Fast statistical methods for mapping synaptic connectivity on dendrites

Liam Paninski

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

June 3, 2012

Co-PI: R. Yuste. Joint work with A. Pakman, J. Huggins, E. Pnevmatikakis.

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)

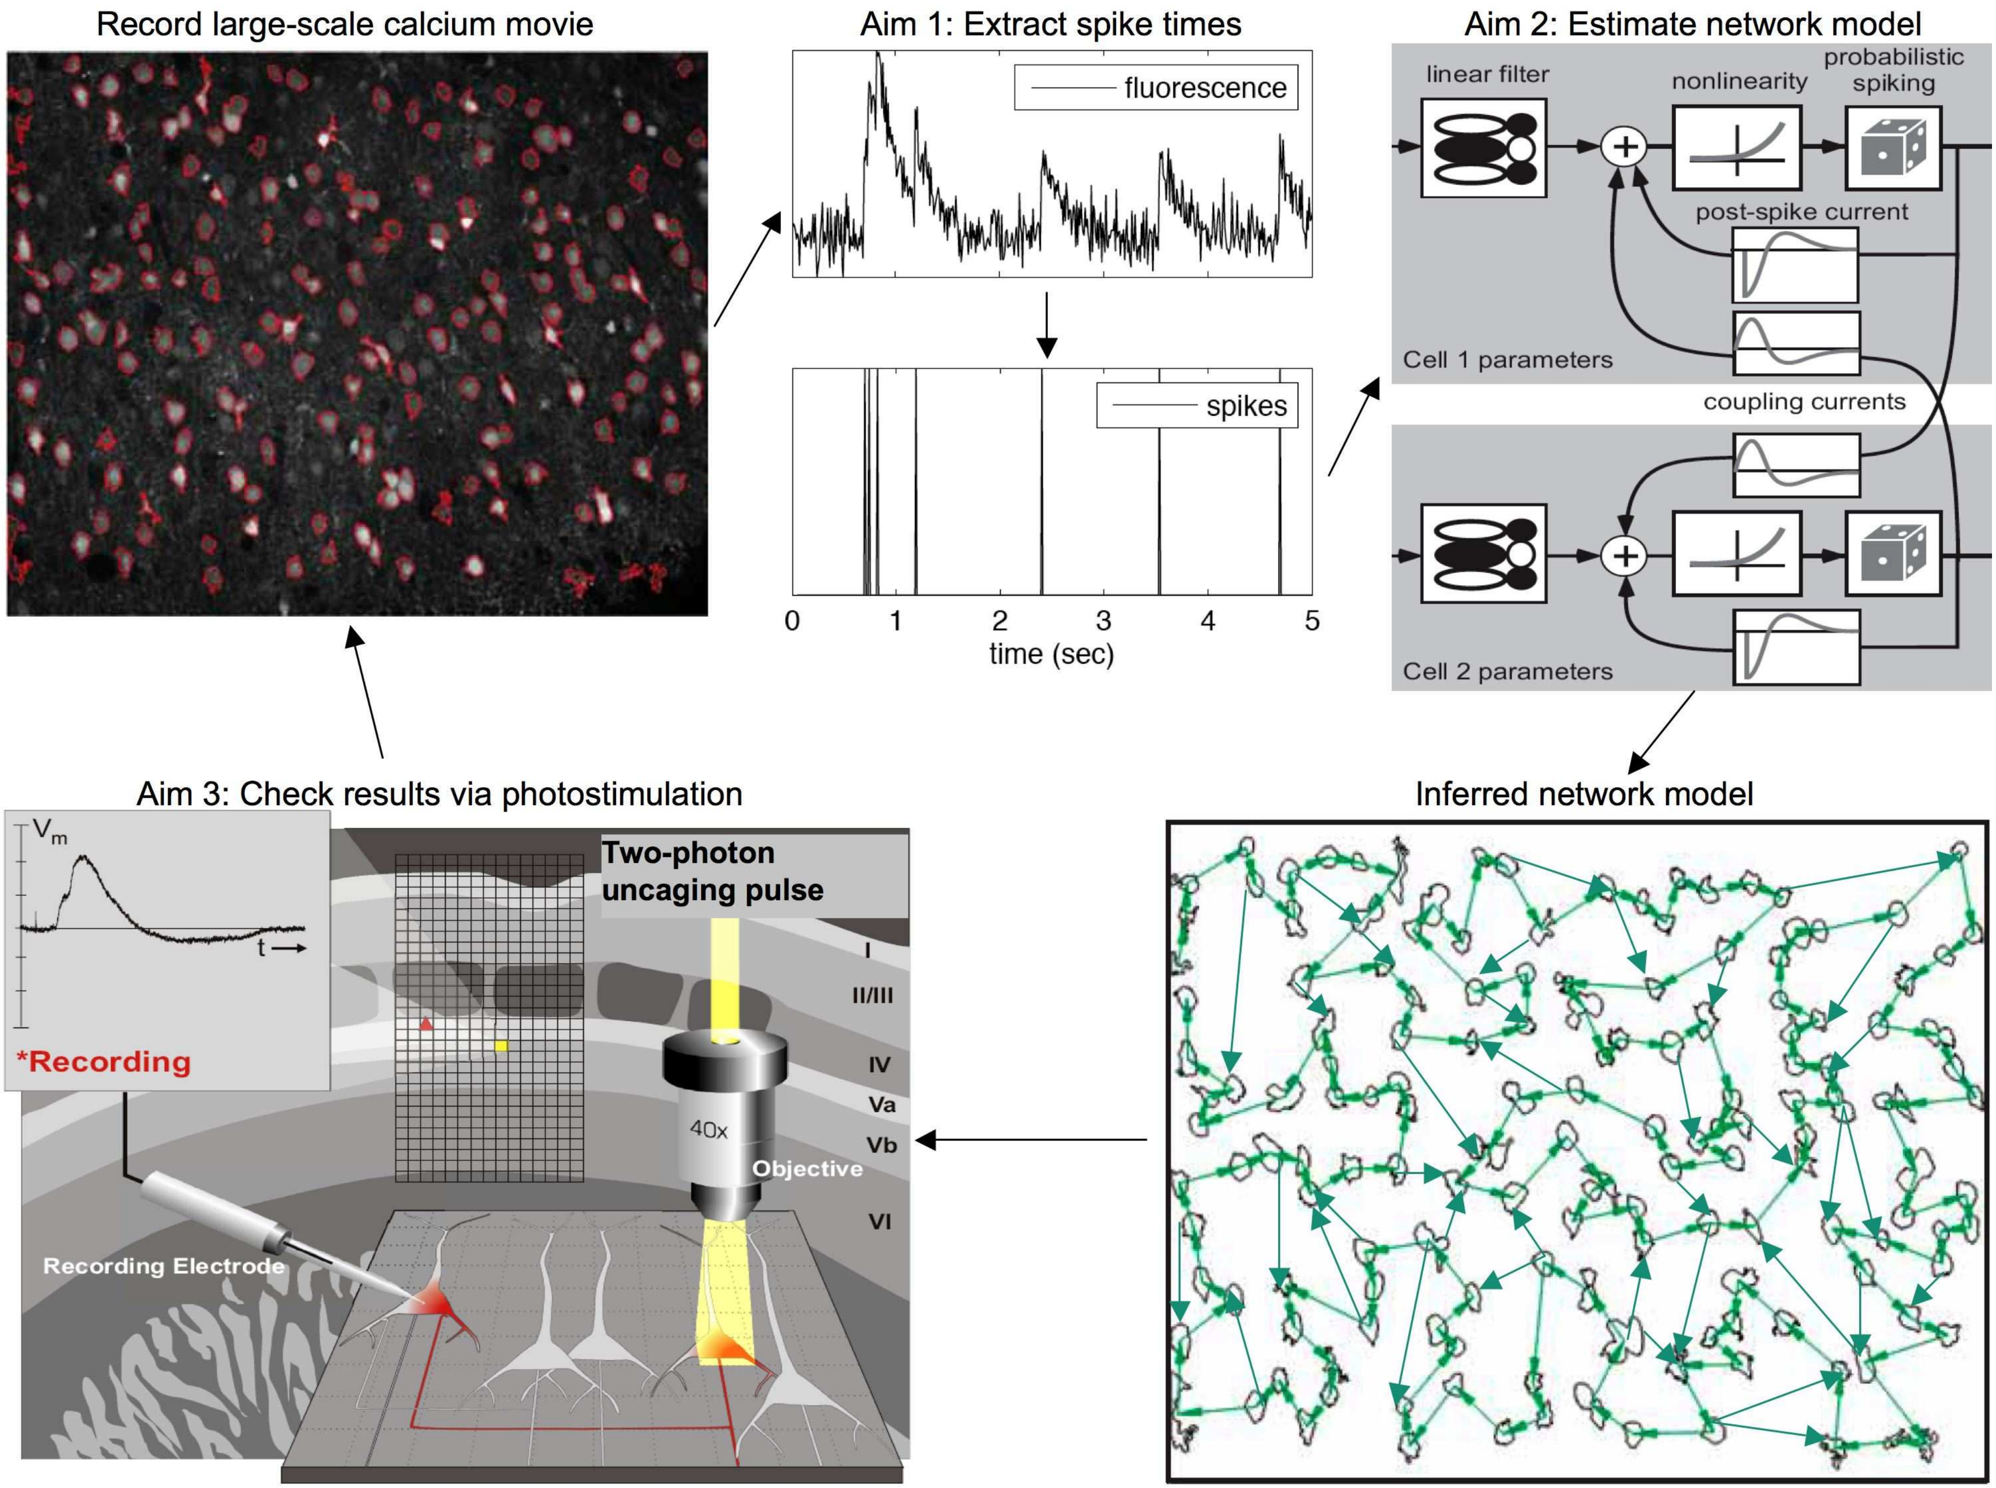

Record large-scale calcium movie
Aim 1: Extract spike times
fluorescence
spikes
0
1
2
3
4
5
time (sec)
Aim 2: Estimate network model
linear filter
nonlinearity
probabilistic spiking
post-spike current
Cell 1 parameters
coupling currents
Cell 2 parameters
Inferred network model
Aim 3: Check results via photostimulation
V_m
t
*Recording
Two-photon uncaging pulse
I
II/III
IV
Va
Vb
VI
40x
Objective
Recording Electrode

