Coding and computation by neural ensembles in the primate retina

Liam Paninski

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

June 7, 2010

co-PI's: E. Simoncelli (NYU), E.J. Chichilnisky (Salk)

— with J. Pillow (UT Austin), G. Field, J. Gauthier, J. Shlens (Salk), A. Litke (UCSC), E. Lalor (TC Dublin), S. Koyama (CMU), **Y. Ahmadian**, J. Kulkarni, H. Liu, T. Machado, D. Pfau, X. Pitkow, **M. Vidne** (Columbia).

Retinal ganglion neuronal data

Preparation: dissociated macaque retina

— extracellularly-recorded responses of populations of RGCs

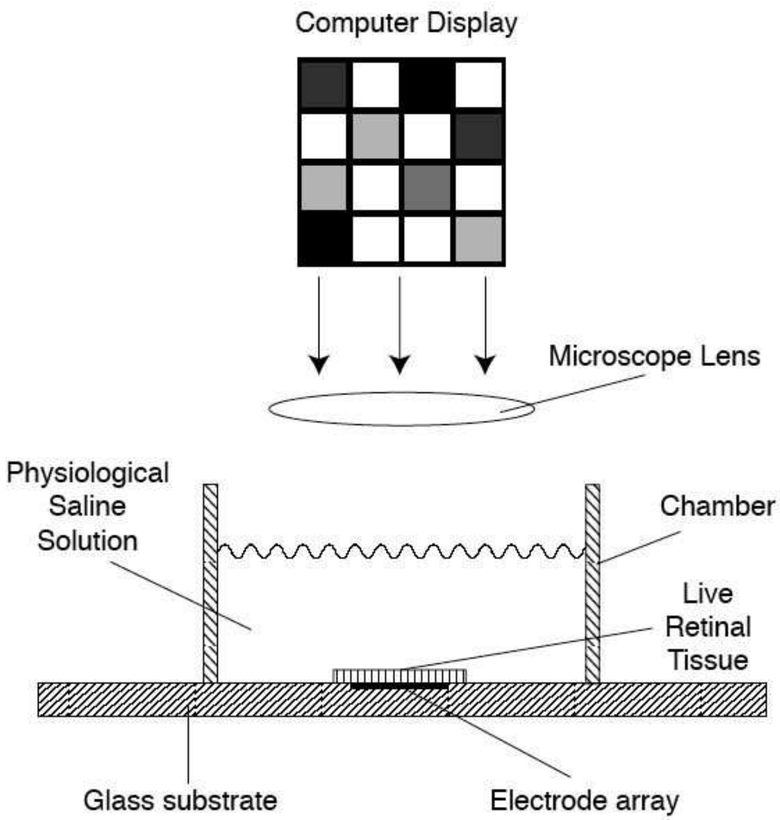

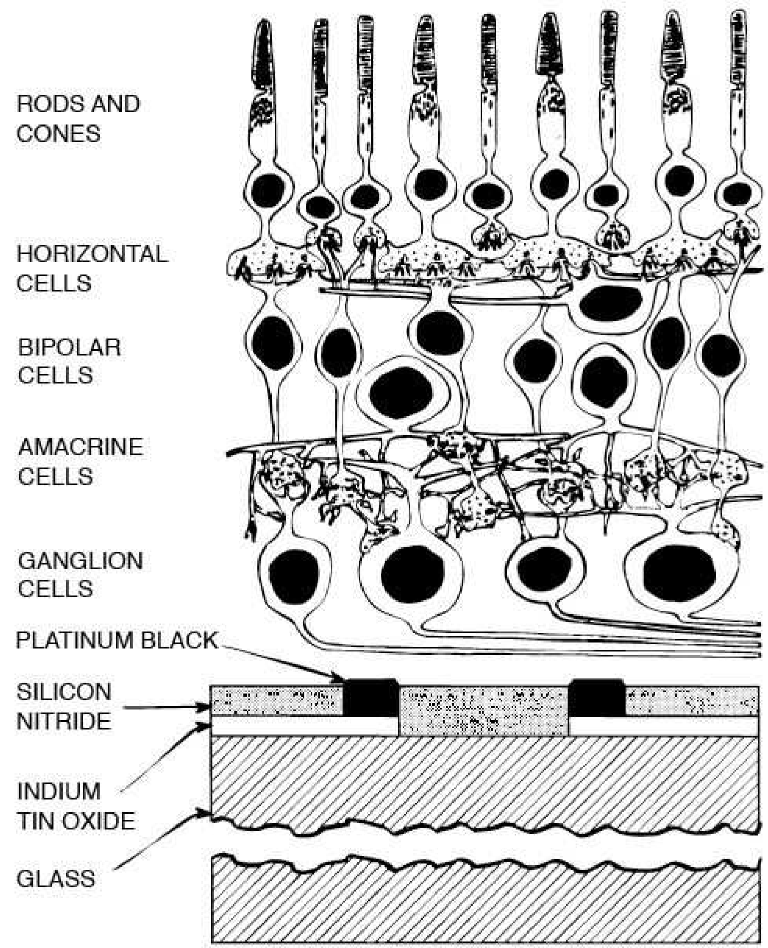

Stimulus: random spatiotemporal visual stimuli (Pillow et al., 2008)

