

# Statistical methods for understanding neural coding and dynamics

Liam Paninski

Department of Statistics and Center for Theoretical Neuroscience  
Columbia University

<http://www.stat.columbia.edu/~liam>

*liam@stat.columbia.edu*

November 18, 2010

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)

# Some exciting open challenges


- inferring biophysical neuronal properties from noisy recordings
- reconstructing the full dendritic spatiotemporal voltage from noisy, subsampled observations
- estimating subthreshold voltage given superthreshold spike trains
- extracting spike timing from slow, noisy calcium imaging data
- reconstructing presynaptic conductance from postsynaptic voltage recordings
- inferring connectivity from large populations of spike trains
- decoding behaviorally-relevant information from spike trains
- optimal control of neural spike timing

— to solve these, we need to combine the two classical branches of computational neuroscience: dynamical systems and neural coding

# Part 1: modeling correlated spiking in retina

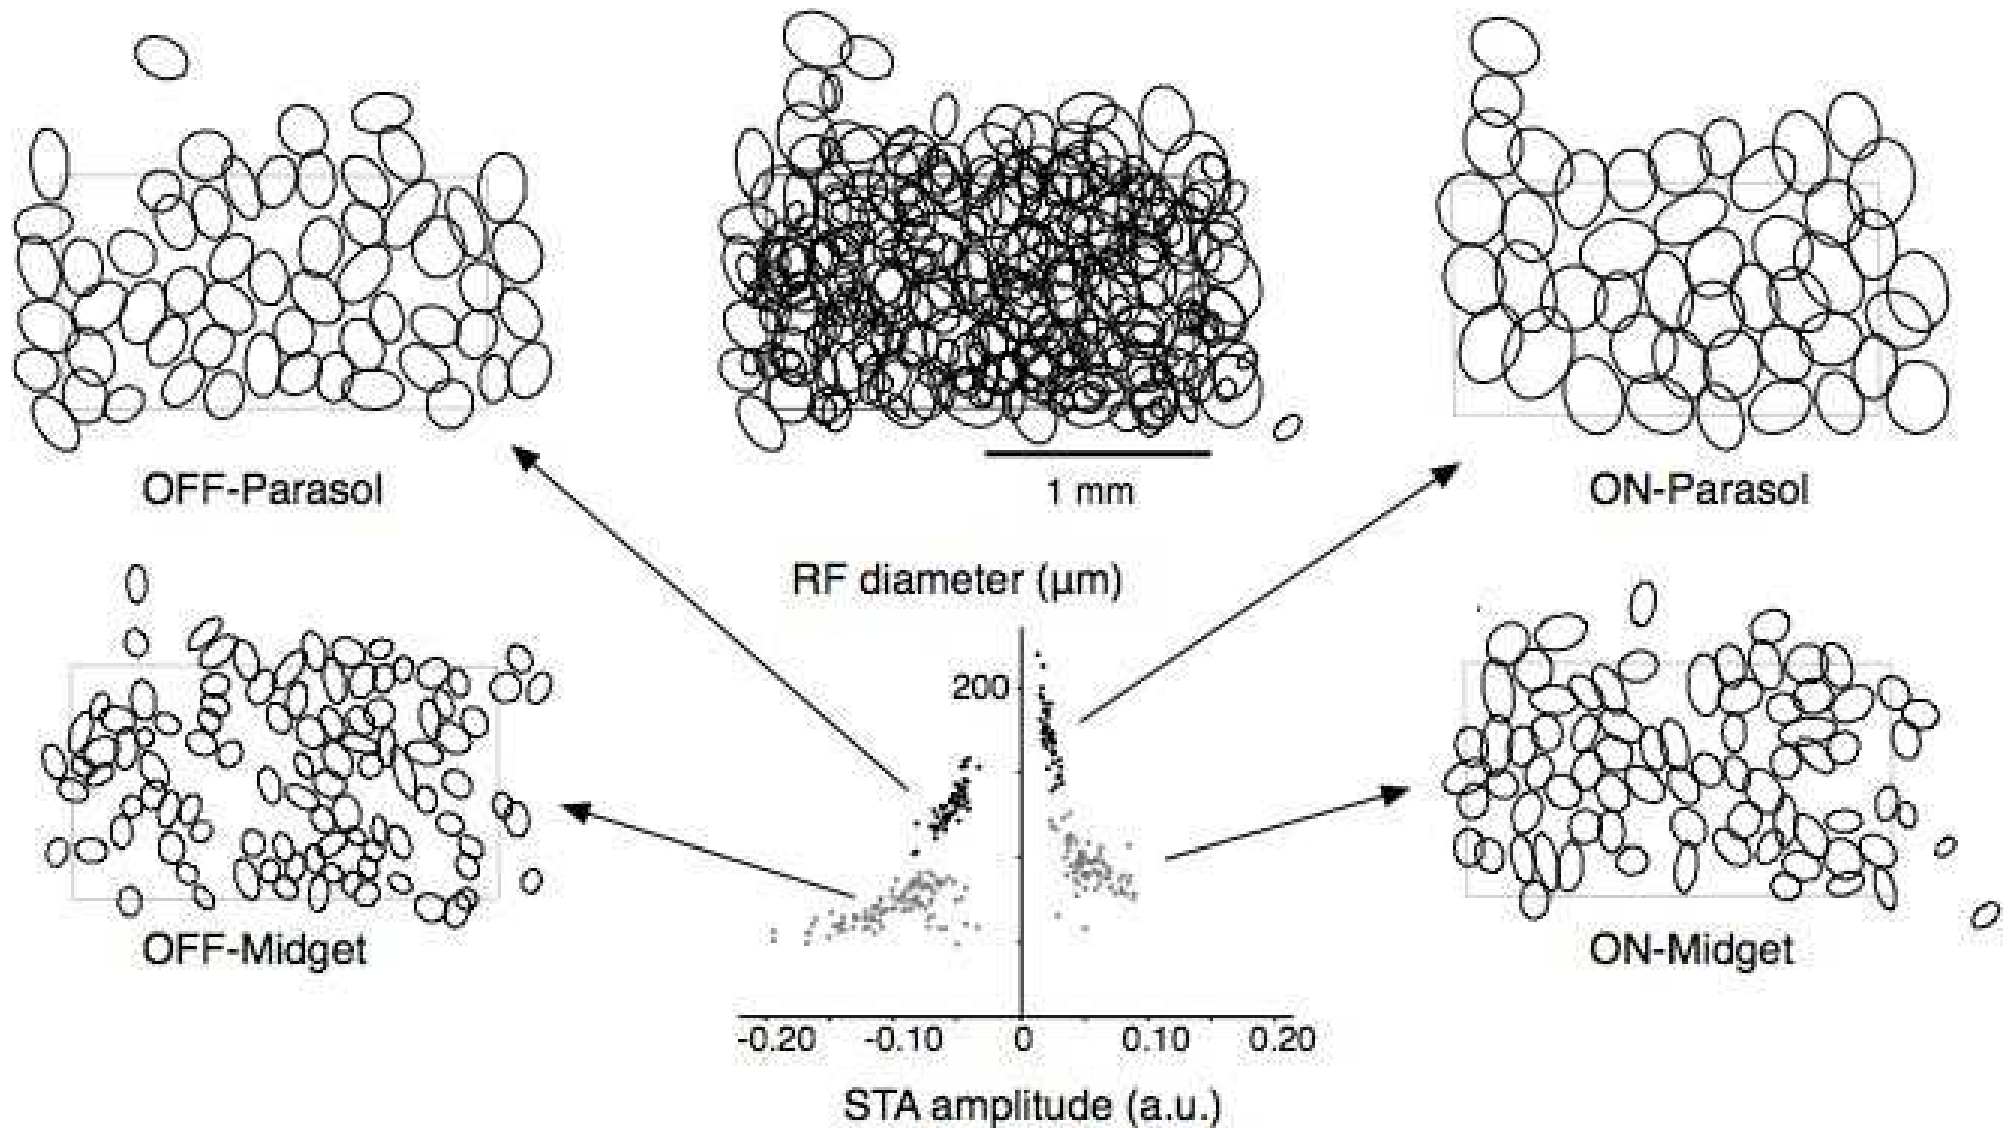
Preparation: dissociated macaque retina

— extracellularly-recorded responses of populations of RGCs

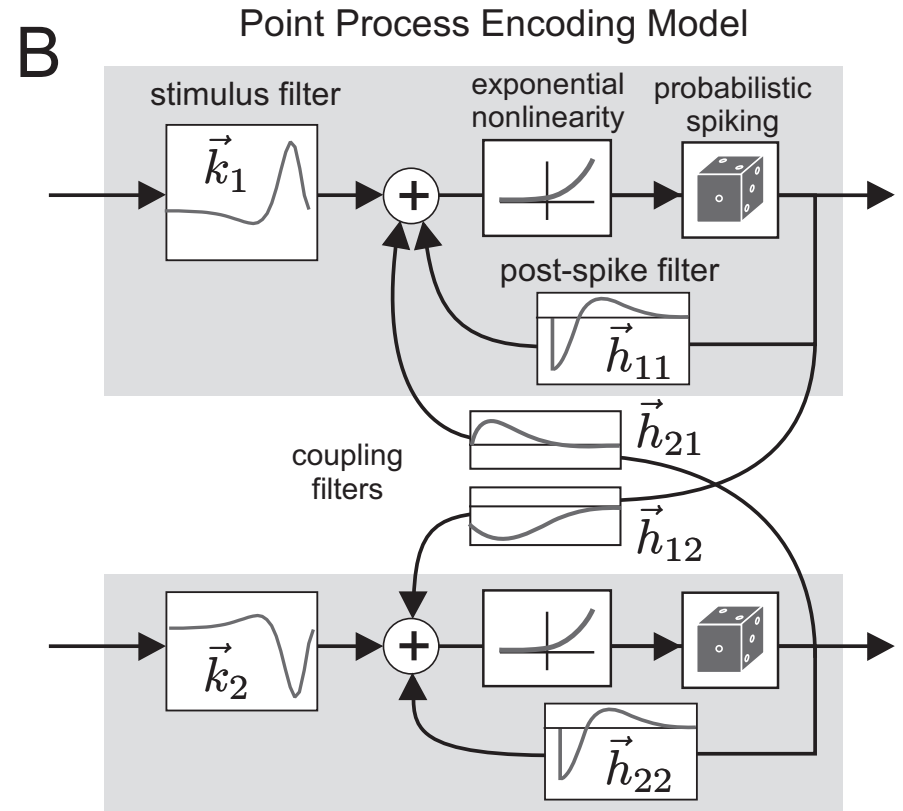
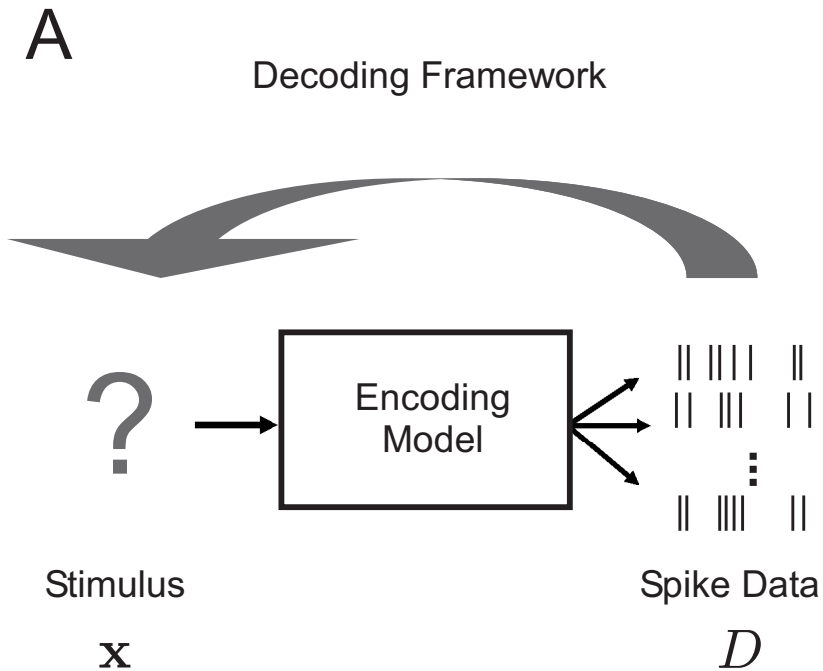


Stimulus: random spatiotemporal visual stimuli (Pillow et al., 2008)

# Receptive fields tile visual space



# Multineuronal point-process model




$$\lambda_i(t) = f \left( b_i + \vec{k}_i \cdot \vec{x}(t) + \sum_{i',j} h_{i',j} n_{i'}(t-j) \right),$$

- likelihood is easy to compute and to maximize (concave optimization) (Paninski, 2004; Paninski et al., 2007; Pillow et al., 2008)
- close connections to noisy integrate-and-fire model
- captures spike timing precision and details of spatiotemporal correlations in retinal ganglion cell network

# Reconsidering the model

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




## Pros:

- Tractable model-fitting and optimal decoding
- Captures response statistics

## Cons:

- Instantaneous coupling filters
- No explicit Common Input

# Considering common input effects




— universal problem in network analysis: can't observe all neurons!