Mapping connectivity at the dendritic level

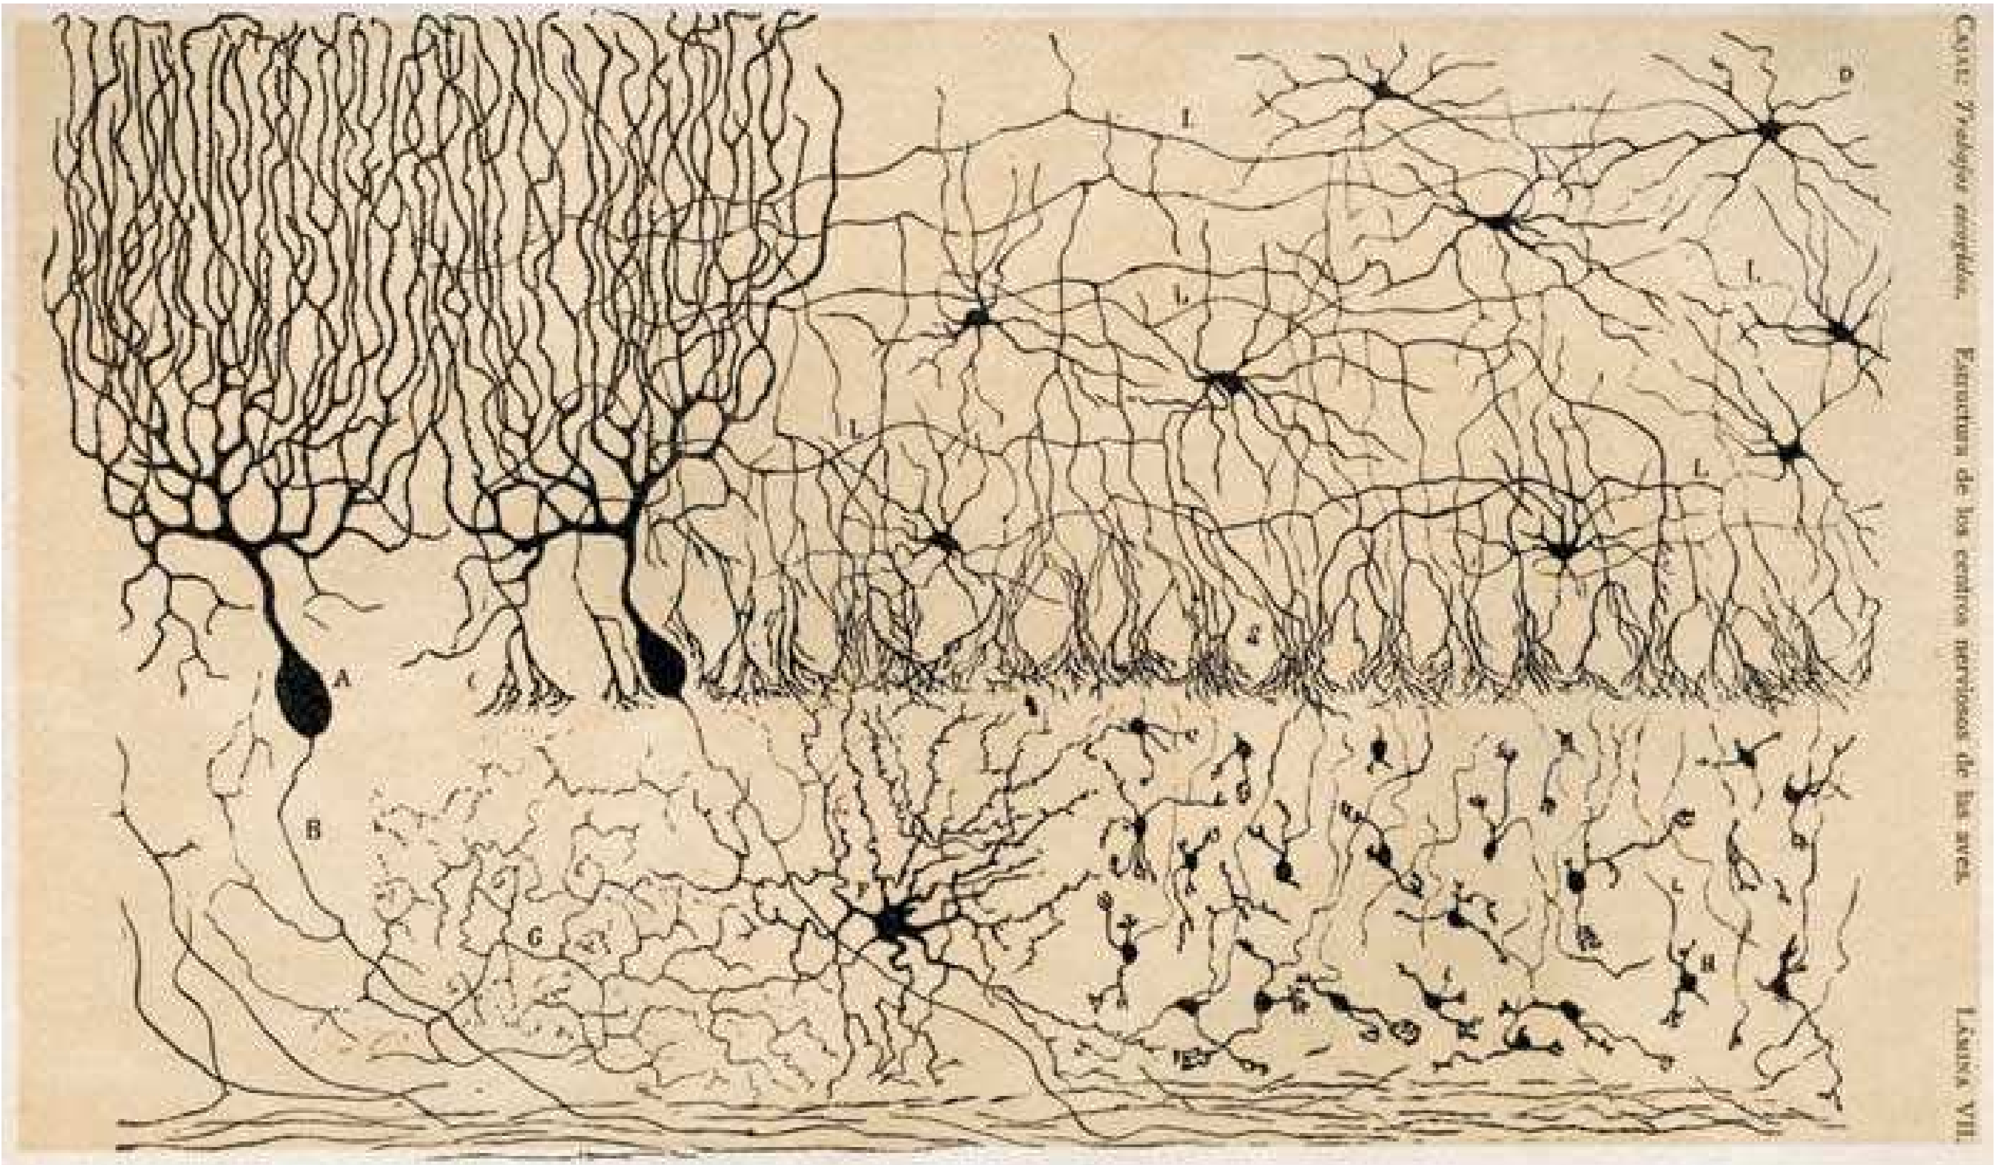

Ramon y Cajal, 1888.

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.