Receptive fields tile visual space

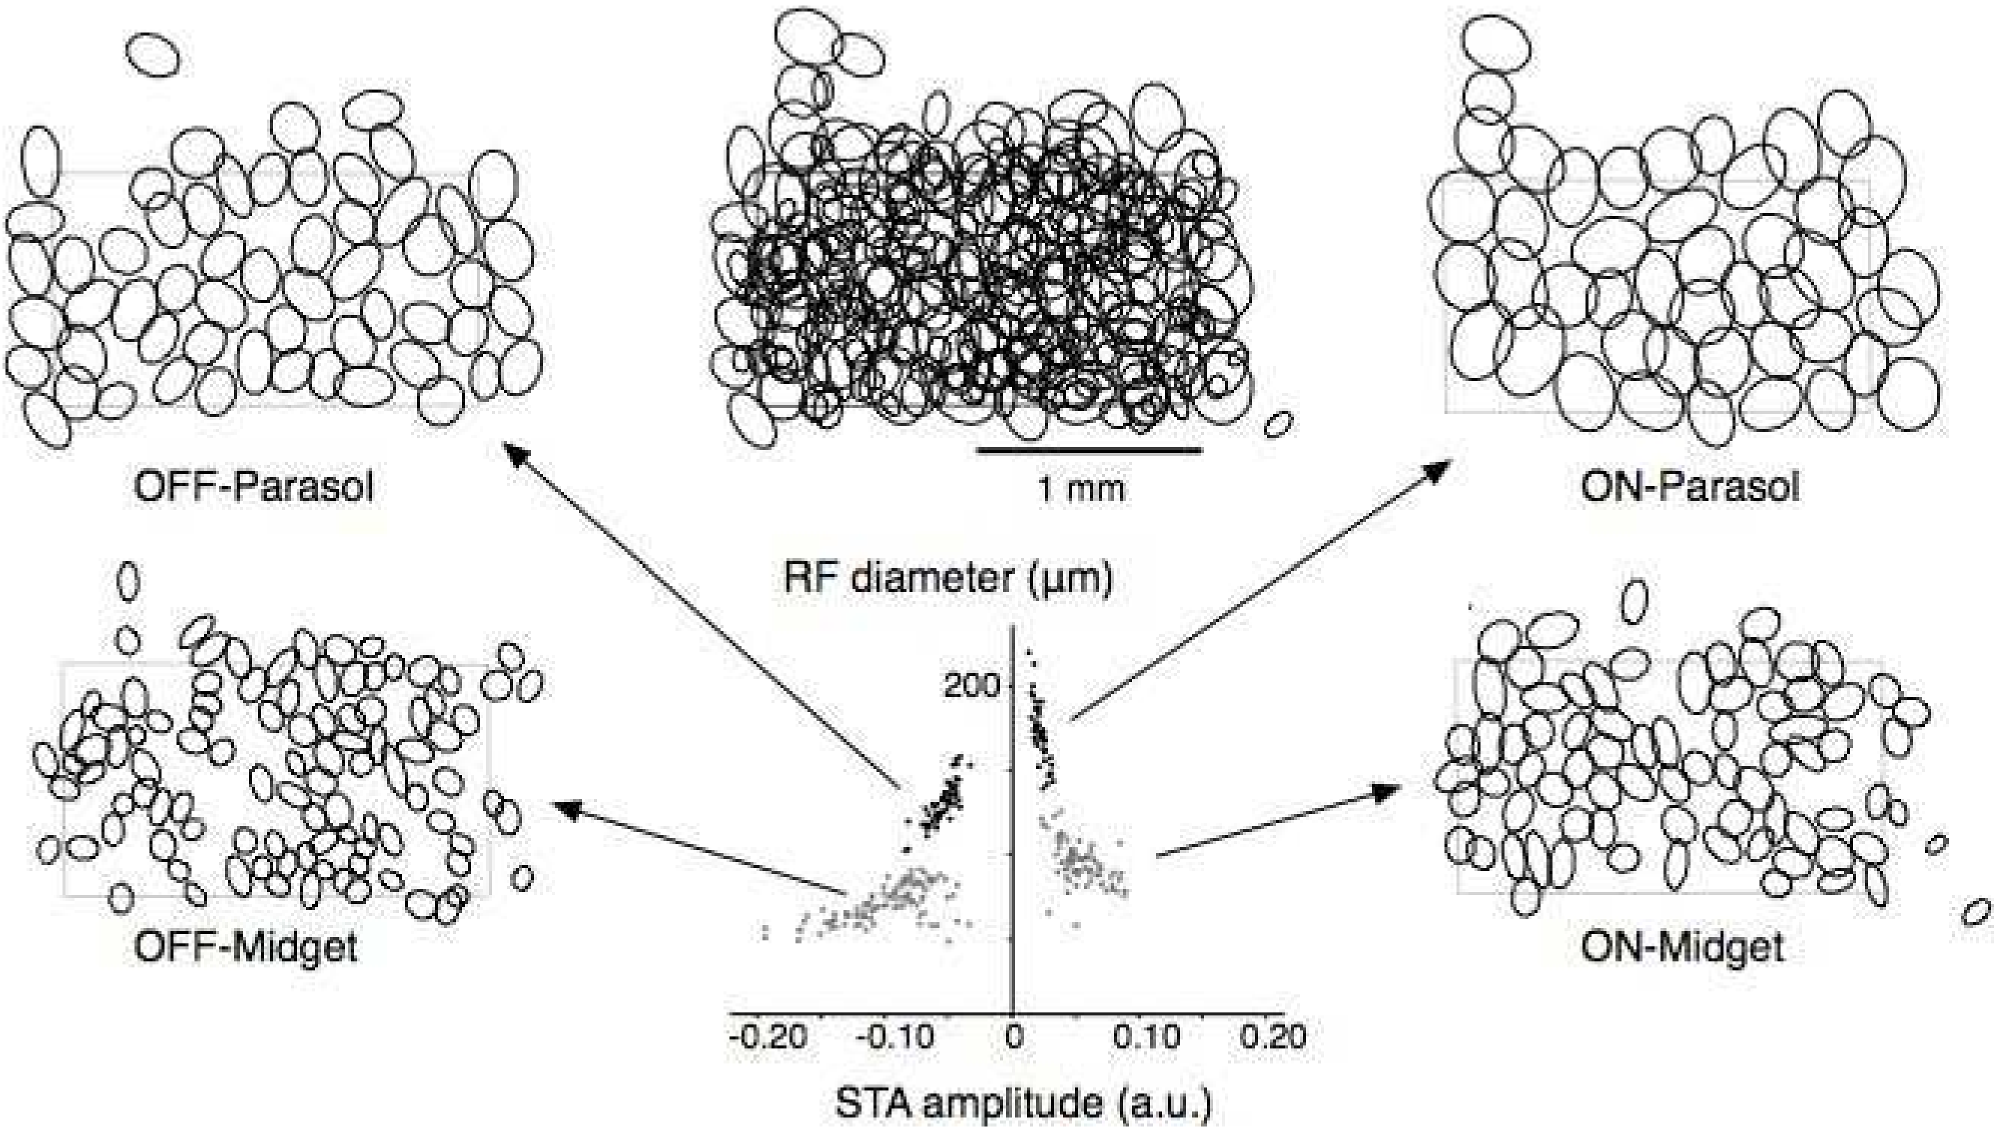

Multineuronal point-process model