## Intracellular findings:

- RGCs receive strongly correlated synaptic input in the absence of modulated light stimuli




Khuc Trong & Rieke Nature Neuro 2008

- ON RGCs are weakly electrically coupled

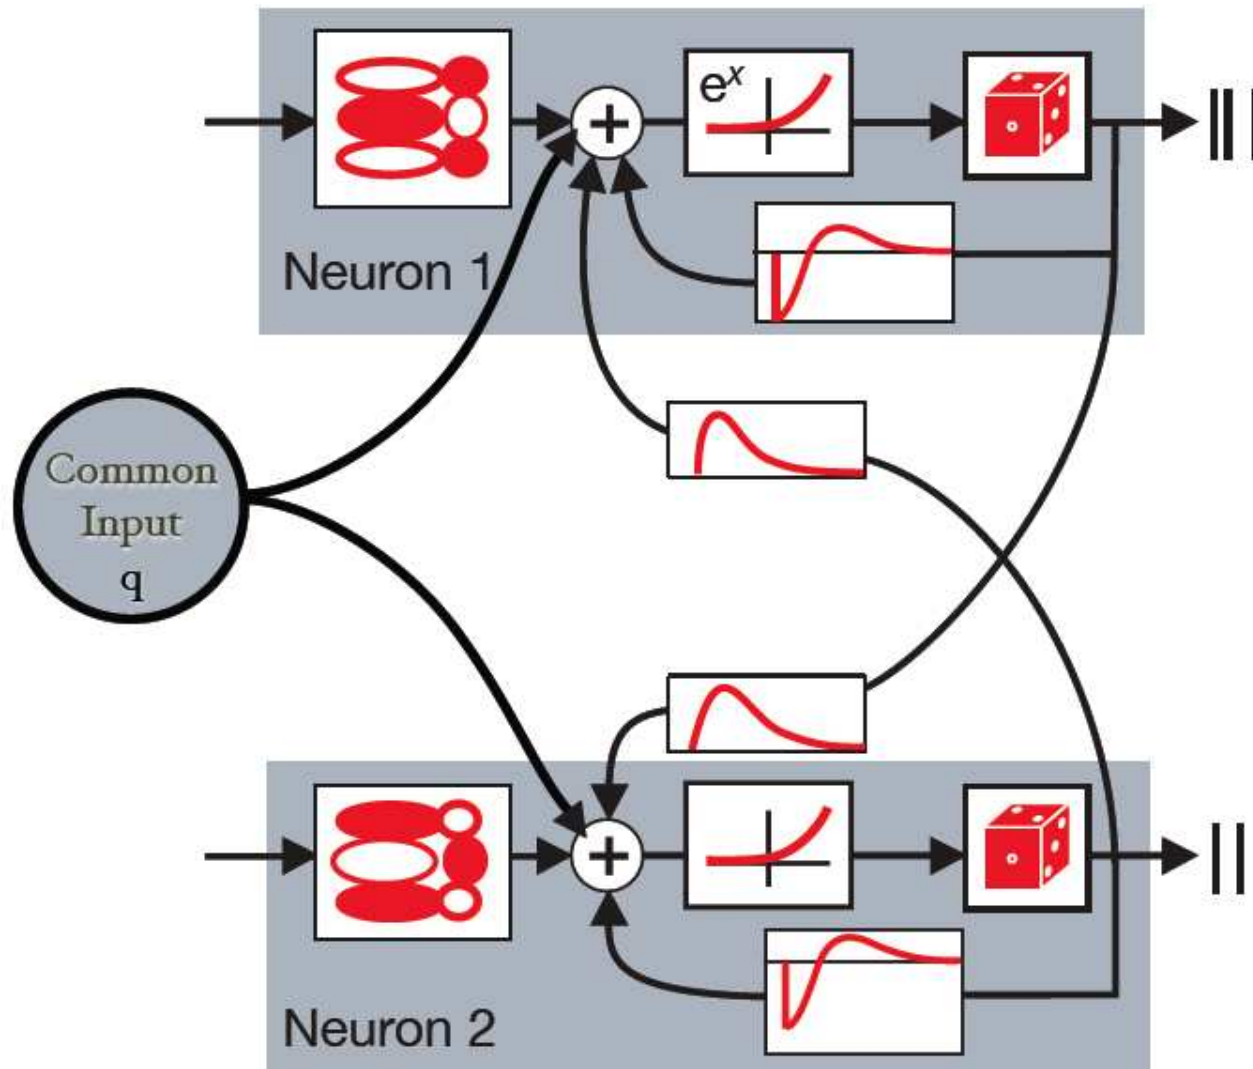


