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

October 5, 2012

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

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)
- optical methods (eg two-photon, FLIM) and optogenetics (channelrhodopsin, viral tracers, "brainbow")
- high-density multielectrode recordings (Litke's 512-electrode retinal readout system; Shepard's 65,536-electrode active array)

Example: neural prosthetics

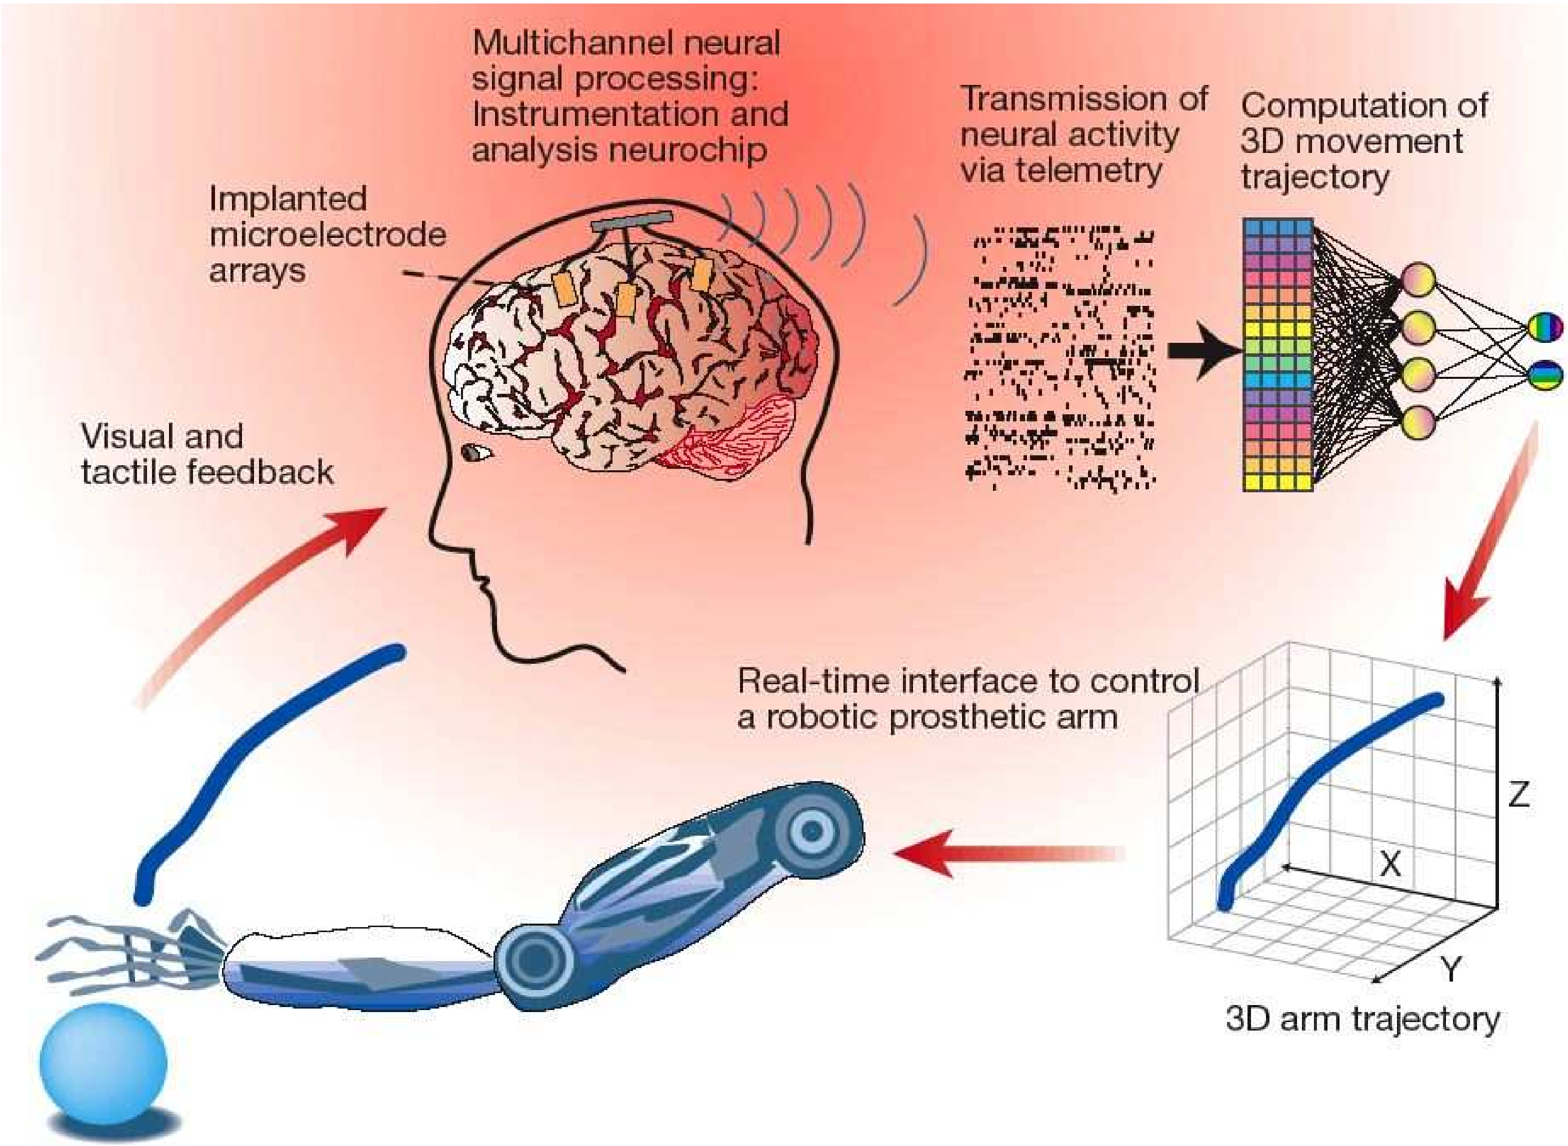

Example: neural prosthetics

(Loading monkey-zombies.mp4)

