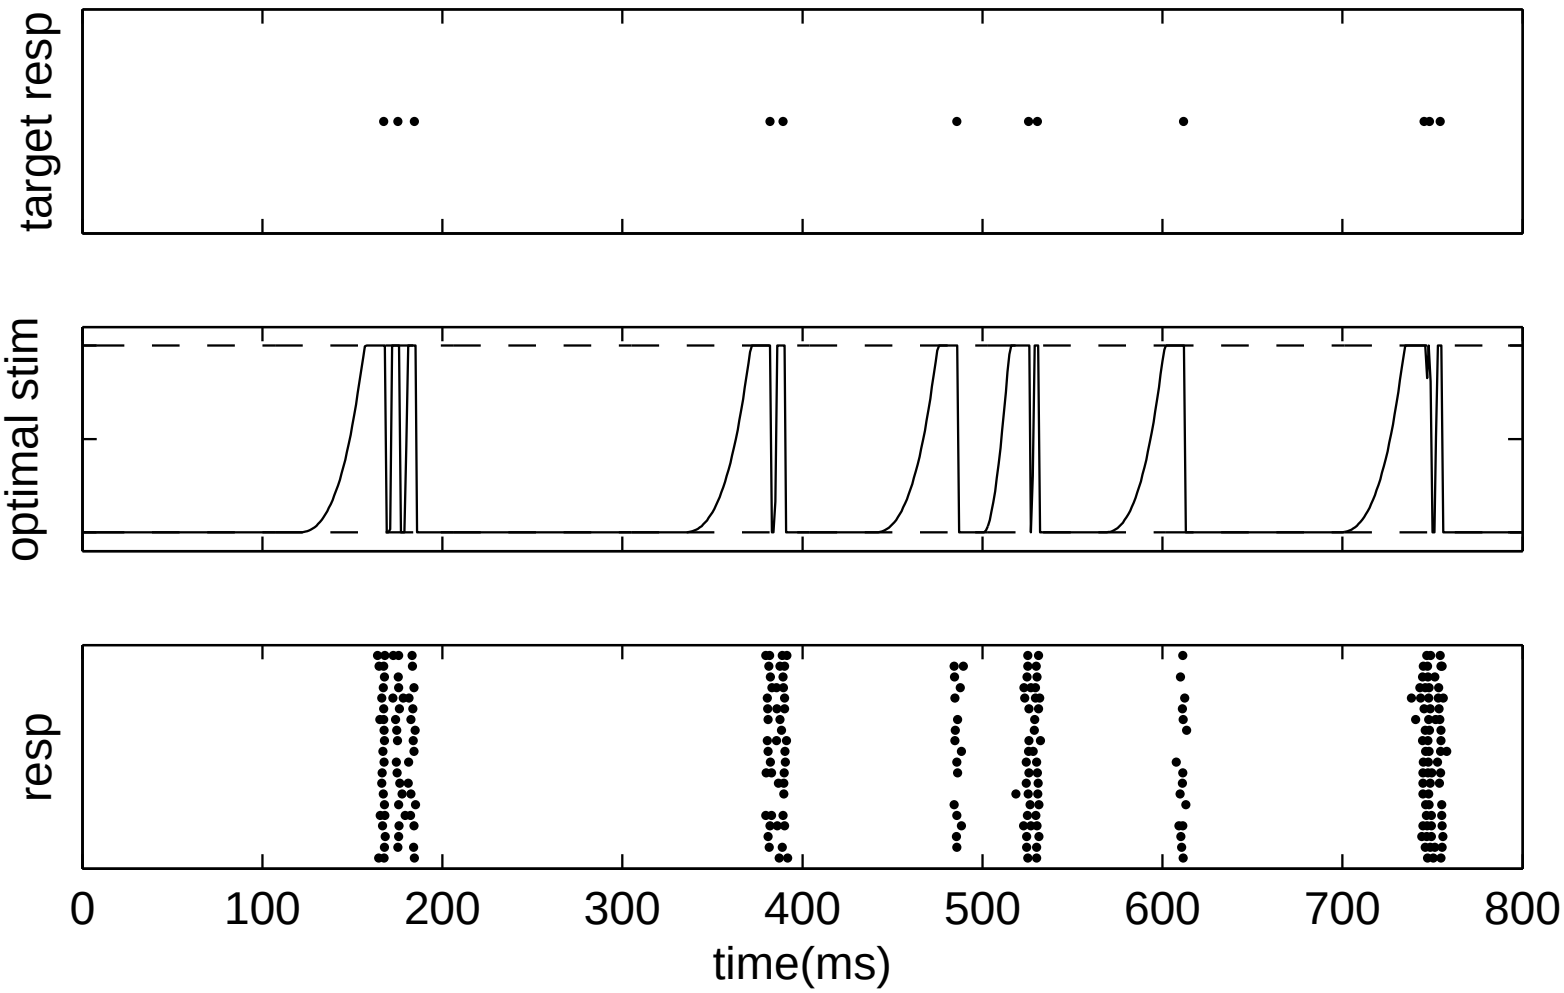

Extension to optical stimulation methods is straightforward (Ahmadian and Paninski, 2010).