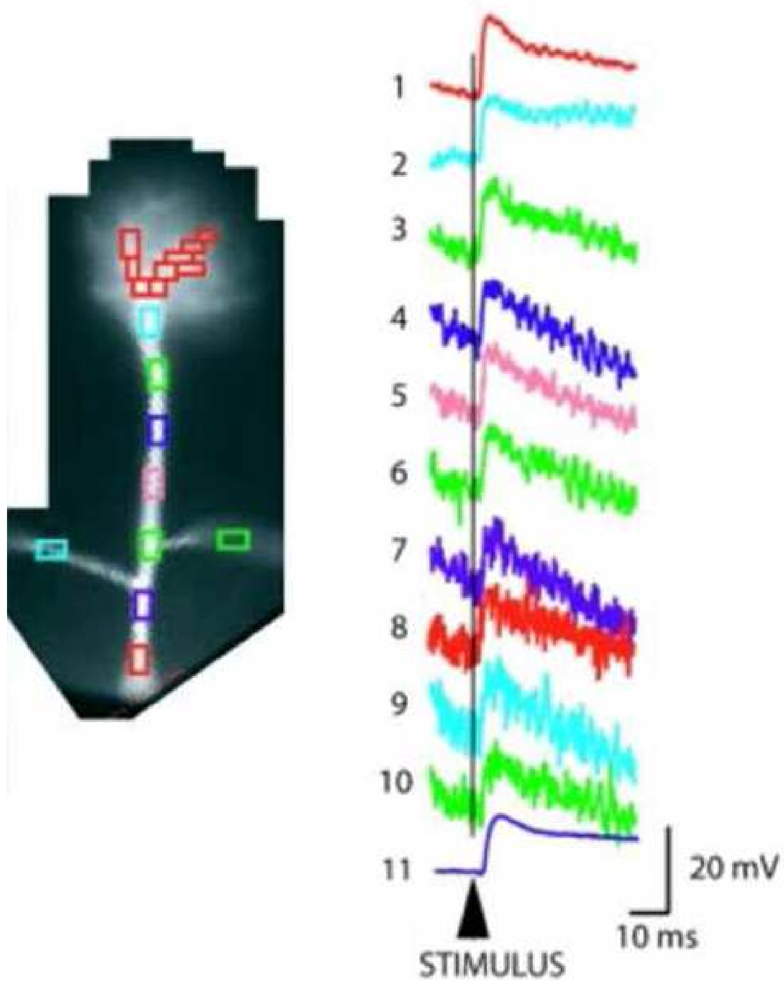

Basic paradigm: compartmental models

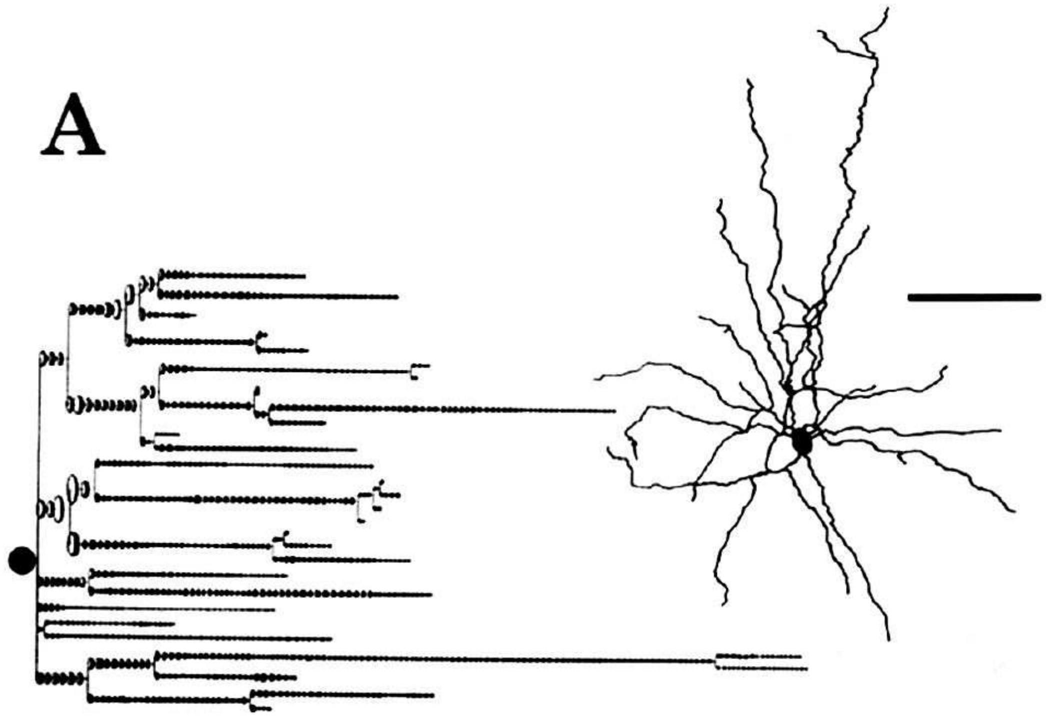

- write neuronal dynamics in terms of equivalent nonlinear, time-varying RC circuits
- leads to a coupled system of stochastic differential equations

Inference of spatiotemporal neuronal state given noisy observations

Variable of interest, q_t, evolves according to a noisy differential equation (e.g., cable equation):

$$dq/dt = f(q) + \epsilon_t.$$

Make noisy observations:

$$y(t) = g(q_t) + \eta_t.$$

We want to infer $E(q_t|Y)$: optimal estimate given observations. We also want errorbars: quantify how much we actually know about q_t.

If $f(.)$ and $g(.)$ are linear, and ϵ_t and η_t are Gaussian, then solution is classical: Kalman filter.

Extensions to nonlinear dynamics, non-Gaussian observations: hidden Markov ("state-space") model, particle filtering (Huys and Paninski, 2009)

Basic idea: Kalman filter

Dynamics and observation equations:

$$d\vec{V}/dt = A\vec{V} + \vec{\epsilon}_t$$

$$\vec{y}_t = B_t\vec{V} + \vec{\eta}_t$$

$V_i(t)$ = voltage at compartment i

A = cable dynamics matrix: includes leak terms ($A_{ii} = -g_l$) and intercompartmental terms ($A_{ij} = 0$ unless compartments are adjacent)

B_t = observation matrix: point-spread function of microscope

Even this case is challenging, since $d = \dim(\vec{V})$ is very large

Standard Kalman filter: $O(d^3)$ computation per timestep (matrix inversion)

(Paninski, 2010): methods for Kalman filtering in just $O(d)$ time: take advantage of sparse tree structure.

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, 2010). Generalizable to many other state-space models (Pnevmatikakis and Paninski, 2011).

Example: inferring voltage from subsampled observations

(Loading low-rank-speckle.mp4)

Example: summed observations

(Loading low-rank-horiz.mp4)

Applications

- Optimal experimental design: which parts of the neuron should we image? Submodular optimization (Huggins and Paninski, 2011)
- 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