Example: retinal ganglion neuronal data

Preparation: dissociated macaque retina

— extracellularly-recorded responses of populations of RGCs

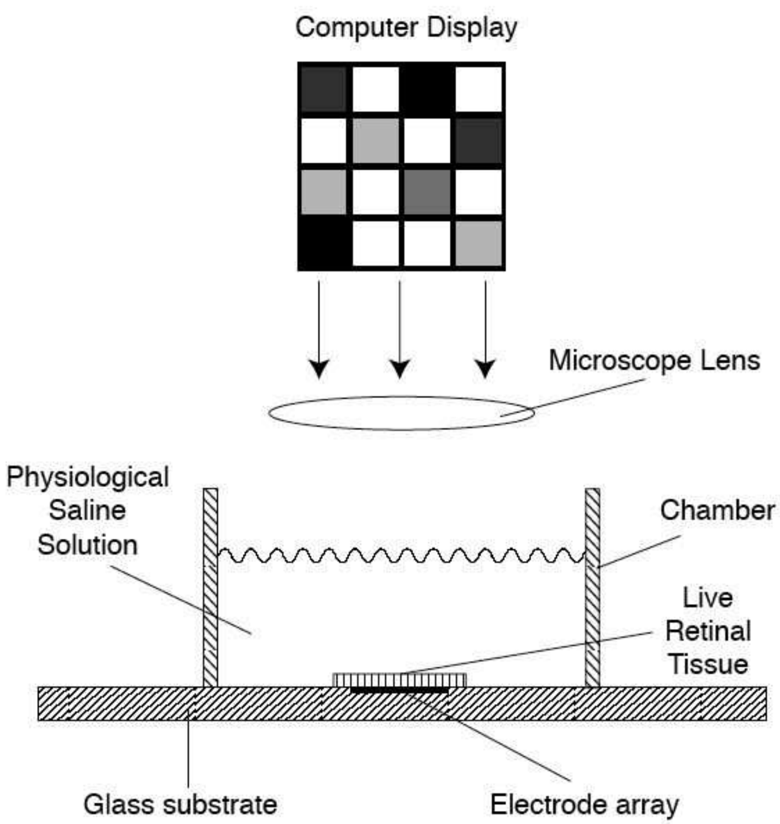

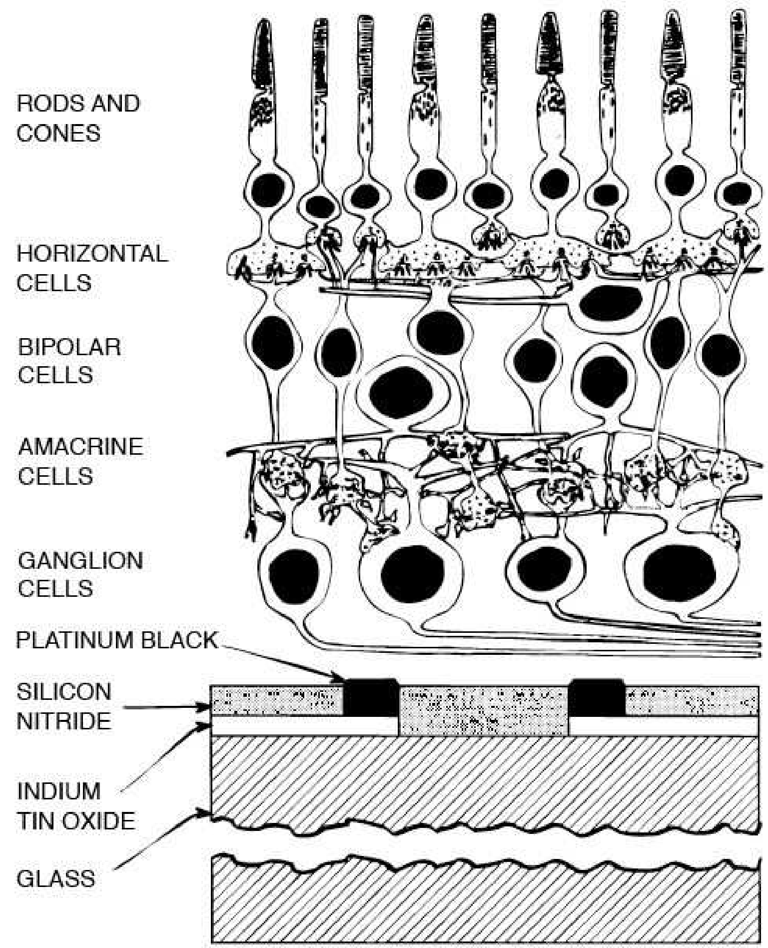

Receptive fields tile visual space

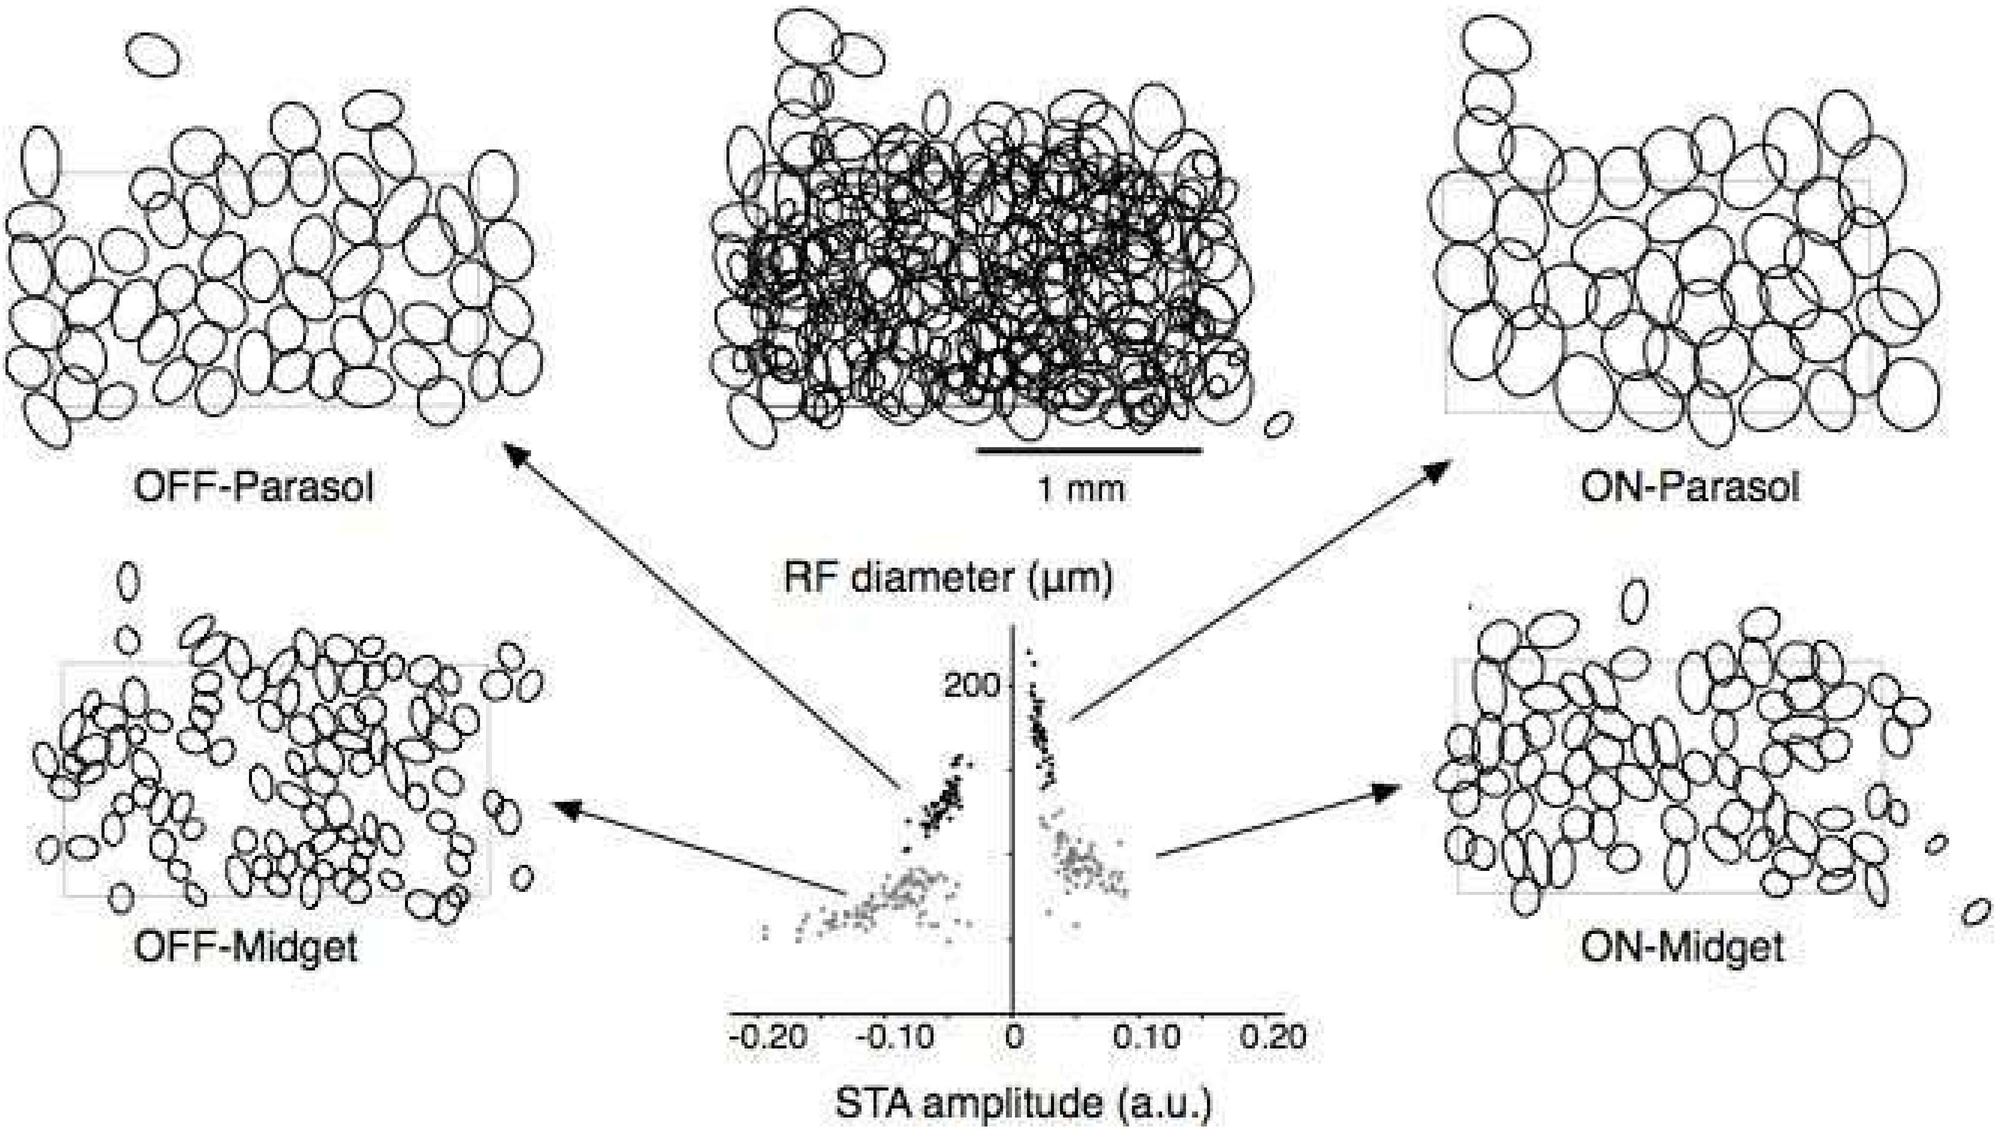

Multineuronal point-process model

$$\lambda_i(t) = \exp\left(k_i \cdot x(t) + h_i \cdot y_i(t) + \sum_{i \neq j} l_{i,j} \cdot y_j(t) + Lq(t)\right)$$

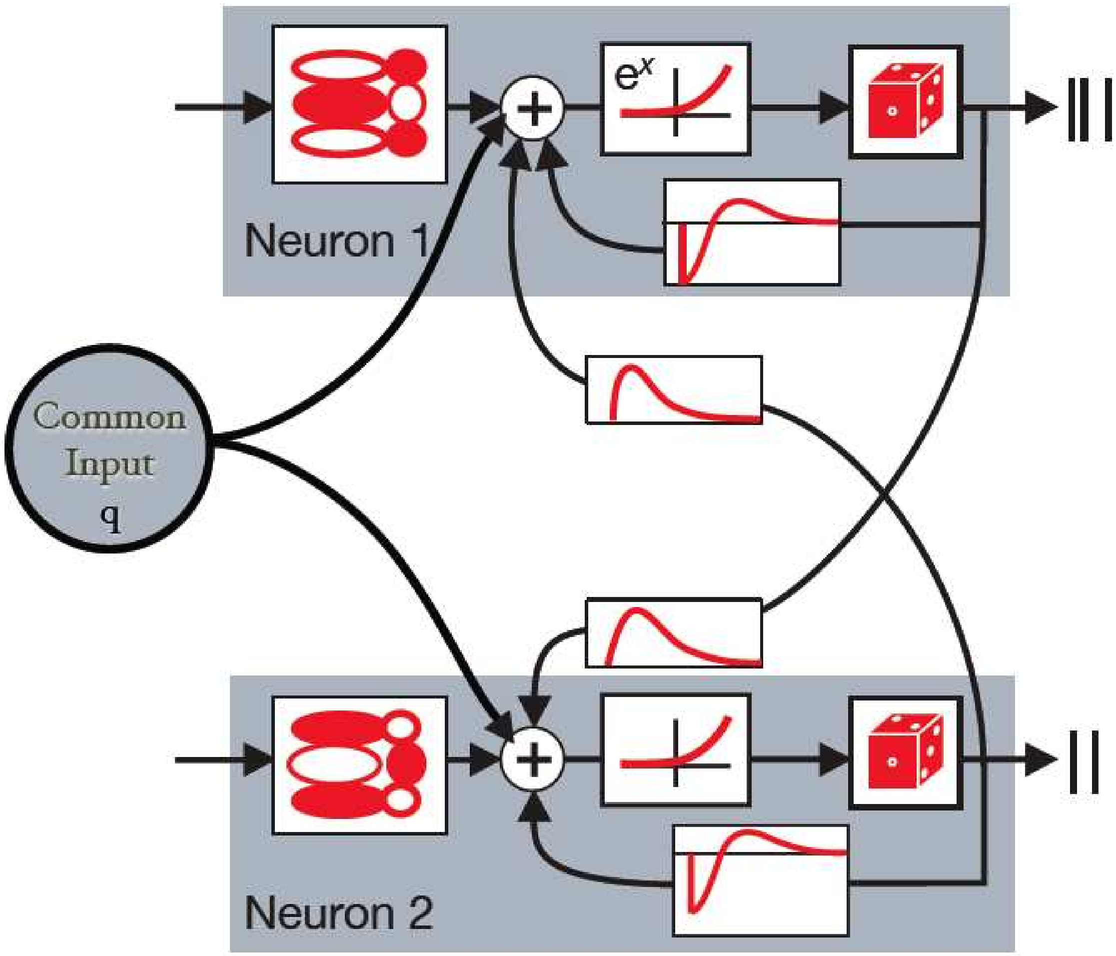

— likelihood is tractable to compute and to maximize (concave optimization)

(Paninski, 2004; Paninski et al., 2007; Pillow et al., 2008; Paninski et al., 2010)