Example: intracellular control of spike timing

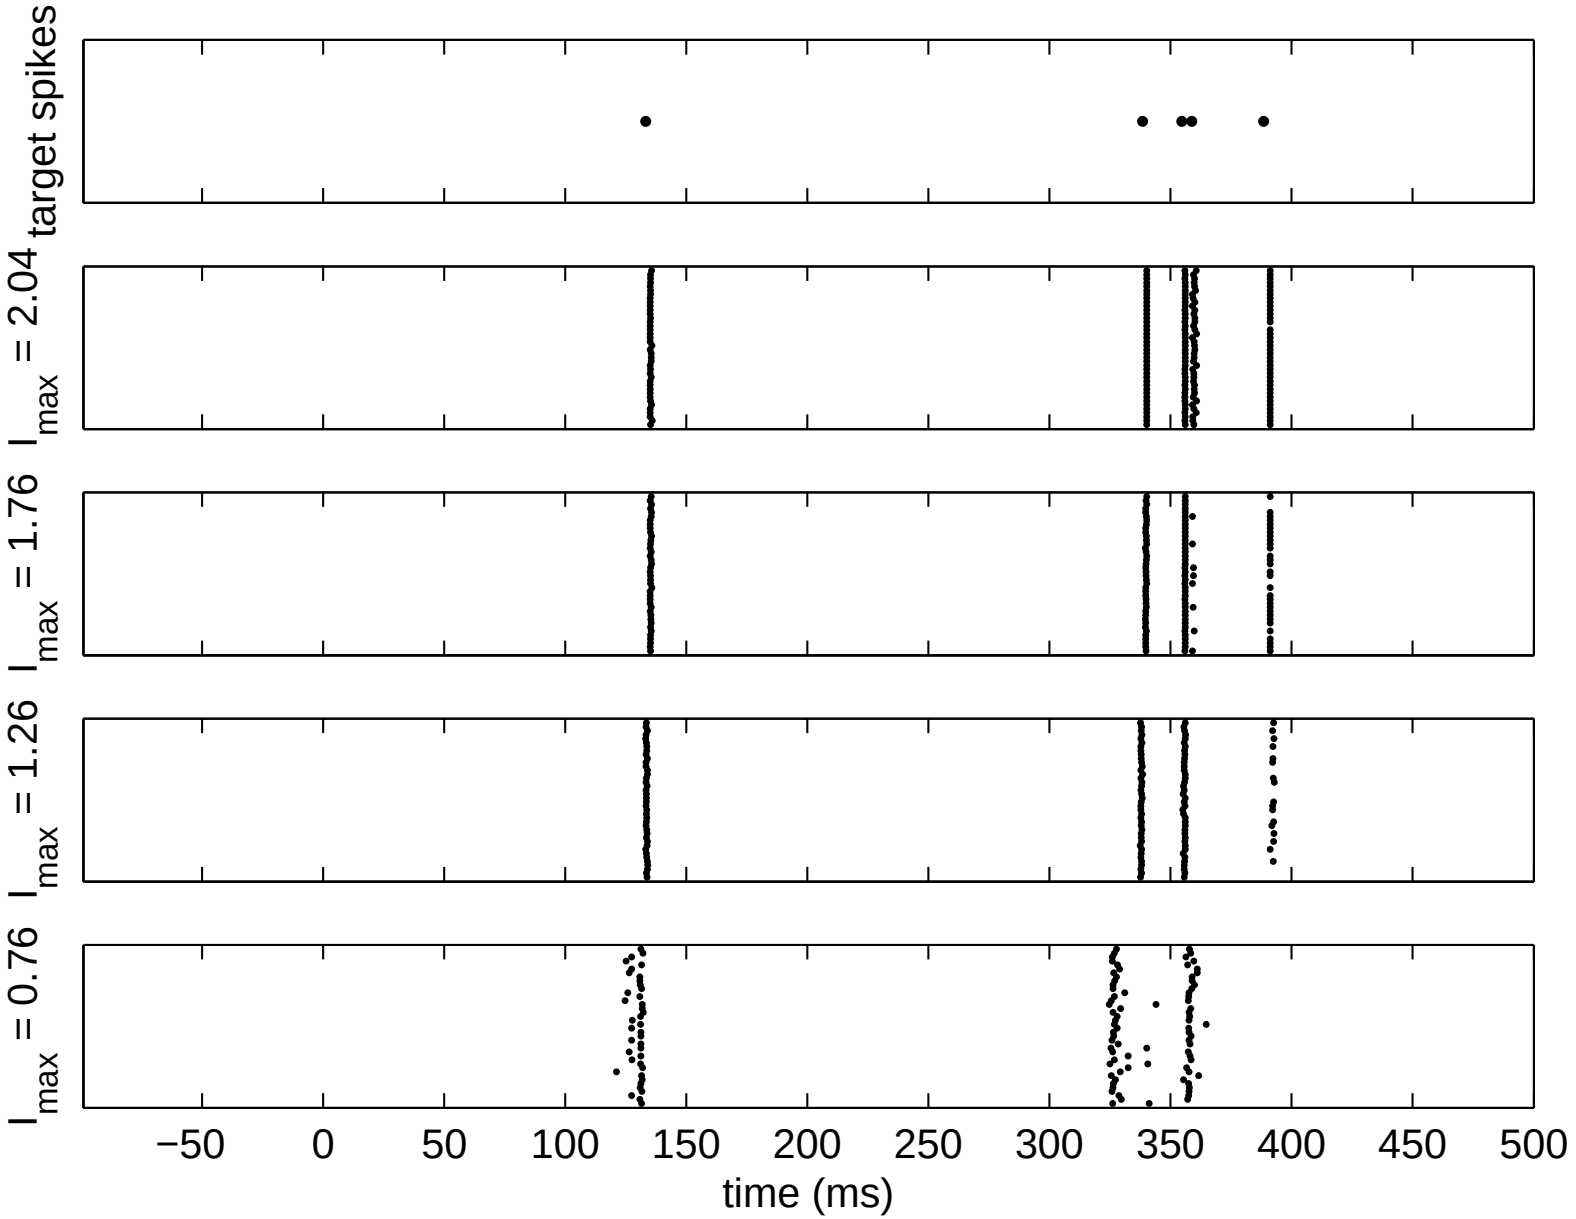

(Ahmadian et al 2010)

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. and Paninski, L. (2010). Optimal control of spike trains. *In preparation.*

Ahmadian, Y., Pillow, J., and Paninski, L. (2010). Efficient Markov Chain Monte Carlo methods for decoding population spike trains. *Under review, Neural Computation.*

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., Ahrens, M., and Paninski, L. (2006). Efficient estimation of detailed single-neuron models. *Journal of Neurophysiology*, 96:872–890.

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

Lewi, J., Butera, R., and Paninski, L. (2009). Sequential optimal design of neurophysiology experiments. *Neural Computation*, 21:619–687.

Neal, R., Beal, M., and Roweis, S. (2003). Inferring state sequences for non-linear systems with embedded hidden Markov models. *NIPS*, 16.

Paninski, L. (2009). Fast Kalman filtering on dendritic trees. *In progress.*

Paninski, L., Ahmadian, Y., Ferreira, D., Koyama, S., Rahnama, K., Vidne, M., Vogelstein, J., and Wu, W. (2010). A new look at state-space models for neural data. *Journal of Computational Neuroscience*, In press.

Paninski, L. and Ferreira, D. (2008). State-space methods for inferring synaptic inputs and weights. *COSYNE.*

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., Babadi, B., Watson, B., Yuste, R., and Paninski, L. (2008). Fast nonnegative deconvolution via tridiagonal interior-point methods, applied to calcium fluorescence data. *Statistical analysis of neural data (SAND) conference.*

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*, In press;
http://www.stat.columbia.edu/∼liam/research/abstracts/vogelstein-bj08-abs.html.