- No electrical coupling seen between OFF RGCs



# Extension: including common input effects

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



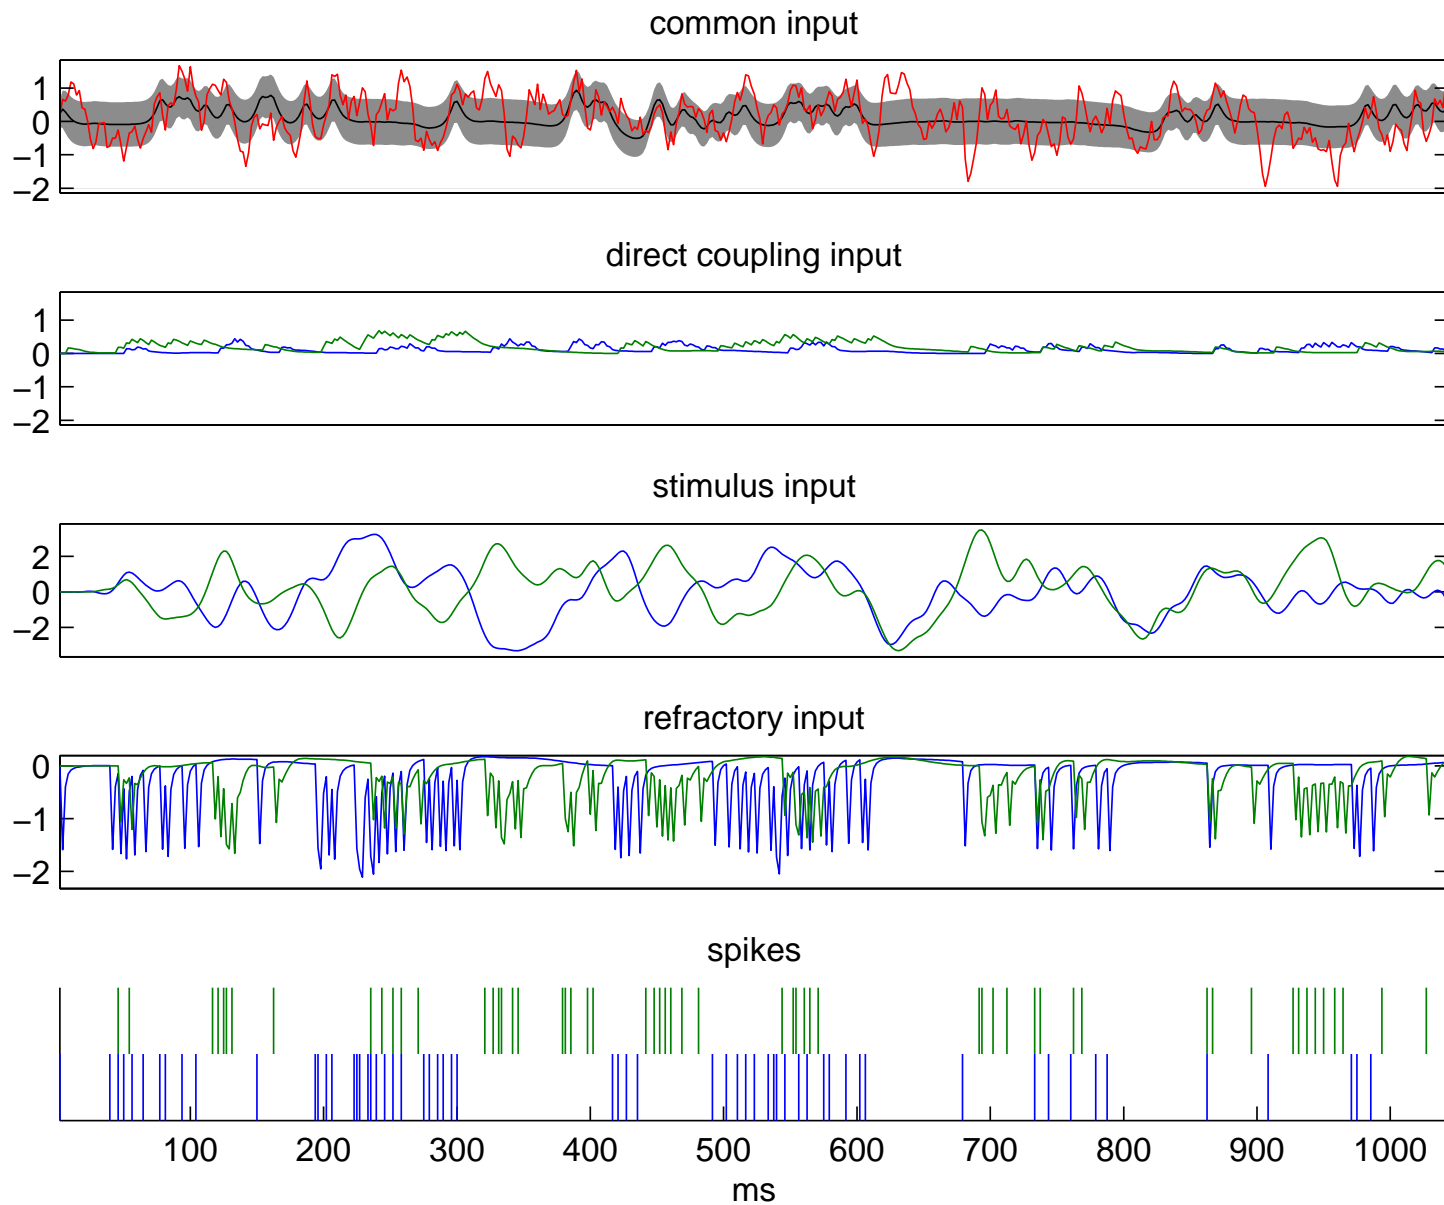
# Direct state-space optimization methods