Predicting single-neuron responses

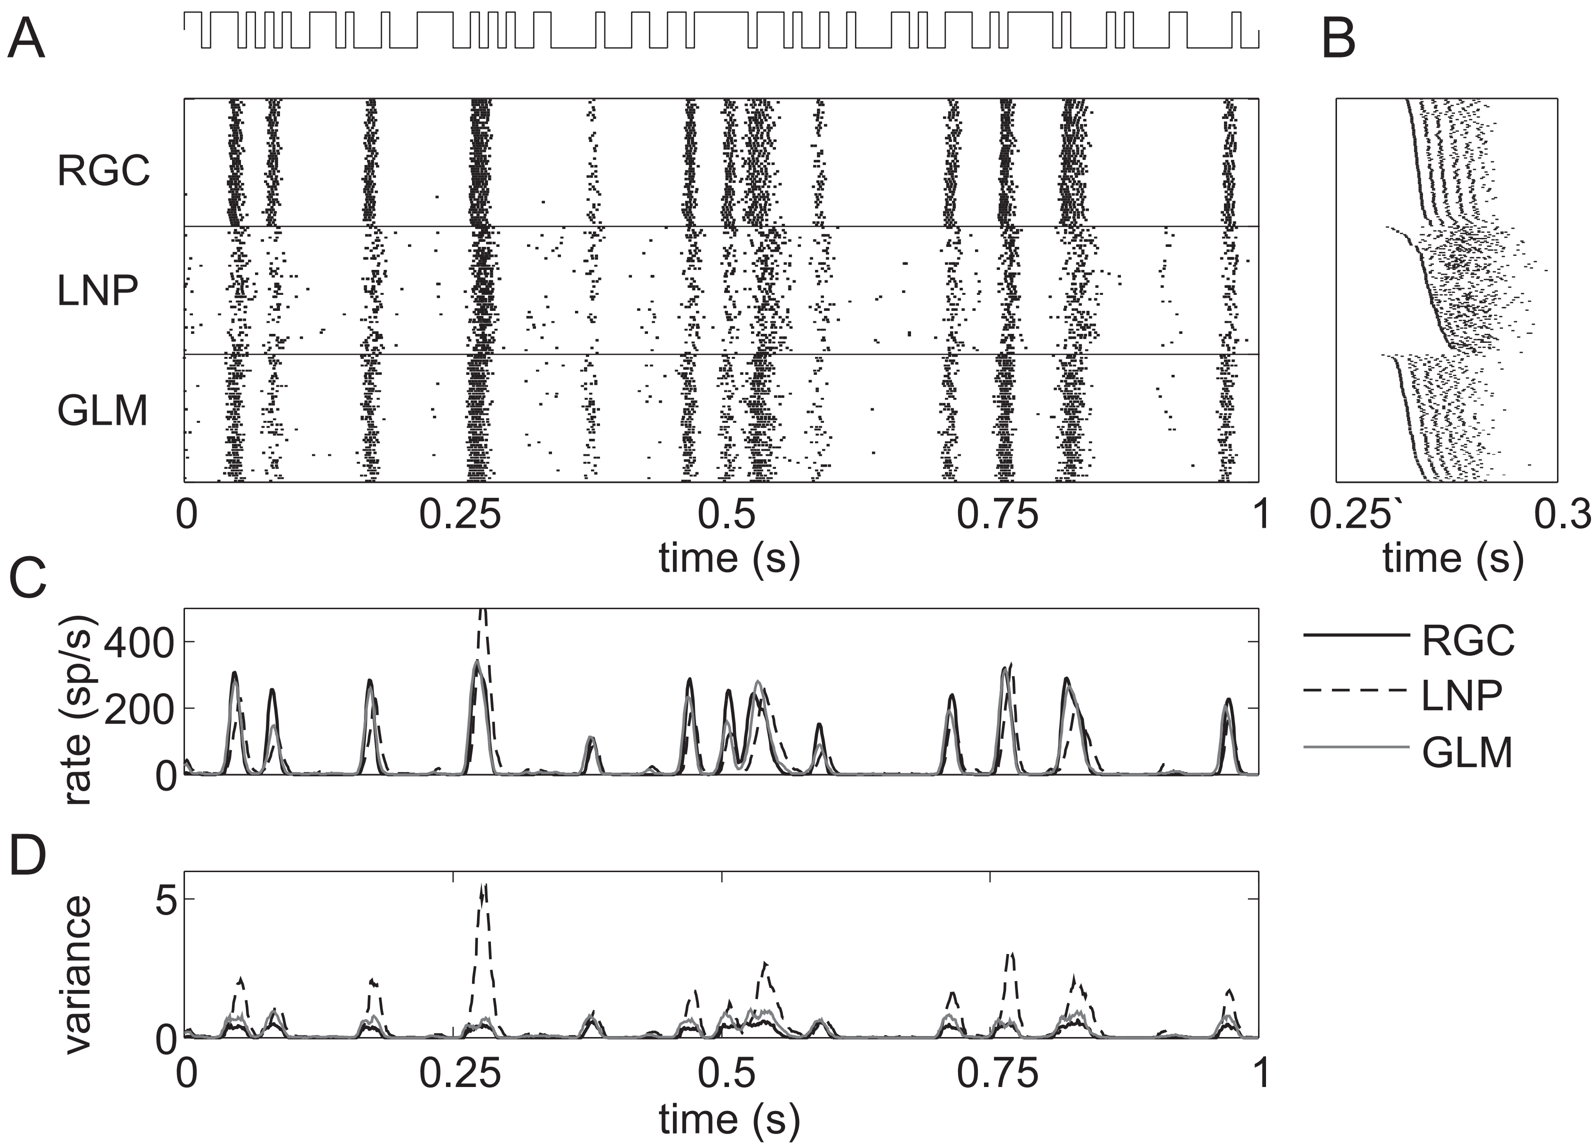

— model captures high precision of retinal responses. Also captures correlations between neurons.

Optimal Bayesian decoding

$E(\vec{x}|spikes) \approx \arg\max_{\vec{x}} \log P(\vec{x}|spikes) = \arg\max_{\vec{x}} [\log P(spikes|\vec{x}) + \log P(\vec{x})]$

(Loading yashar-decode.mp4)

— Computational points:

- $\log P(spikes|\vec{x})$ is concave in $\vec{x}$: concave optimization again.
- Decoding can be done in linear time via standard Newton-Raphson methods, since Hessian of $\log P(\vec{x}|spikes)$ w.r.t. $\vec{x}$ is banded (Pillow et al., 2010).

Optimal Bayesian decoding

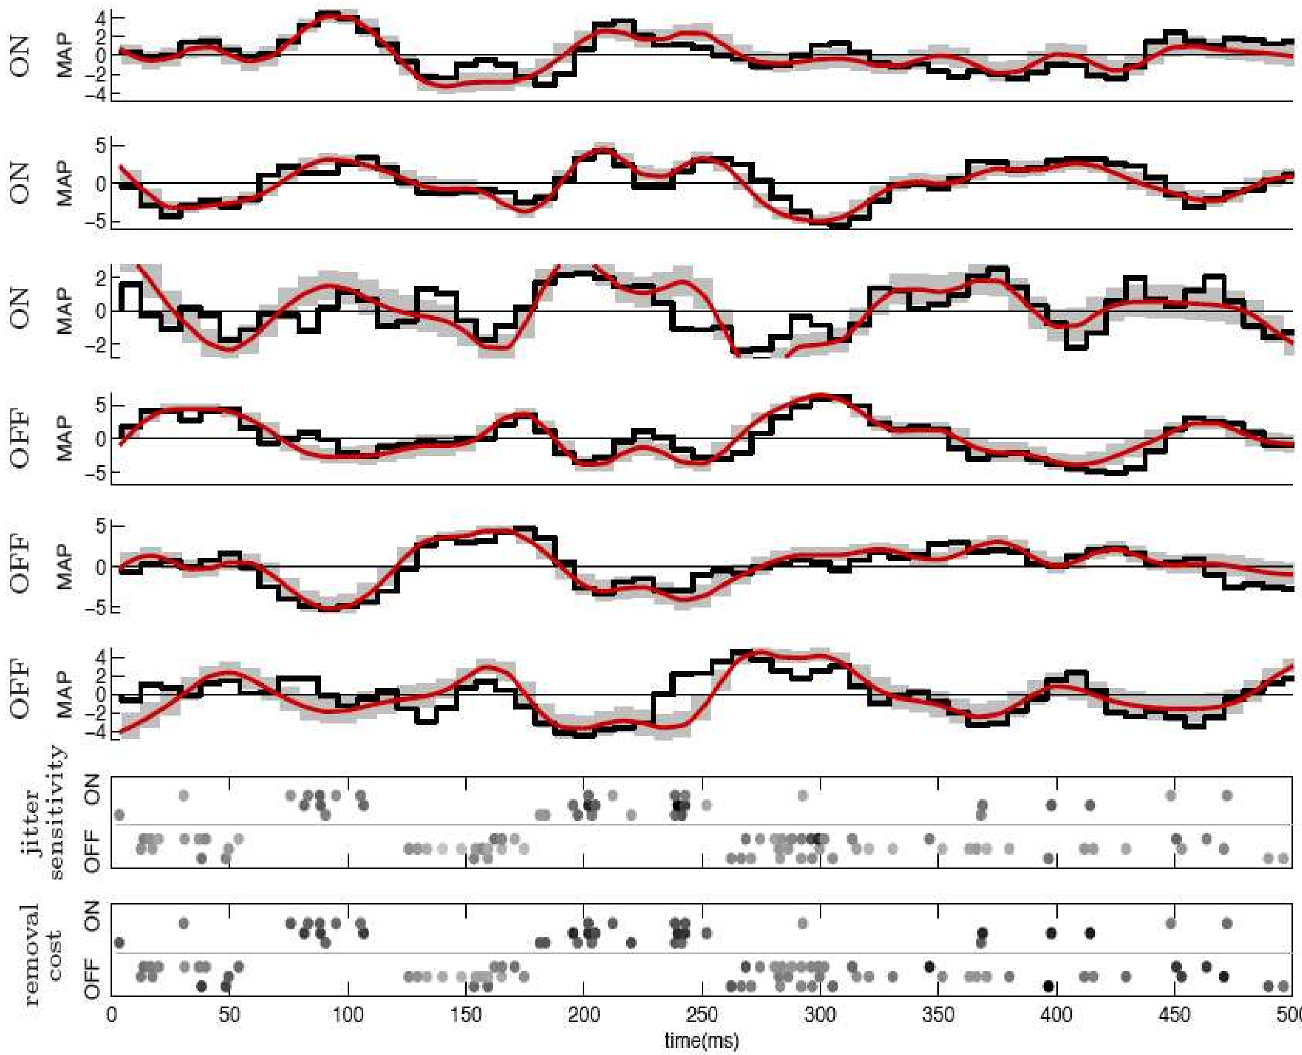

— further applications: decoding velocity signals (Lalor et al., 2009), tracking images perturbed by eye jitter (Pfau et al., 2009)

— paying attention to correlations improves decoding accuracy (Pillow et al., 2008).