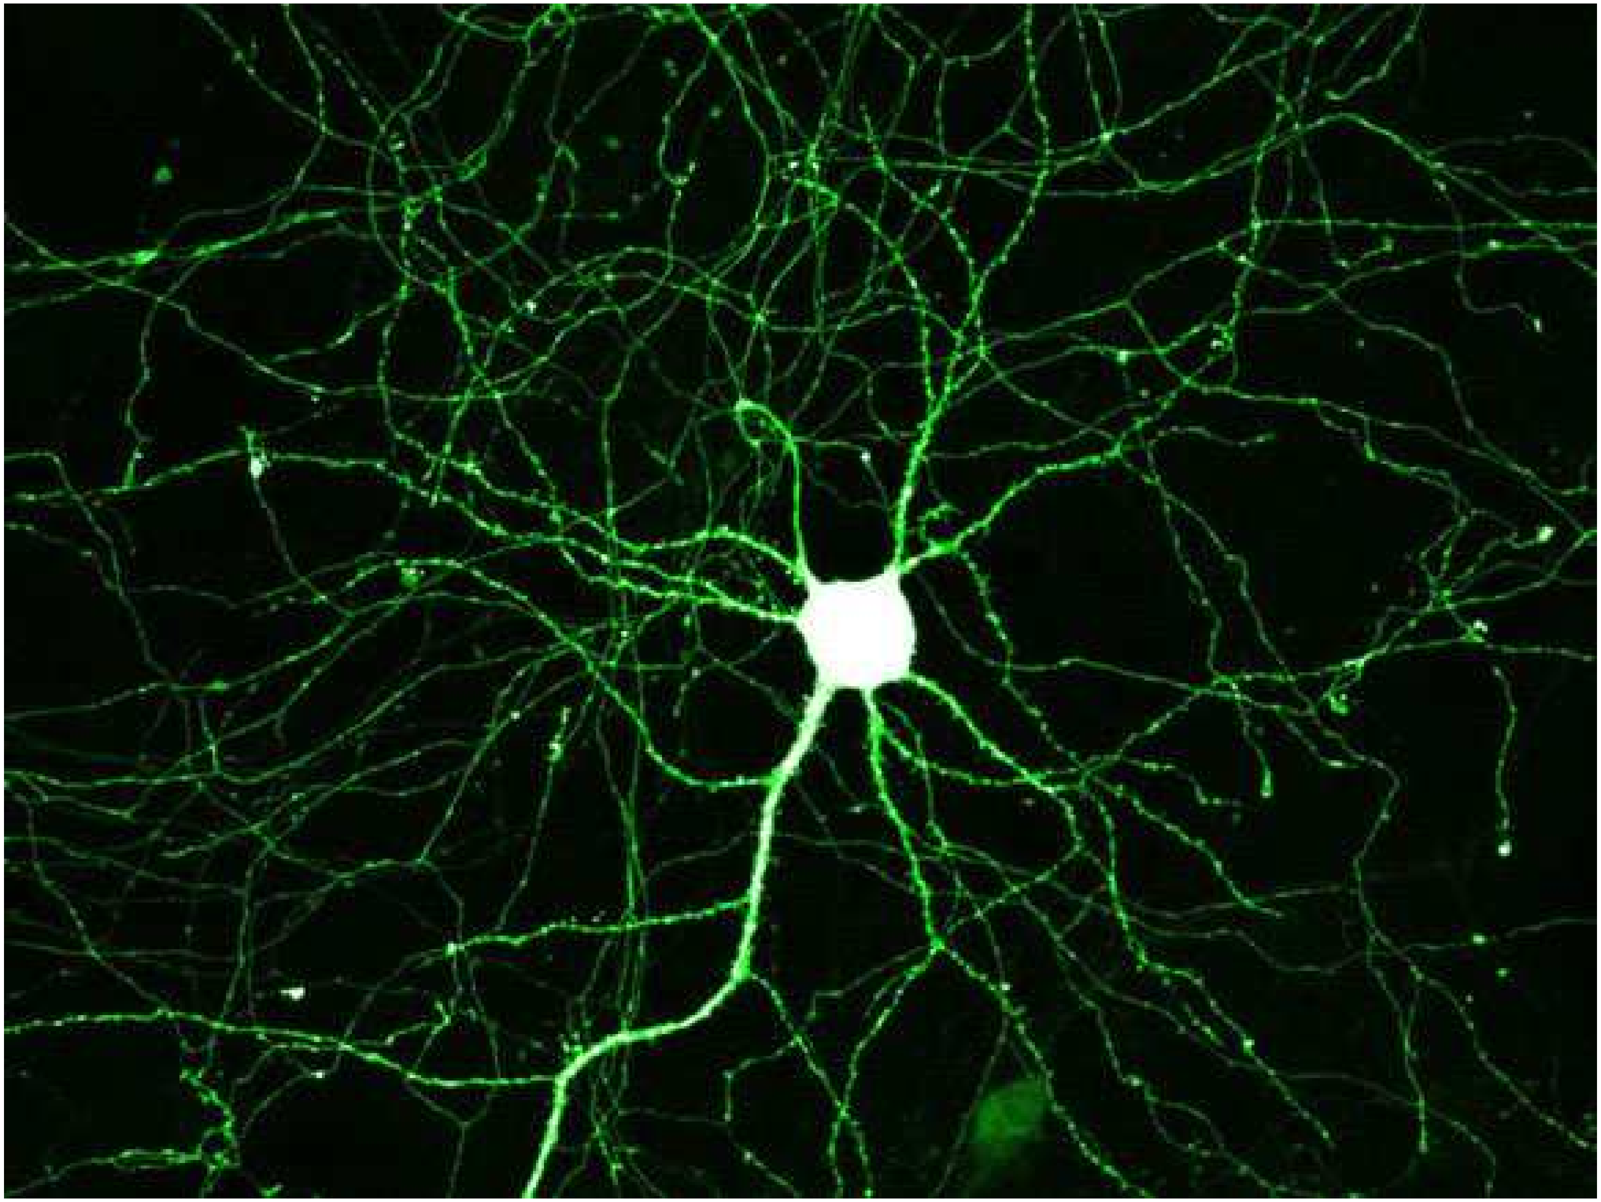

Application: synaptic locations/weights