To fit parameters, optimize approximate marginal likelihood:

$$\begin{aligned}\log p(\textit{spikes}|\theta) &= \log \int p(Q|\theta)p(\textit{spikes}|\theta, Q)dQ \\ &\approx \log p(\hat{Q}_\theta|\theta) + \log p(\textit{spikes}|\hat{Q}_\theta) - \frac{1}{2} \log |J_{\hat{Q}_\theta}| \\ \hat{Q}_\theta &= \arg \max_Q \{ \log p(Q|\theta) + \log p(\textit{spikes}|Q) \}\end{aligned}$$


- $Q$  is a very high-dimensional latent (unobserved) “common input” term. Taken to be a Gaussian process here with autocorrelation time  $\approx 5$  ms (Khuc-Trong and Rieke, 2008).
- correlation strength specified by one parameter per cell pair.
- all terms can be computed in  $O(T)$  via banded matrix methods (Paninski et al., 2010).

# Inferred common input effects are strong




— note that inferred direct coupling effects are now relatively small.

# Common-input-only model captures x-corrs



— single and triple-cell activities captured well, too (Vidne et al., 2009)

# Inferring cone locations




— cone locations and color identity can be inferred accurately with high spatial-resolution stimuli via maximum a posteriori estimates (Field et al., 2010).

# Inferring cone-to-RGC effective connectivity


*ON parasol*

*retina 1,*



*OFF parasol*

50 μm



RODS AND CONES

HORIZONTAL CELLS

BIPOLAR CELLS

AMACRINE CELLS


GANGLION CELLS

PLATINUM BLACK

SILICON NITRIDE


INDIUM TIN OXIDE

GLASS




# Part 2: applications to cortex


Record large-scale calcium movie




Aim 1: Extract spike times




Aim 2: Estimate network model



Aim 3: Check results via photostimulation




Inferred network model






# Model-based estimation of spike rates

## data



## schematic



## equations

$$F_t = \alpha C_t + \beta + \sigma \varepsilon_t, \quad \varepsilon_t \stackrel{iid}{\sim} \mathcal{N}(0, 1)$$
$$C_t = -(1 - \Delta/\tau)C_{t-1} + n_t$$
$$n_t \sim \text{poisson}(\lambda\Delta)$$

Note: each component here can be generalized easily (Vogelstein et al., 2009).

# Fast maximum a posteriori (MAP) filter

Start by writing out the posterior:

$$\begin{aligned}\log p(C|F) &= \log p(C) + \log p(F|C) + \text{const.} \\ &= \sum_t \log p(C_{t+1}|C_t) + \sum_t \log p(F_t|C_t) + \text{const.}\end{aligned}$$


Three basic observations:

- If  $\log p(C_{t+1}|C_t)$  and  $\log p(F_t|C_t)$  are concave, then so is  $\log p(C|F)$ .
- Hessian  $H$  of  $\log p(C|F)$  is tridiagonal:  $\log p(F_t|C_t)$  contributes a diag term, and  $\log p(C_{t+1}|C_t)$  contributes a tridiag term (Paninski et al., 2010).
- $C$  is a linear function of  $n$ .