Inferring cone maps

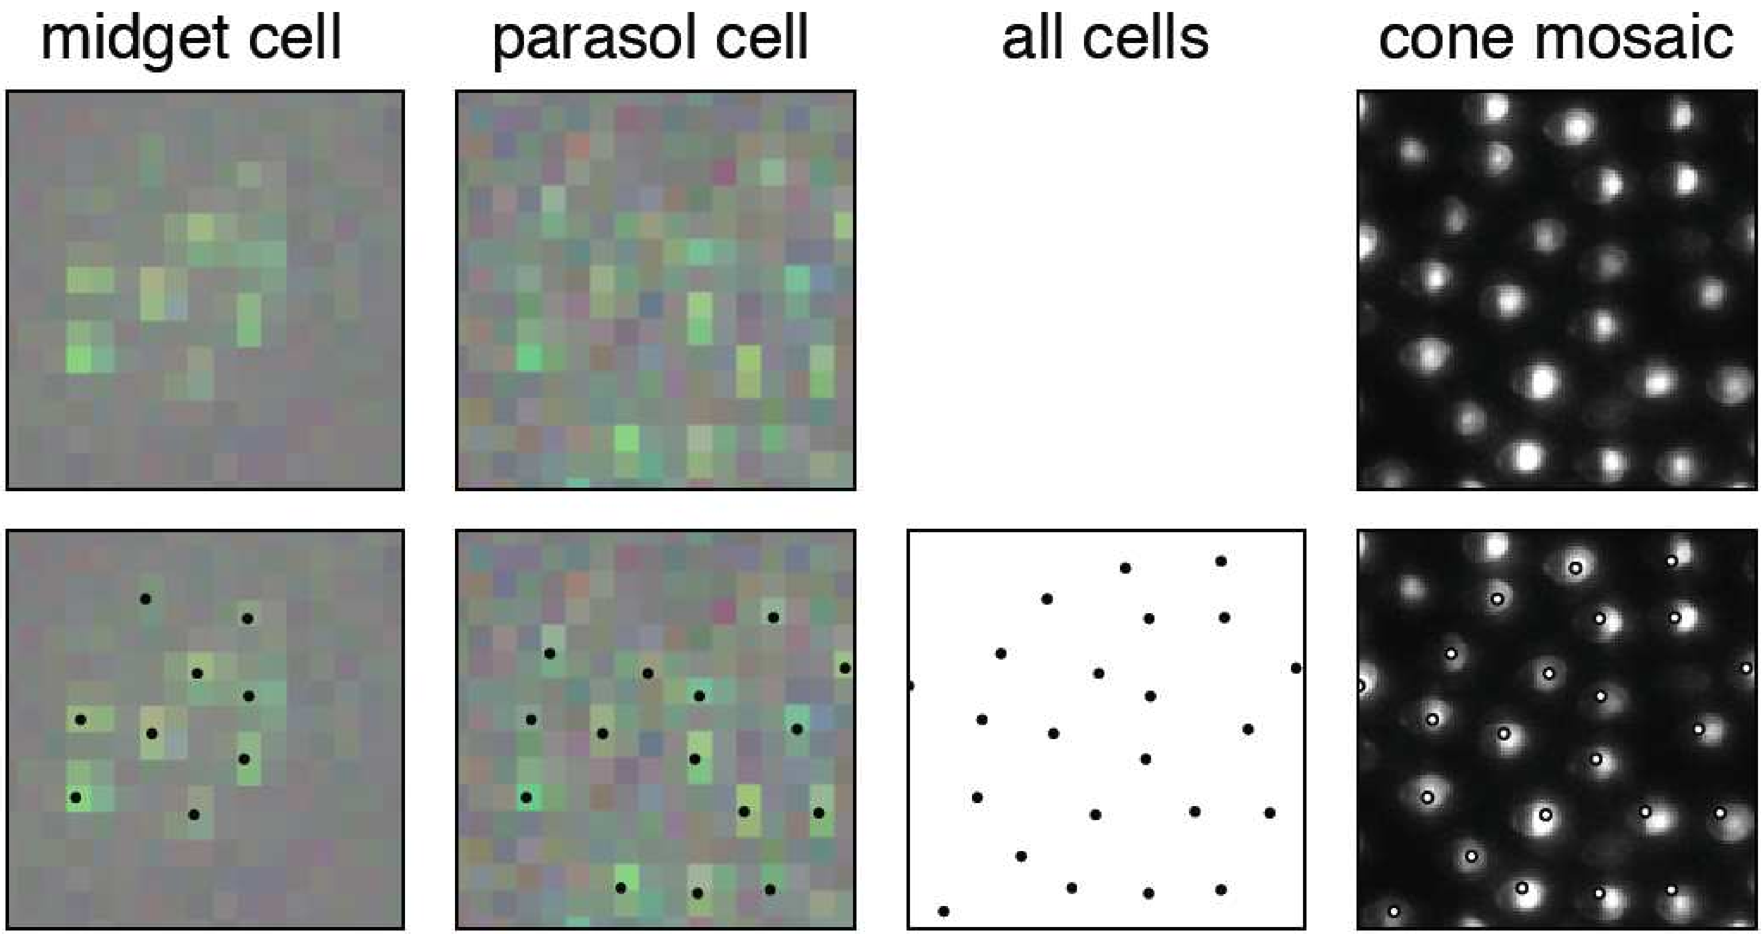

Inferring cone maps

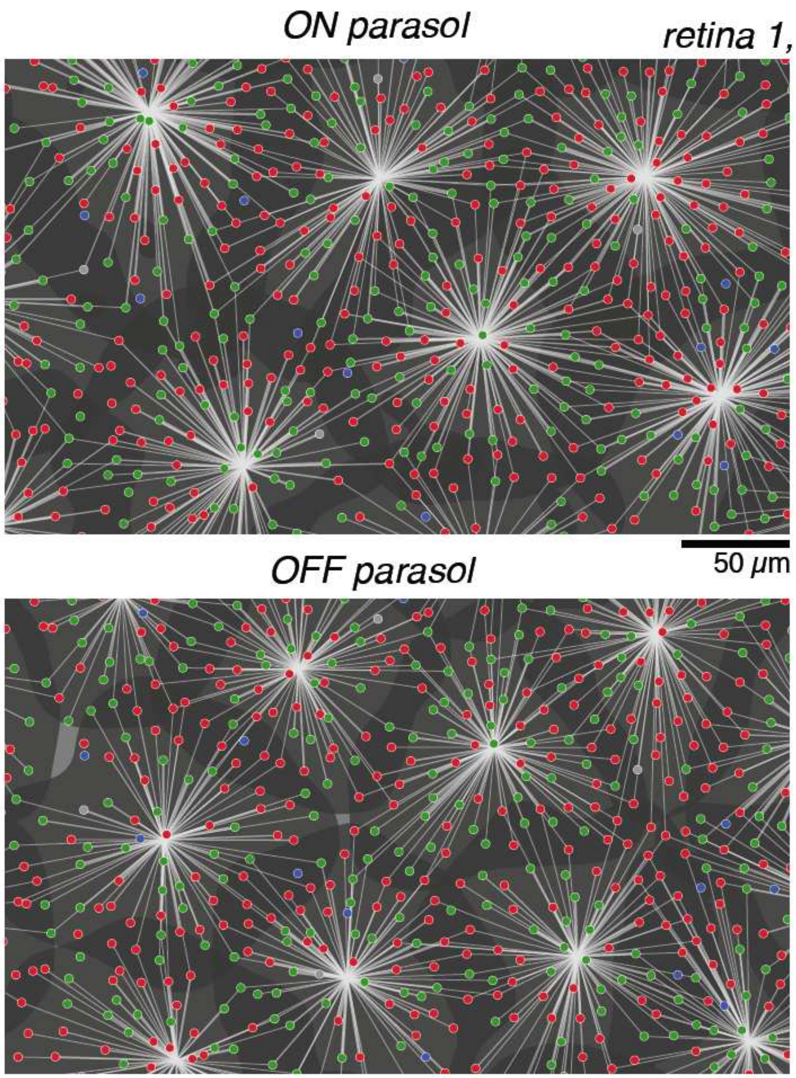

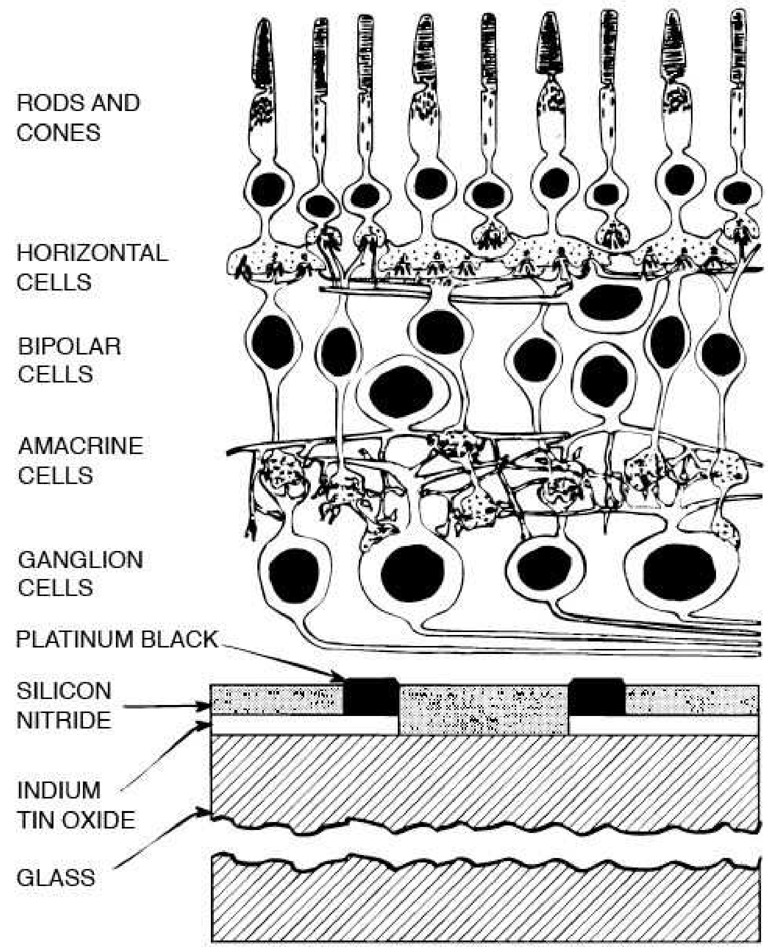

— cone locations and color identity inferred accurately with high-resolution stimuli; Bayesian approach integrates information over multiple simultaneously recorded neurons (Field et al., 2010).