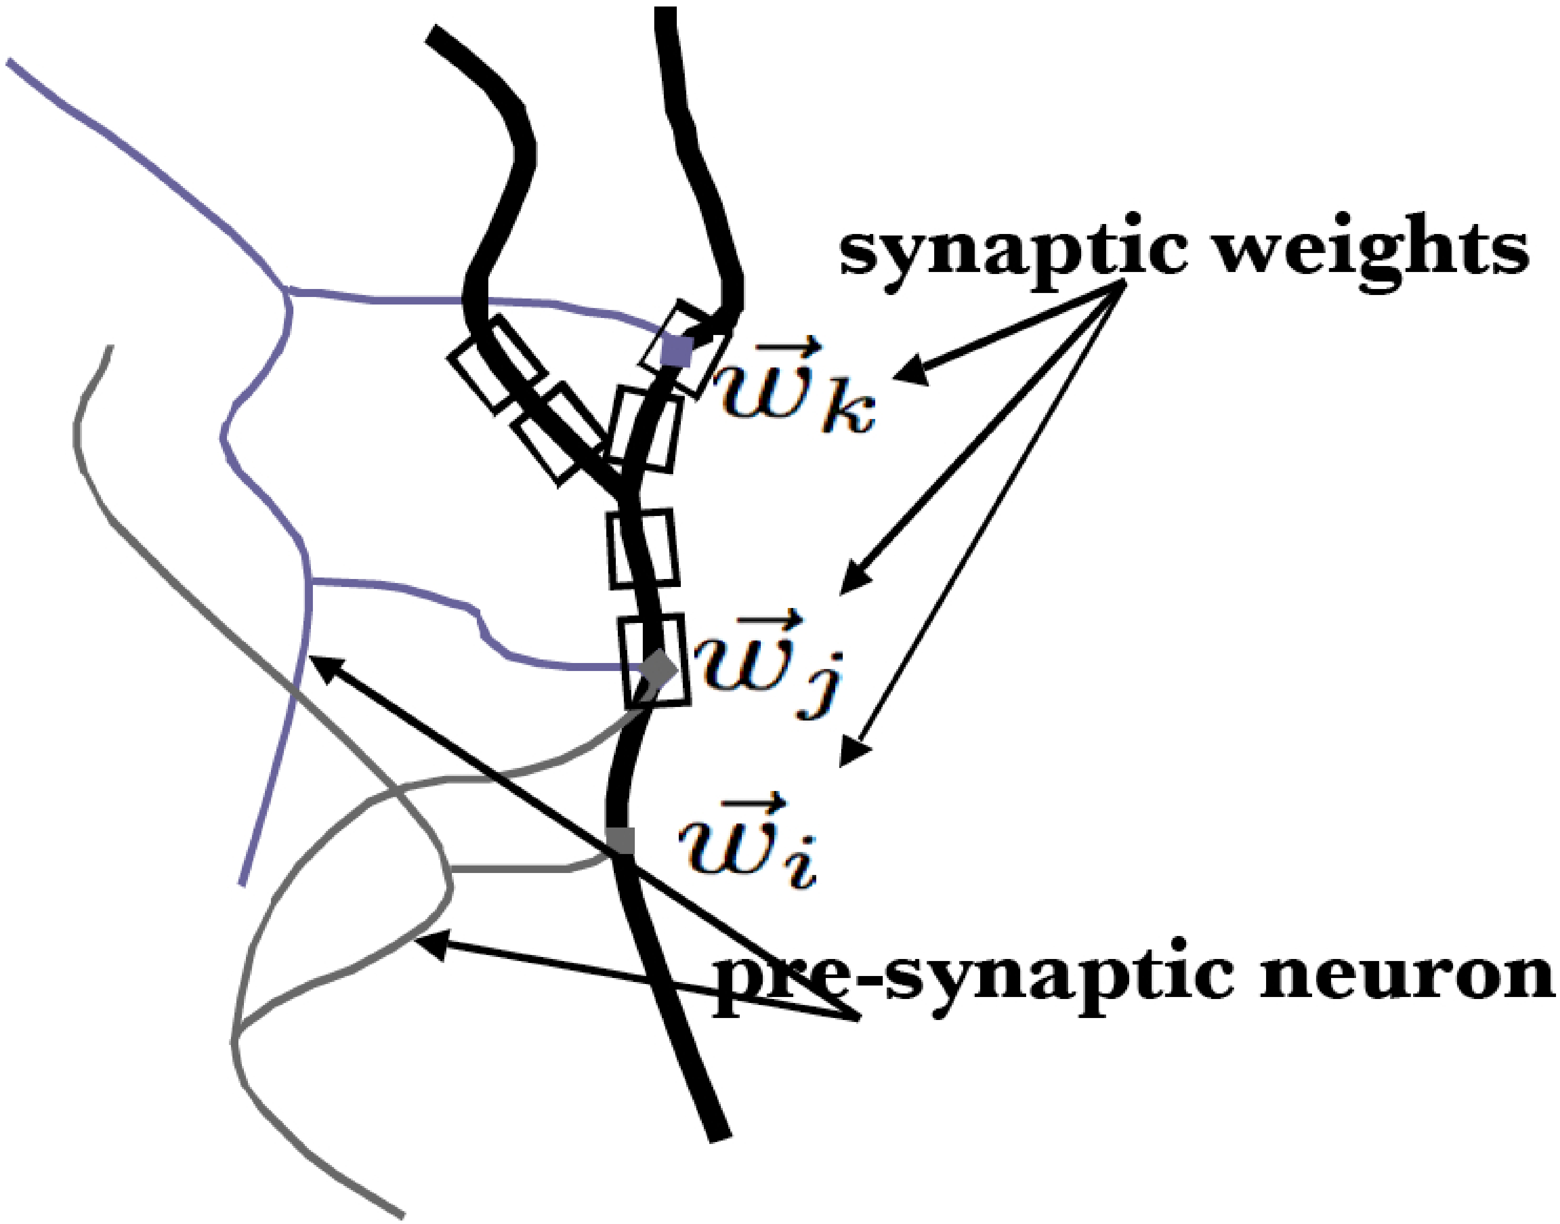

Application: synaptic locations/weights

Including known terms:

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

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

Loglikelihood is quadratic; W is a sparse vector. Adapt standard sparse regression methods from machine learning (Efron et al., 2004).

Total computation time: $O(dTk)$; $d = \#$ compartments, $T = \#$ timesteps, $k = \#$ nonzero weights.

Example: toy neuron

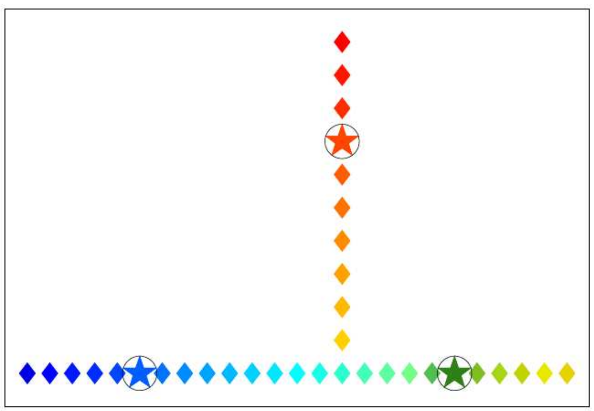

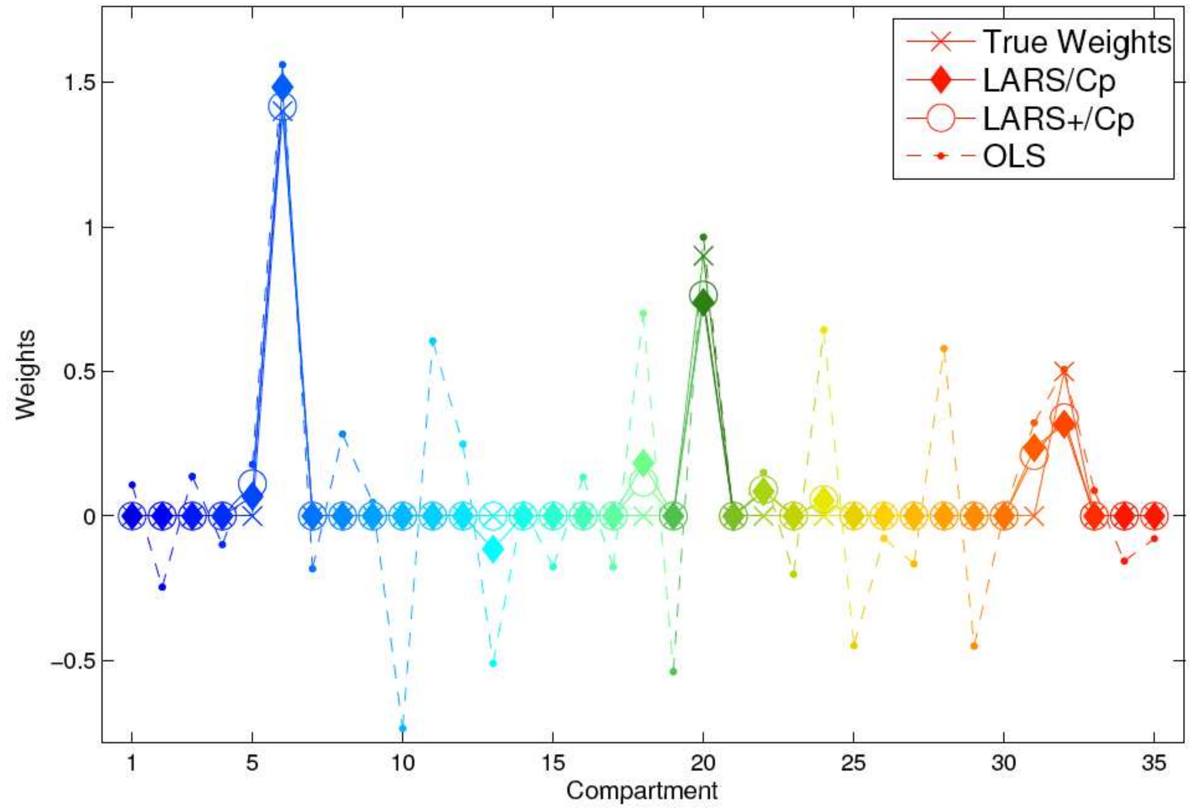

Example: toy neuron

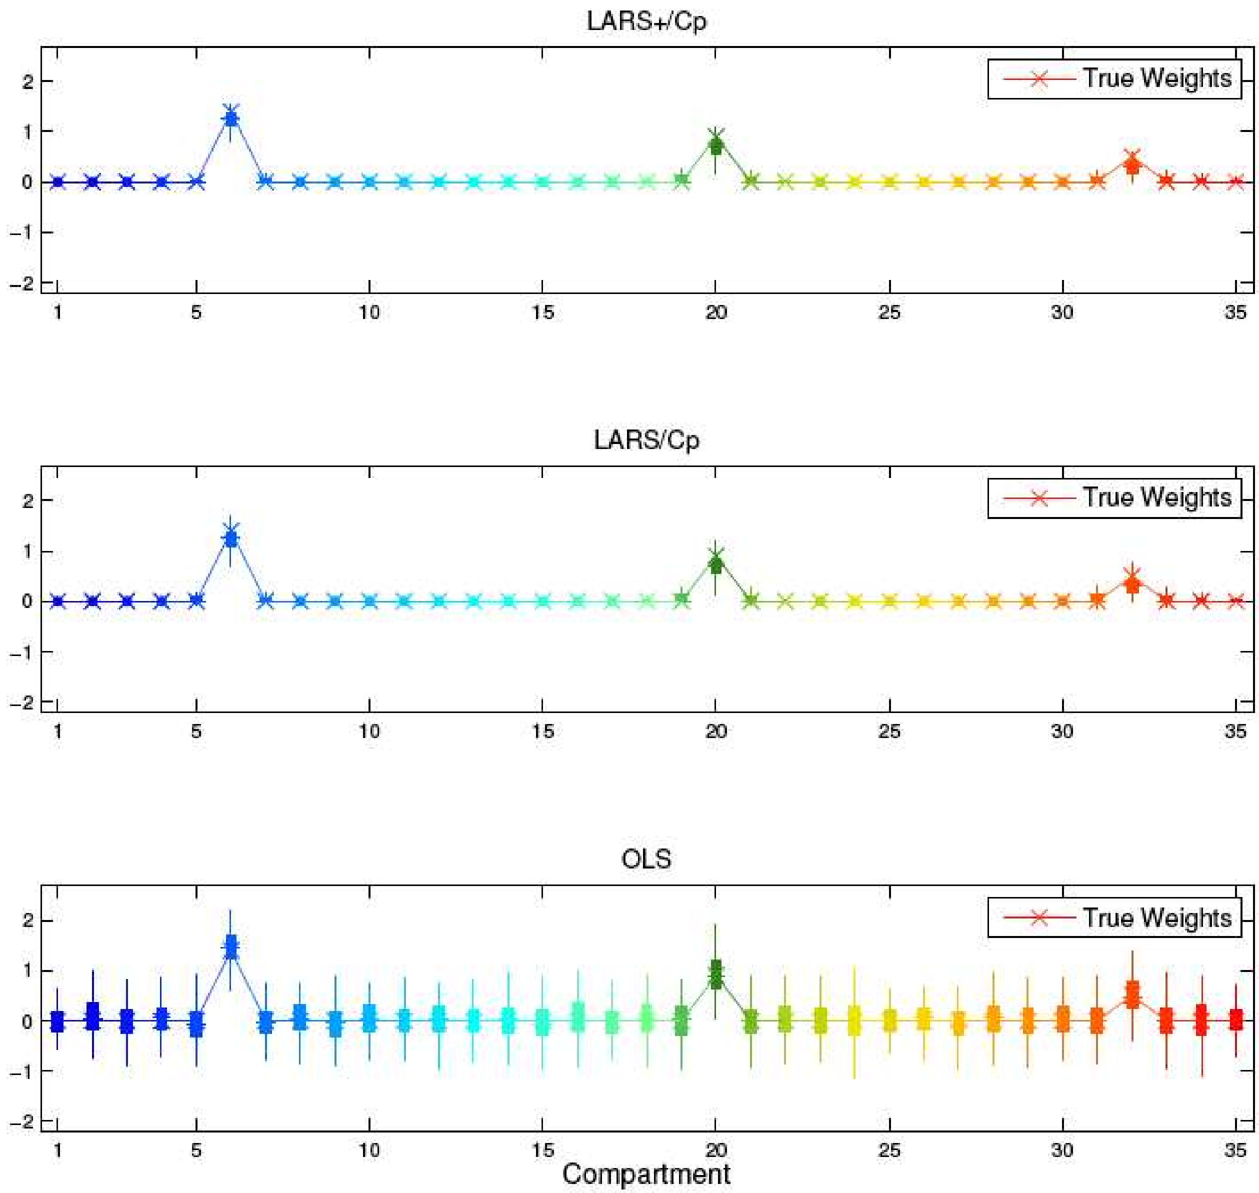

Example: real neural geometry

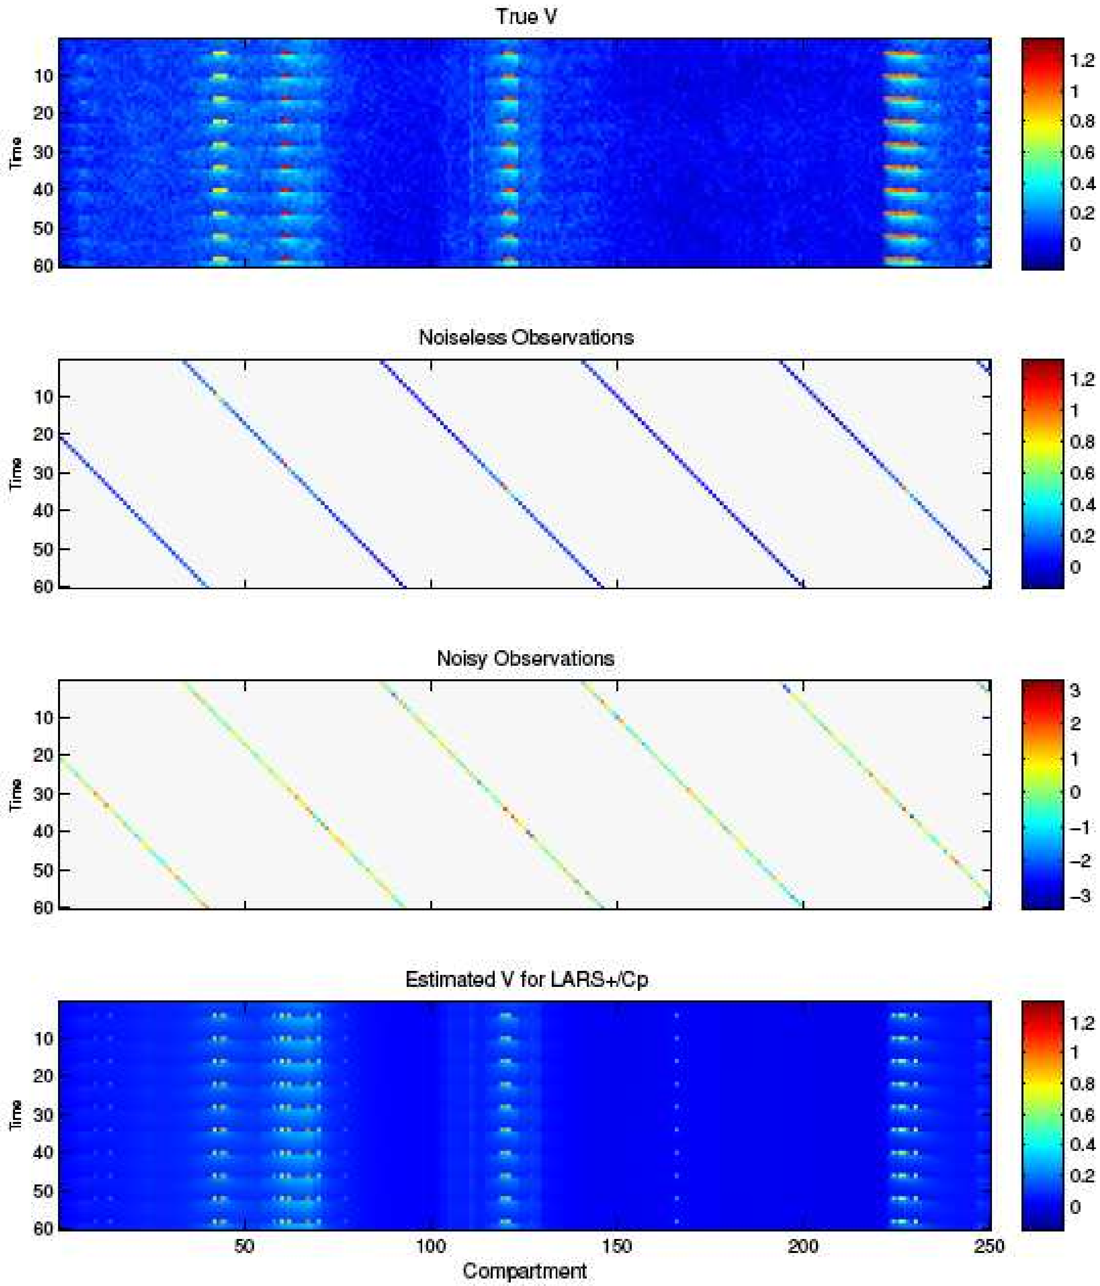

Example: real neural geometry

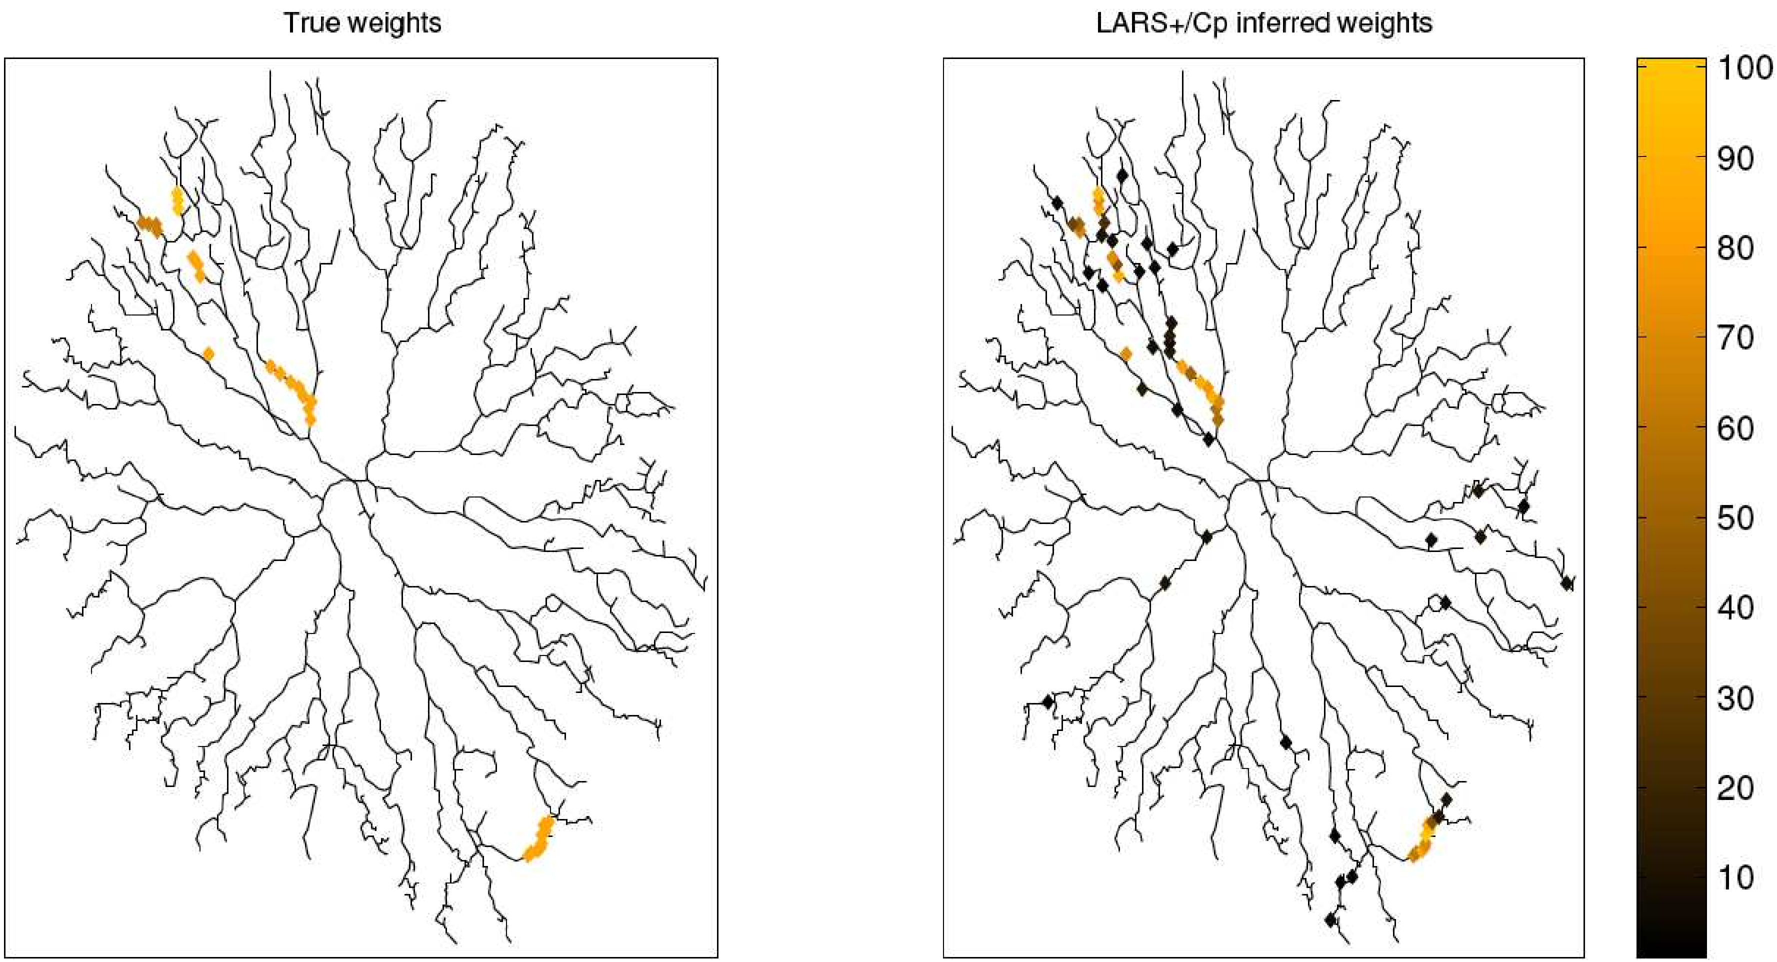

700 timesteps observed; 40 compartments (of > 2000) observed per timestep

Note: random access scanning essential here: results are poor if we observe the same compartments at each timestep.