Newton's method: iteratively solve  $HC_{dir} = \nabla$ . Tridiagonal solver requires  $O(T)$  time. Can include nonneg constraint  $n_t \geq 0$  (Koyama and Paninski, 2009).

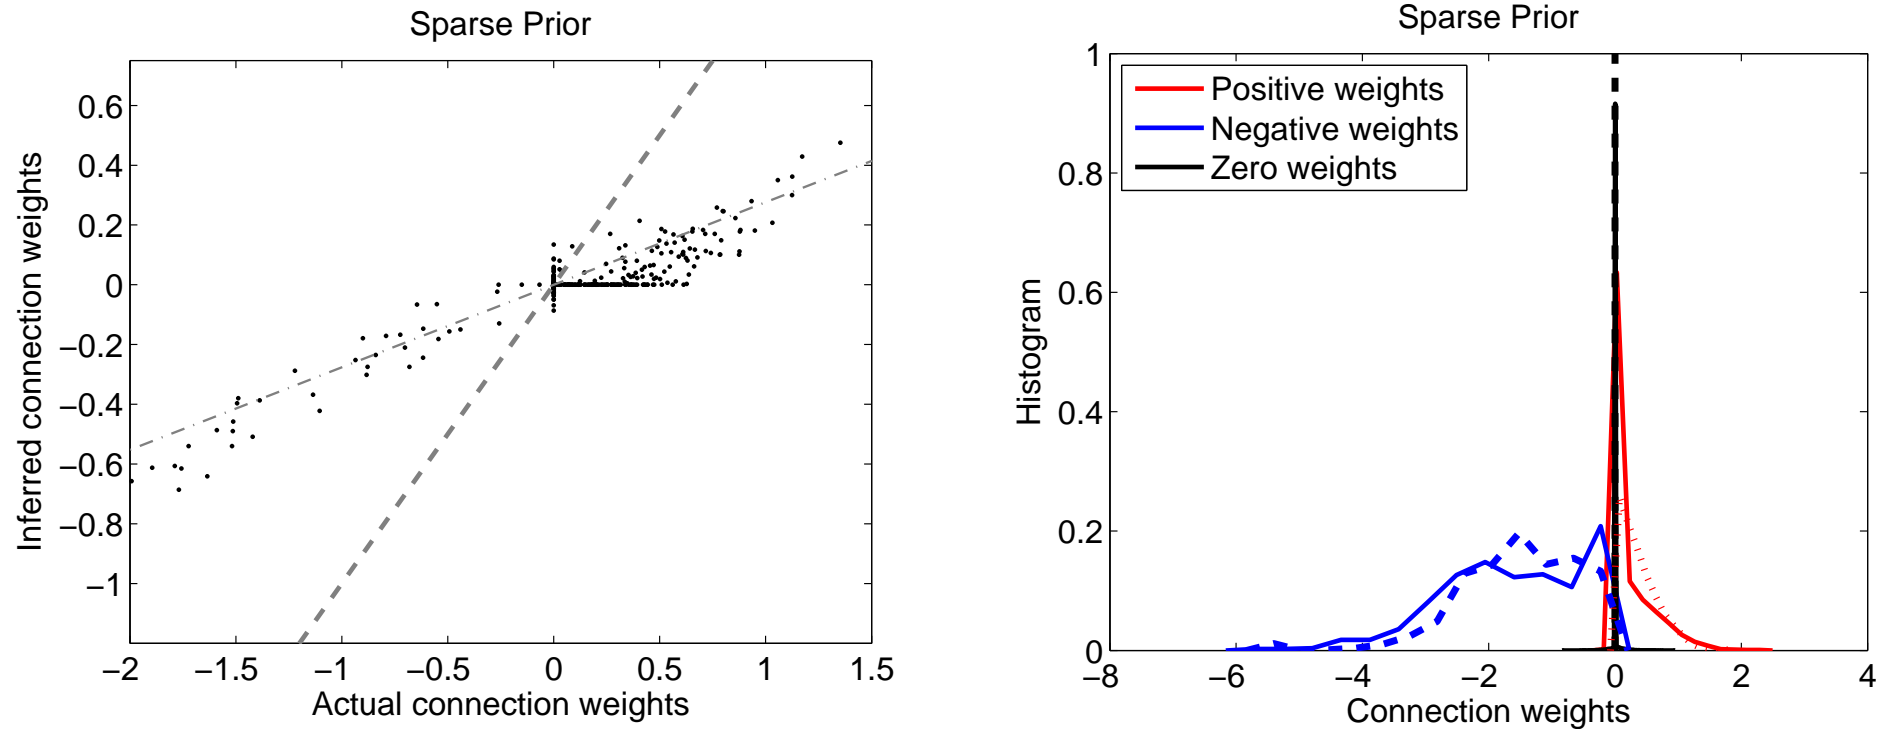
— Two orders of magnitude faster than particle filter: can process data from  $\approx 100$  neurons in real time on a laptop (Vogelstein et al., 2010).