Another major challenge: circuit inference

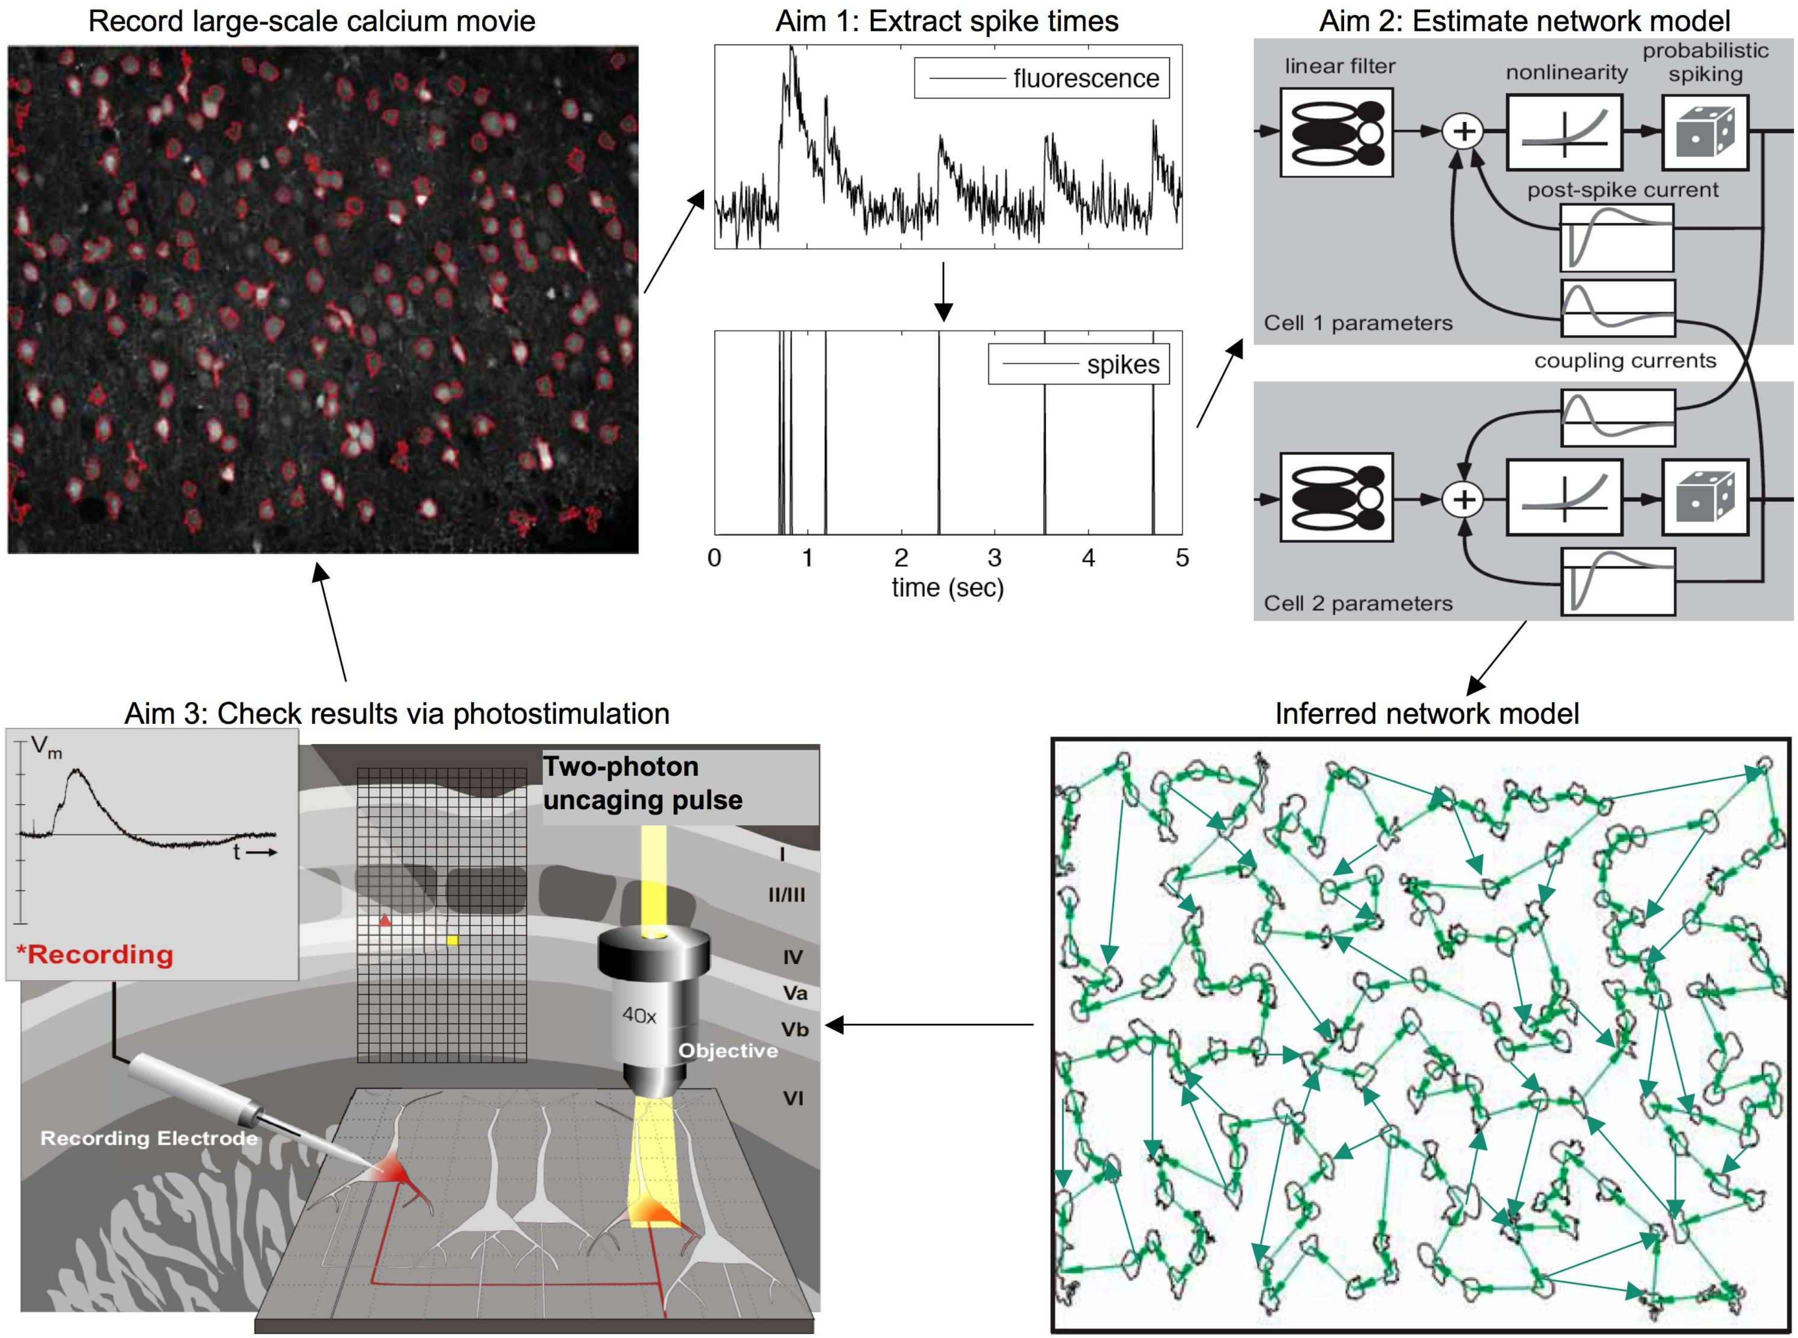

Challenge: slow, noisy calcium data

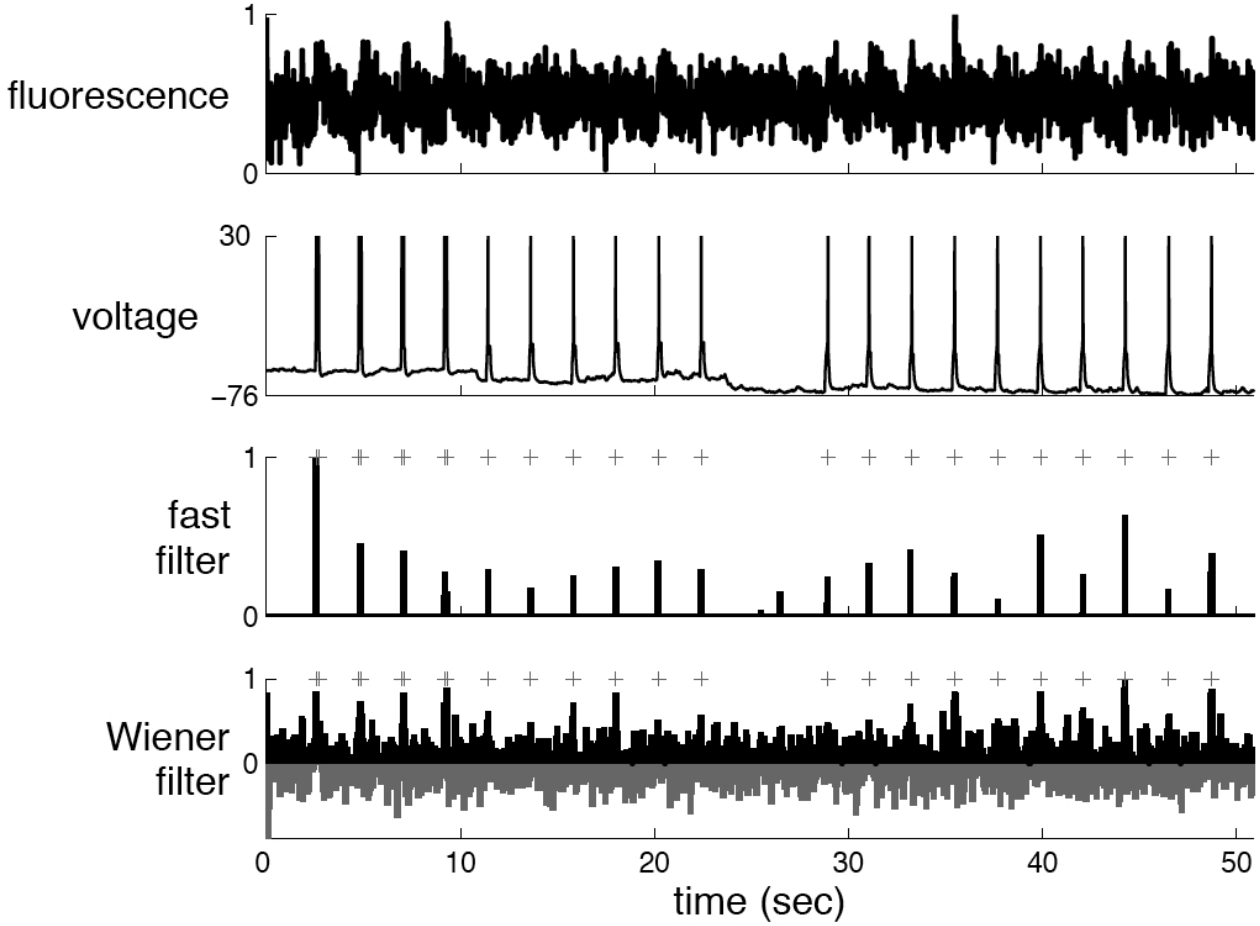

First-order model:

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

— $\tau \approx 100$ ms; nonnegative deconvolution problem. Can be solved by new fast methods (Vogelstein et al., 2009; Vogelstein et al., 2010; Mishchenko et al., 2010).