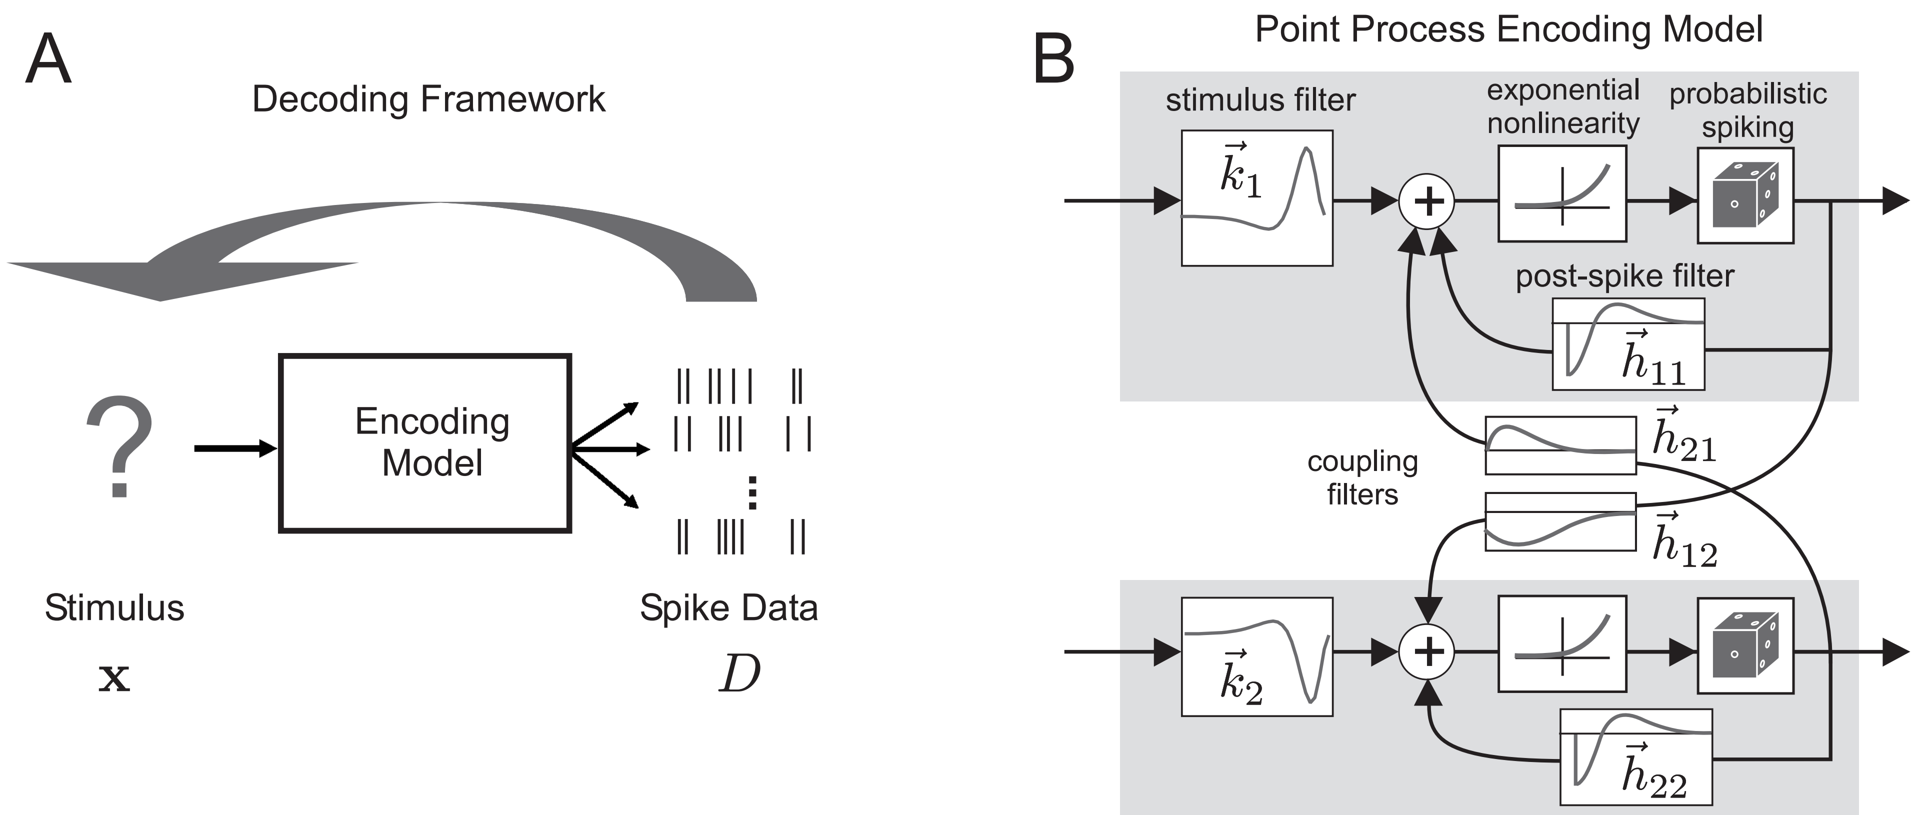

$$\lambda_i(t) = f\left(b_i + \vec{k}_i \cdot \vec{x}(t) + \sum_{i',j} h_{i',j} n_{i'}(t-j)\right),$$

— 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

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; Ahmadian et al., 2010).