# Example: nonnegative MAP filtering



— nonnegative deconvolution is a recurring problem (Vogelstein et al., 2010)  
(e.g., deconvolution of PSPs in intracellular recordings (Paninski et al., 2010))

# Simulated circuit inference



— conductance-based integrate-and-fire networks with biologically plausible connectivity matrices, imaging speed, SNR (Mishchenko et al., 2009).

Good news: MAP connections are inferred with the correct sign, in just a couple minutes of compute time, if network is fully observed. Current work focusing on improved Monte Carlo sampling methods, to better quantify uncertainty in unobserved neurons (Mishchenko and Paninski, 2010).

# Optimal control of spike timing

To test our results, we want 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:


$$\lambda_t = f(V_t + h_t)$$
$$V_{t+dt} = V_t + dt(-gV_t + aI_t) + \sqrt{dt}\sigma\epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, 1).$$

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 tridiagonal  $\implies O(T)$  optimization.

— again, can be done in real time (Ahmadian et al., 2010).

# Simulated electrical control of spike timing



# Example: intracellular control of spike timing




(Ahmadian et al., 2010)

# Optical conductance-based control of spiking

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


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





# Part 3: spatiotemporal filtering on dendrites

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; Knopfel et al., 2006)

# Basic paradigm: the Kalman filter

Variable of interest,  $q_t$ , evolves according to a noisy differential equation (Markov process):

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

Make noisy observations:

$$y_t = g(q_t) + \eta_t.$$

We want to infer  $E(q_t|Y)$ : optimal estimate given observations.

Problem: Kalman filter requires  $O(d^3T)$  time ( $d = \dim(q)$ ).

Reduction to  $O(dT)$ : exploit tree structure of dendrite (Paninski, 2010). Can be applied to voltage- or calcium-sensitive imaging data (Pnevmatikakis et al, 2010).

# 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? (Huggins and Paninski, 2010)
- 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 (Huggins and Paninski, 2011)

# Application: synaptic locations/weights



# Application: synaptic locations/weights


Including known terms:

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

$U_j(t)$  = 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 et al., 2010; Huggins and Paninski, 2011)



# Conclusions

- GLM and state-space approaches provide flexible, powerful methods for answering key questions in neuroscience
- Concave optimizations, banded matrix methods make computations very tractable — real-time, in many cases
- Co-development of experiment and analysis: exciting time for statistical neuroscience

# 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.
- Field et al. (2010). Mapping a neural circuit: A complete input-output diagram in the primate retina. *Under review*.
- Huggins, J. and Paninski, L. (2010). Optimal experimental design for sampling voltage on dendritic trees. *Under review*.
- Huggins, J. and Paninski, L. (2011). A fast method for detecting synapse locations on dendritic trees. *In preparation*.
- Huys, Q., Ahrens, M., and Paninski, L. (2006). Efficient estimation of detailed single-neuron models. *Journal of Neurophysiology*, 96:872–890.
- 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.
- Koyama, S. and Paninski, L. (2009). Efficient computation of the map path and parameter estimation in integrate-and-fire and more general state-space models. *Journal of Computational Neuroscience*, In press.
- Mishchenko, Y. and Paninski, L. (2010). Efficient methods for sampling spike trains in networks of coupled neurons. *In preparation*.
- 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 dendritic trees. *Journal of Computational Neuroscience*, In press.
- 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., 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.
- Pillow, J., Shlens, J., Paninski, L., Sher, A., Litke, A., Chichilnisky, E., and Simoncelli, E. (2008). Spatiotemporal correlations and visual signaling in a complete neuronal population. *Nature*, 454:995–999.
- Vidne, M., Kulkarni, J., Ahmadian, Y., Pillow, J., Shlens, J., Chichilnisky, E., Simoncelli, E., and Paninski, L. (2009). Inferring functional connectivity in an ensemble of retinal ganglion cells sharing a common