Spatiotemporal Bayesian spike estimation

(Loading Tim-data.mp4)

Simulated circuit inference

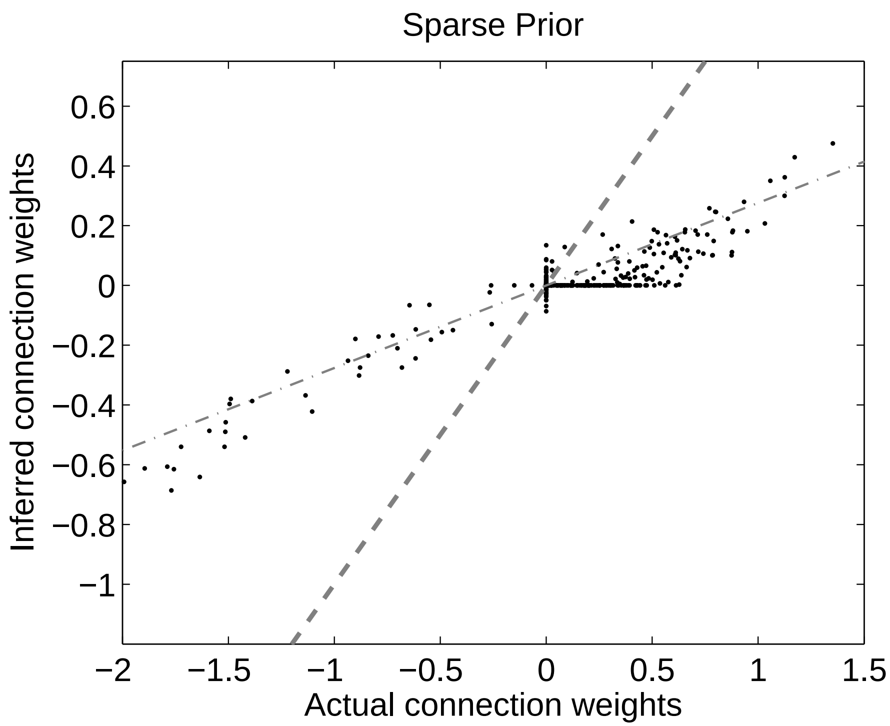

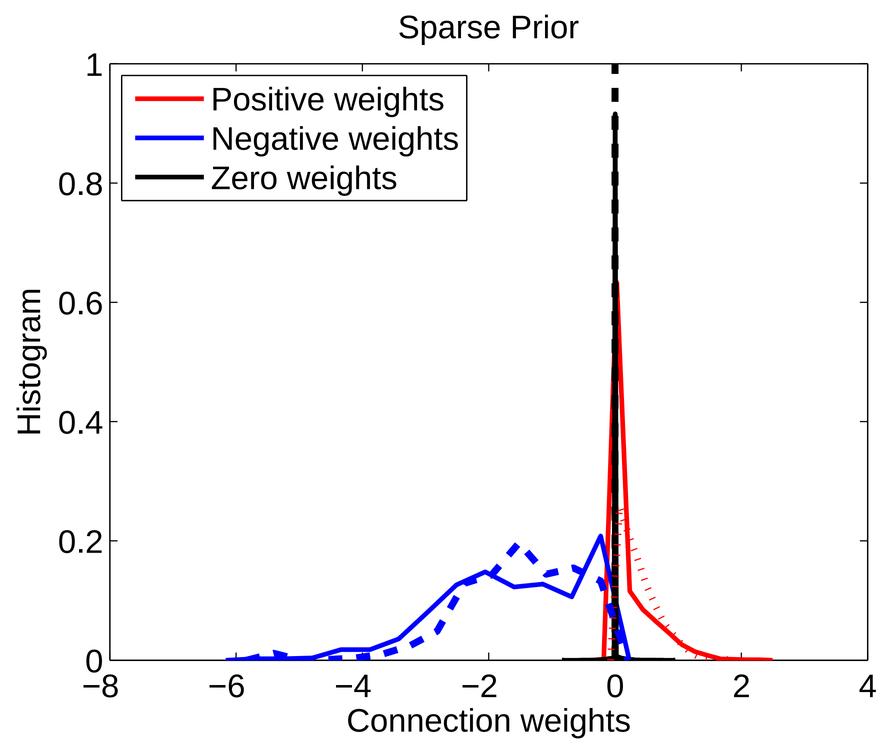

— Connections are inferred with the correct sign in conductance-based integrate-and-fire networks with biologically plausible connectivity matrices (Mishchencko et al., 2009).

Good news: connections are inferred with the correct sign. Fast enough to estimate connectivity in real time (T. Machado). Next step: close the loop.