Work in progress

- Quantifying robustness w.r.t. errors in tree reconstruction, parameter settings, etc.
- Combining fast Kalman filter with particle filter to model strongly nonlinear dendrites
- Exploiting local tree structure: distant compartments nearly uncoordinated ($\rightarrow$ factorized particle filter)
- Incorporating calcium measurements (Pnevmatikakis et al., 2011)

Conclusions

- Modern statistical approaches provide flexible, powerful methods for answering key questions in neuroscience
- Close relationships between biophysics and statistical modeling
- Modern optimization methods make computations very tractable; suitable for closed-loop experiments
- Experimental methods progressing rapidly; many new challenges and opportunities for breakthroughs based on statistical ideas

References

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.

Efron, B., Hastie, T., Johnstone, I., and Tibshirani, R. (2004). Least angle regression. *Annals of Statistics*, 32:407–499.

Huggins, J. and Paninski, L. (2011). Optimal experimental design for sampling voltage on dendritic trees. *J. Comput. Neuro.*, In press.

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.

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.

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

Pnevmatikakis, E., Kelleher, K., Chen, R., Josic, K., Saggau, P., and Paninski, L. (2011). Fast nonnegative spatiotemporal calcium smoothing in dendritic trees. *COSYNE*.

Pnevmatikakis, E. and Paninski, L. (2011). Fast interior-point inference in high-dimensional sparse, penalized state-space models. *Submitted*.