A final challenge: understanding dendrites

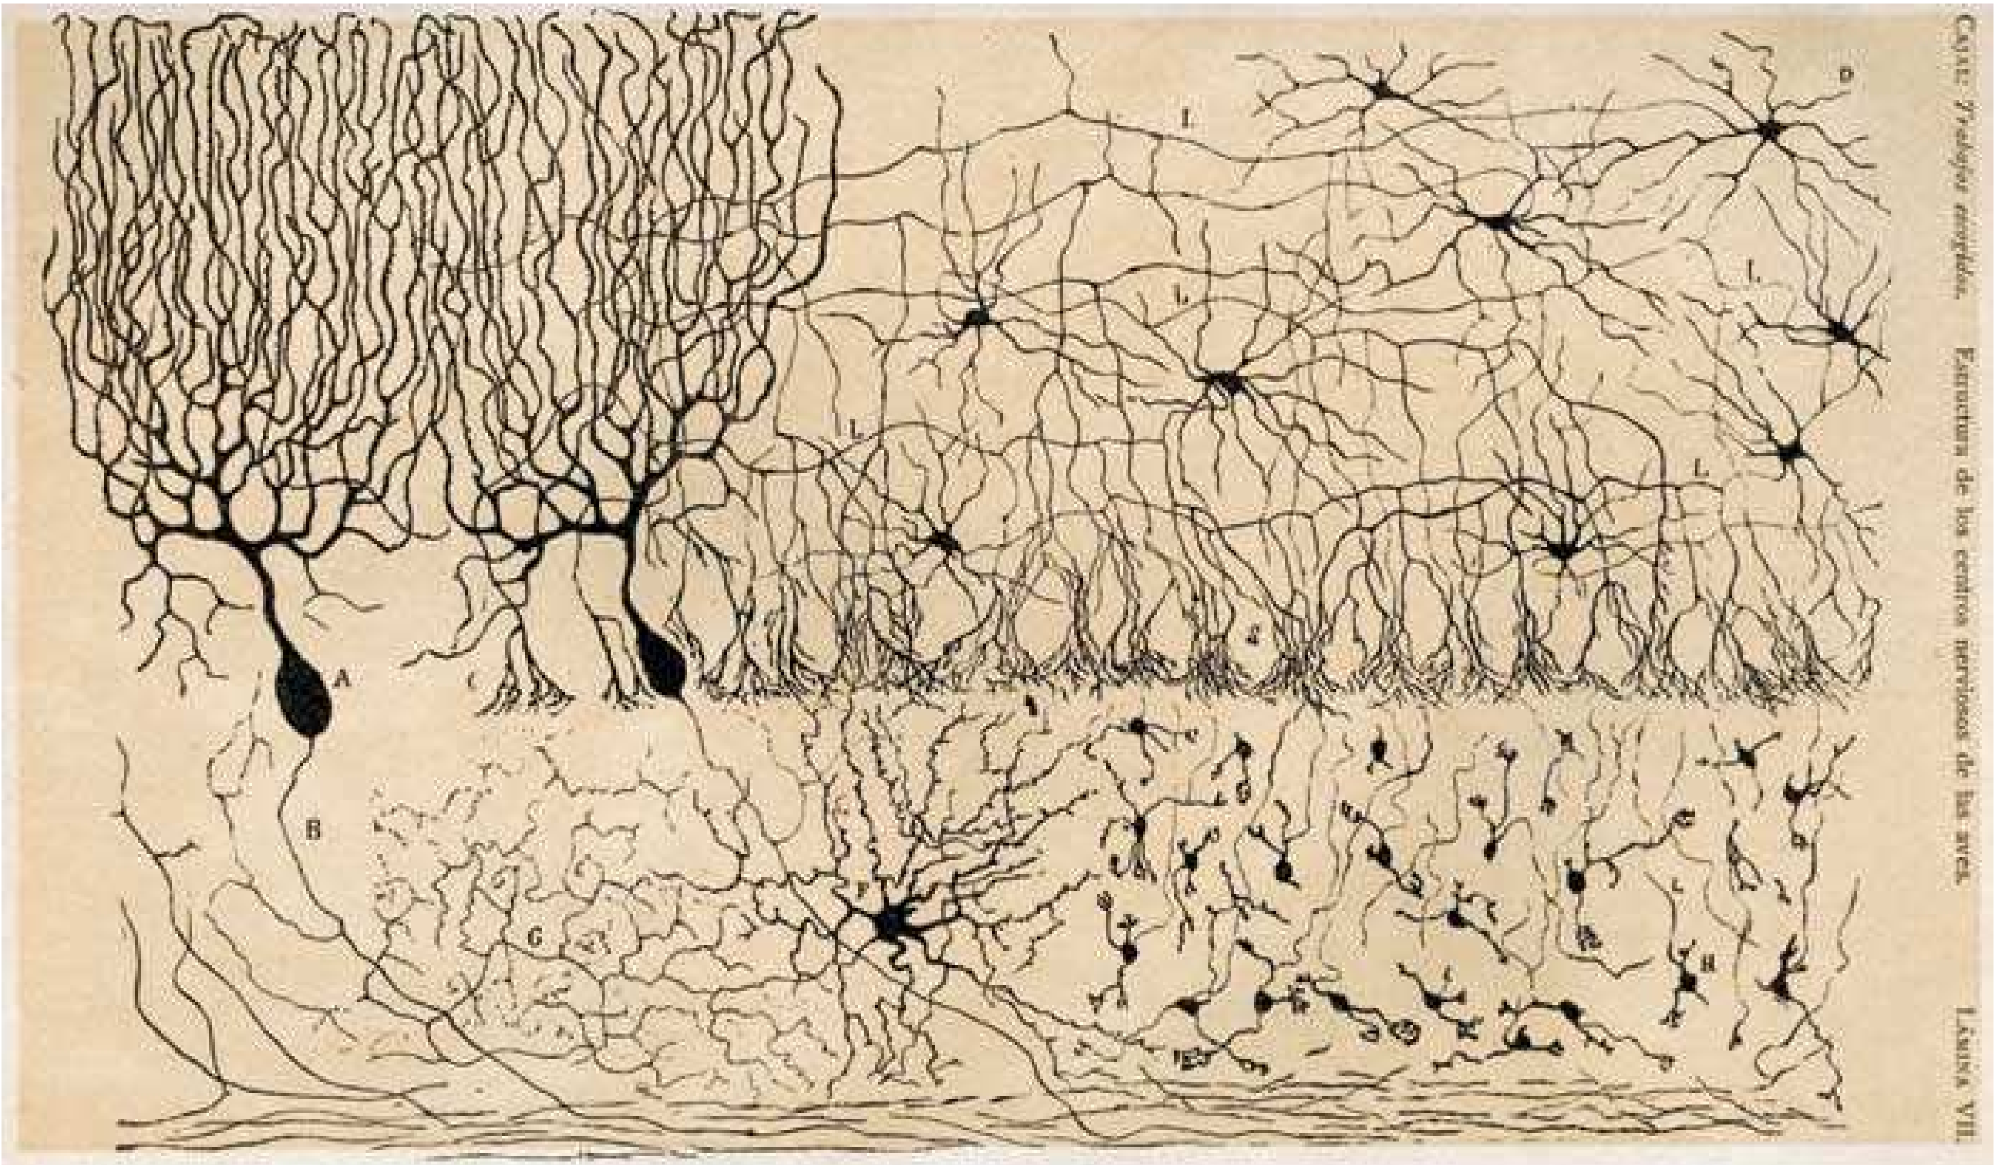

Ramon y Cajal, 1888.

A spatiotemporal 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.

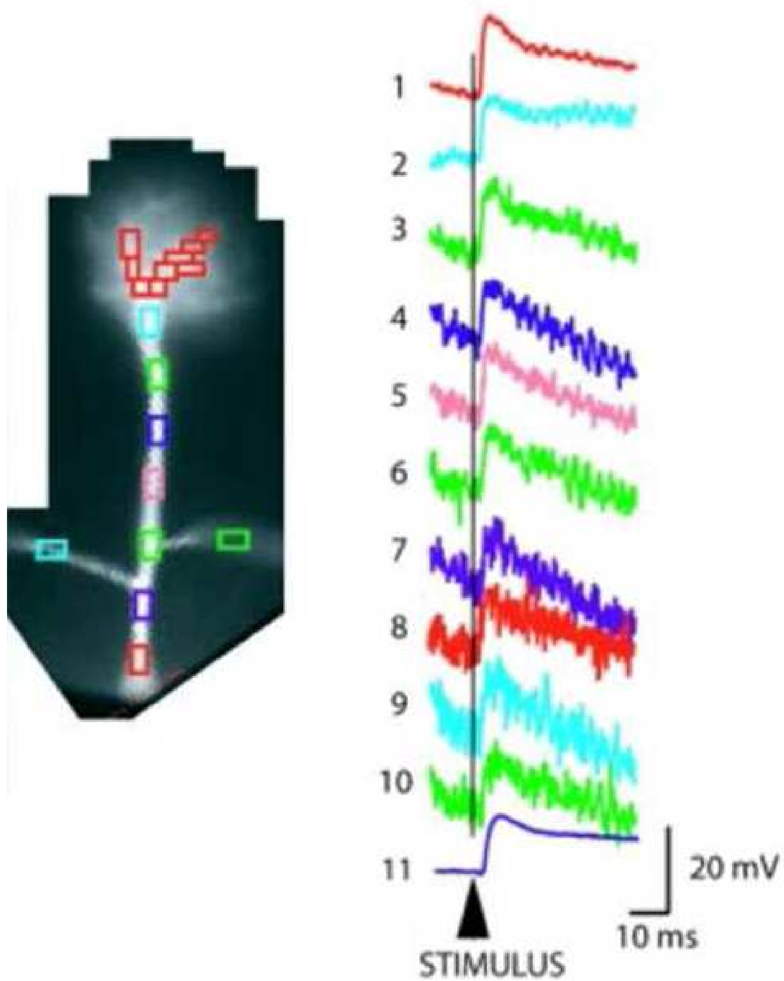

Inference of spatiotemporal neuronal state given noisy observations

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

$$dV/dt = f(V) + \epsilon_t.$$

Make noisy observations:

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

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

If $f(.)$ and $g(.)$ are linear, and ϵ_t and η_t are Gaussian, then solution is classical: Kalman filter. (Many generalizations available; e.g., (Huys and Paninski, 2009).)

Even Kalman case is challenging, since $d = \dim(\vec{V})$ is very large: computation of Kalman filter requires $O(d^3)$ computation per timestep

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

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 (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

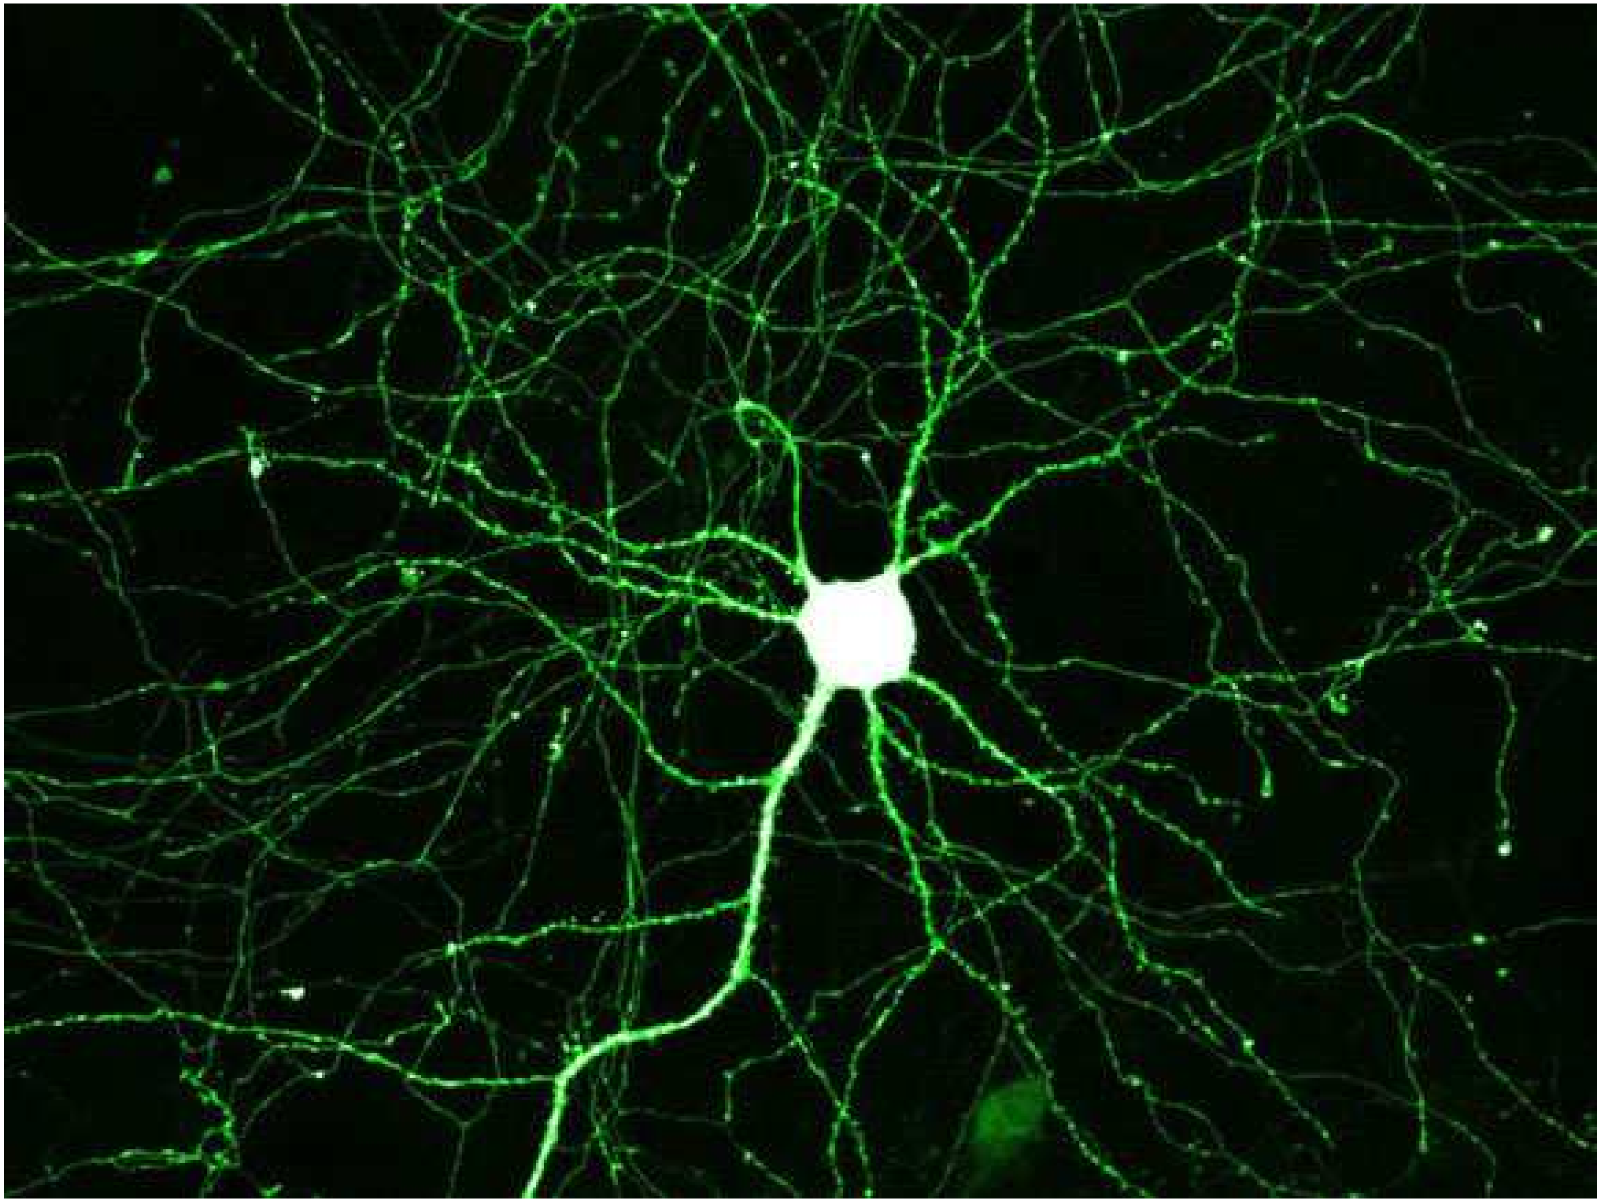

Application: synaptic locations/weights

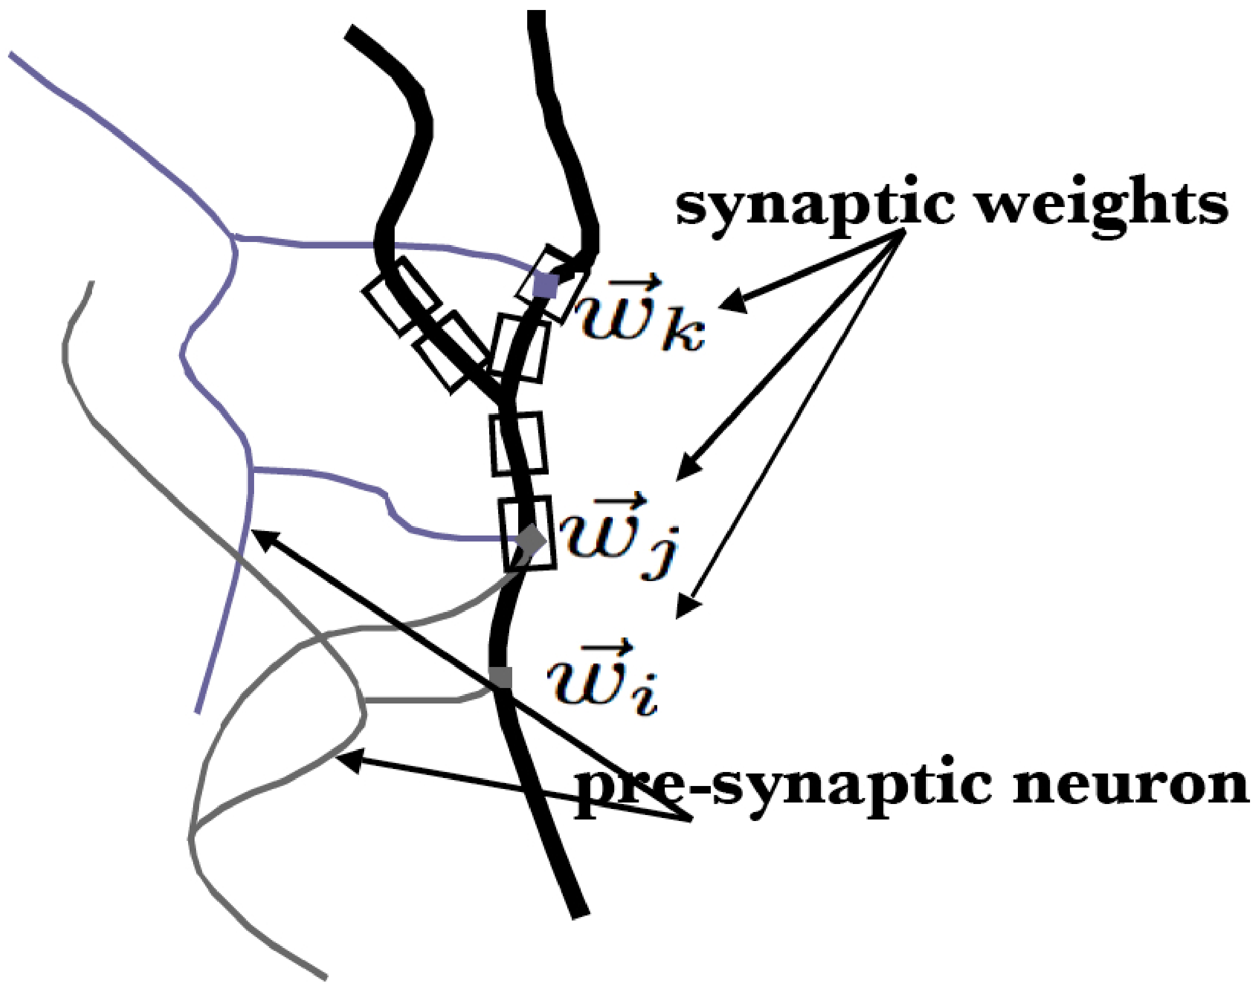

Cast as sparse regression problem $\Longrightarrow$ fast solution (Pakman et al., 2012)

Example: inferring dendritic synaptic maps

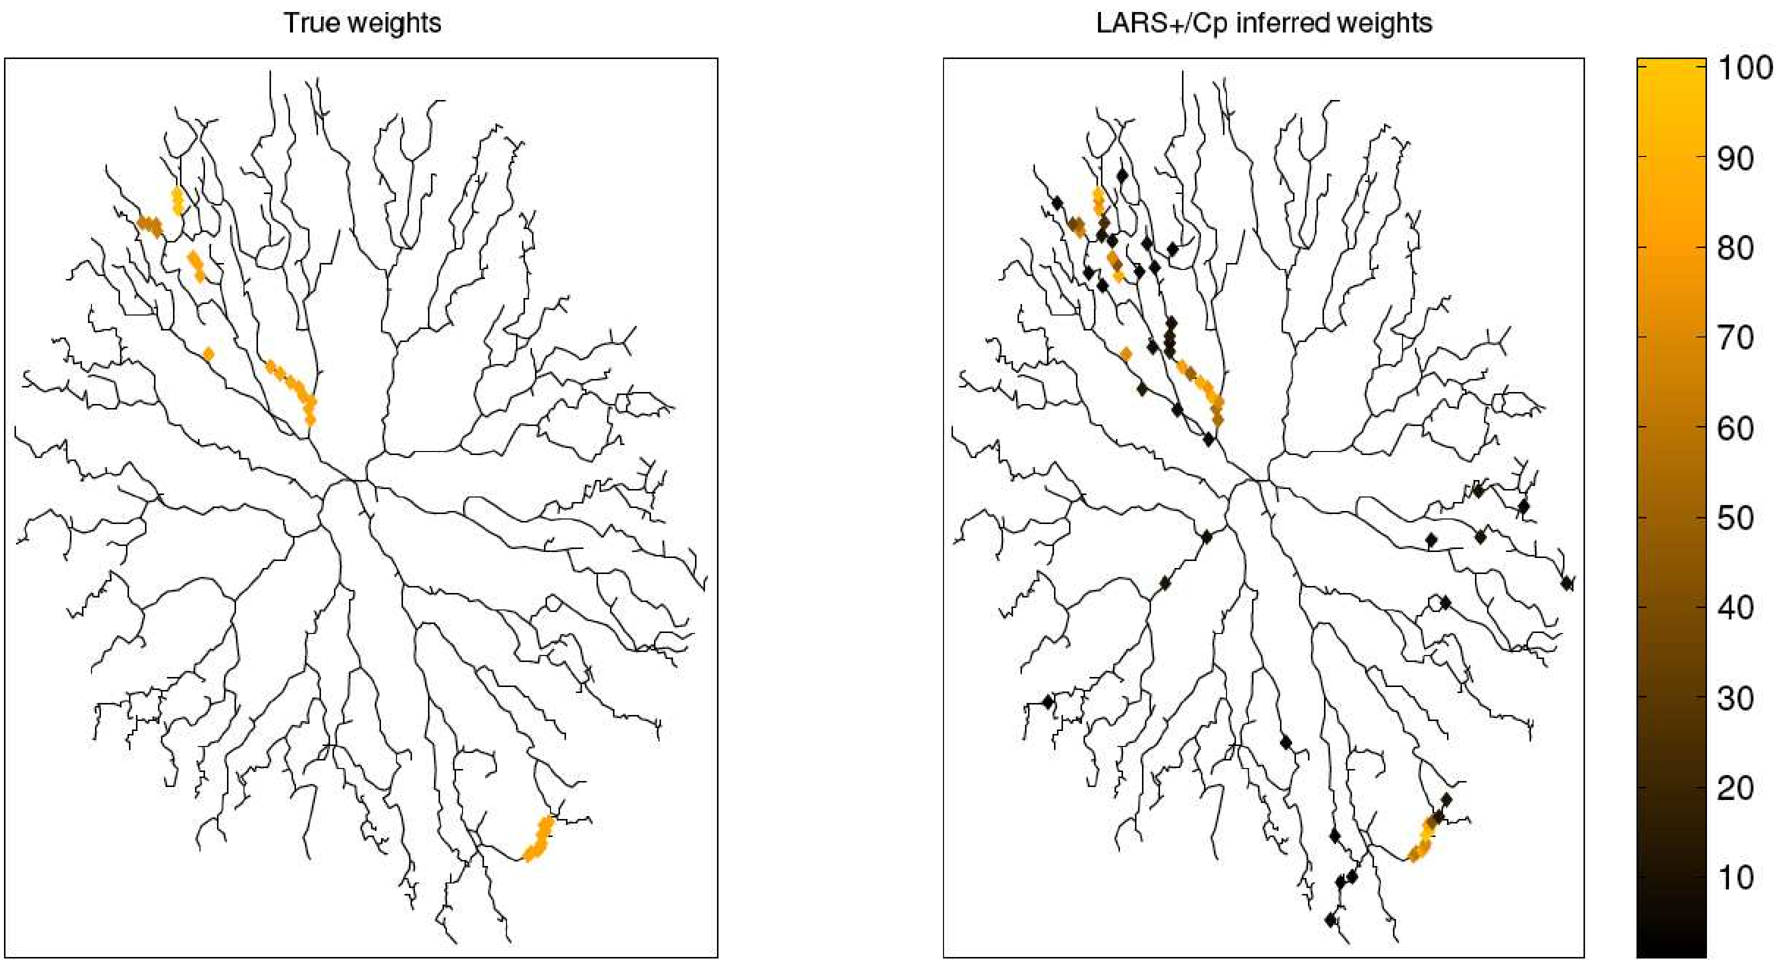

700 timesteps observed; 40 compartments (of > 2000) observed per timestep

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

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.

Field et al. (2010). Mapping a neural circuit: A complete input-output diagram in the primate retina. *Under review.*

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.

Lalor, E., Ahmadian, Y., and Paninski, L. (2009). The relationship between optimal and biologically plausible decoding of stimulus velocity in the retina. *Journal of the Optical Society of America A*, 26:25–42.

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. (2004). Maximum likelihood estimation of cascade point-process neural encoding models. *Network: Computation in Neural Systems*, 15:243–262.

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

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*, 29:107–126.

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.

Pfau, D., Pitkow, X., and Paninski, L. (2009). A Bayesian method to predict the optimal diffusion coefficient in random fixational eye movements. *Conference abstract: Computational and systems neuroscience.*

Pillow, J., Ahmadian, Y., and Paninski, L. (2010). Model-based decoding, information estimation, and change-point detection in multi-neuron spike trains. *In press, Neural Computation.*

Pillow, J., Shlens, J., Paninski, L., Sher, A., Litke, A., Chichilnisky, E., and Simoncelli, E. (2008).