Optimal Bayesian decoding

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

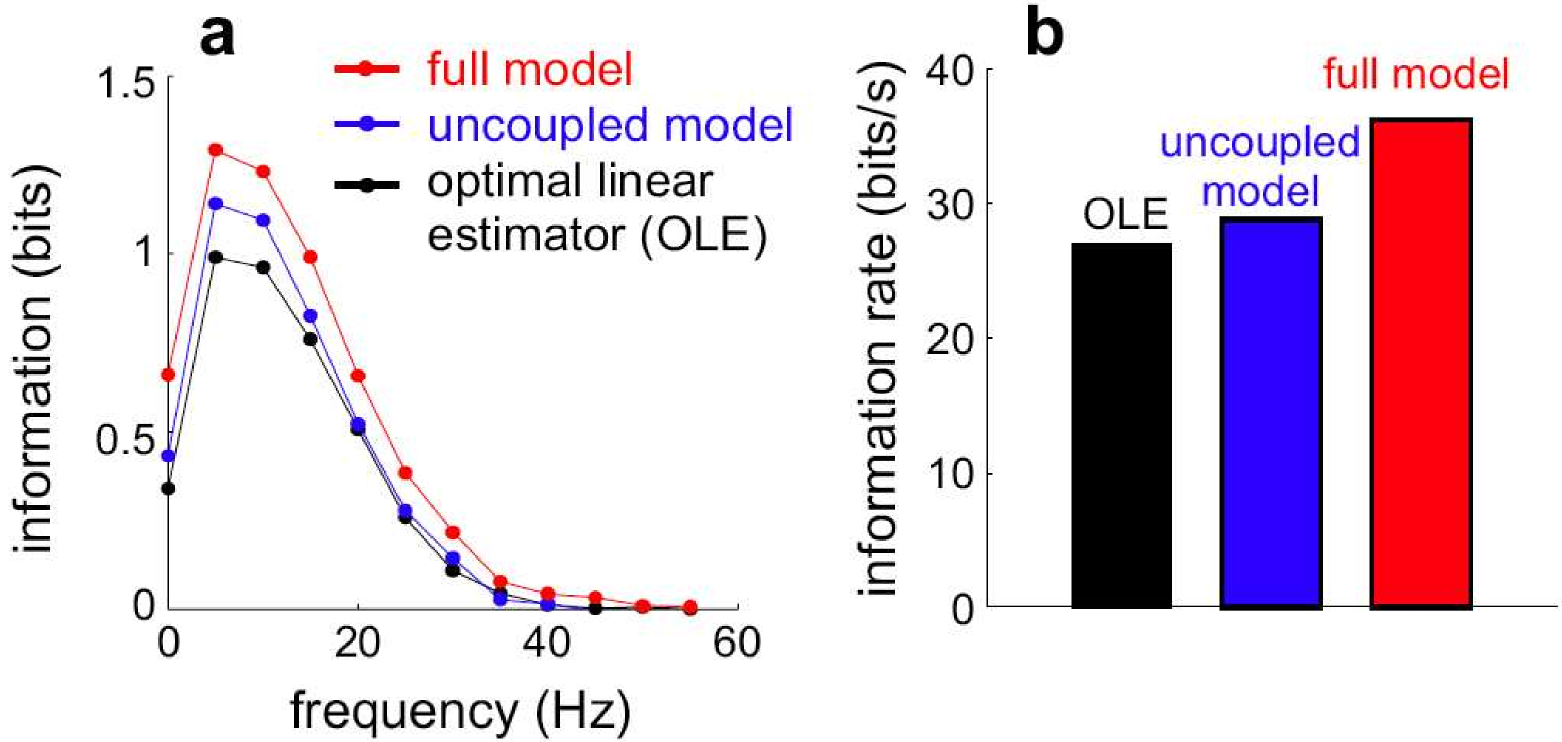

— Computational points:

- $\log P(spikes|\vec{x})$ 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; Ahmadian et al., 2010).

— Biological point: paying attention to correlations improves decoding accuracy.

Application: how important is timing?

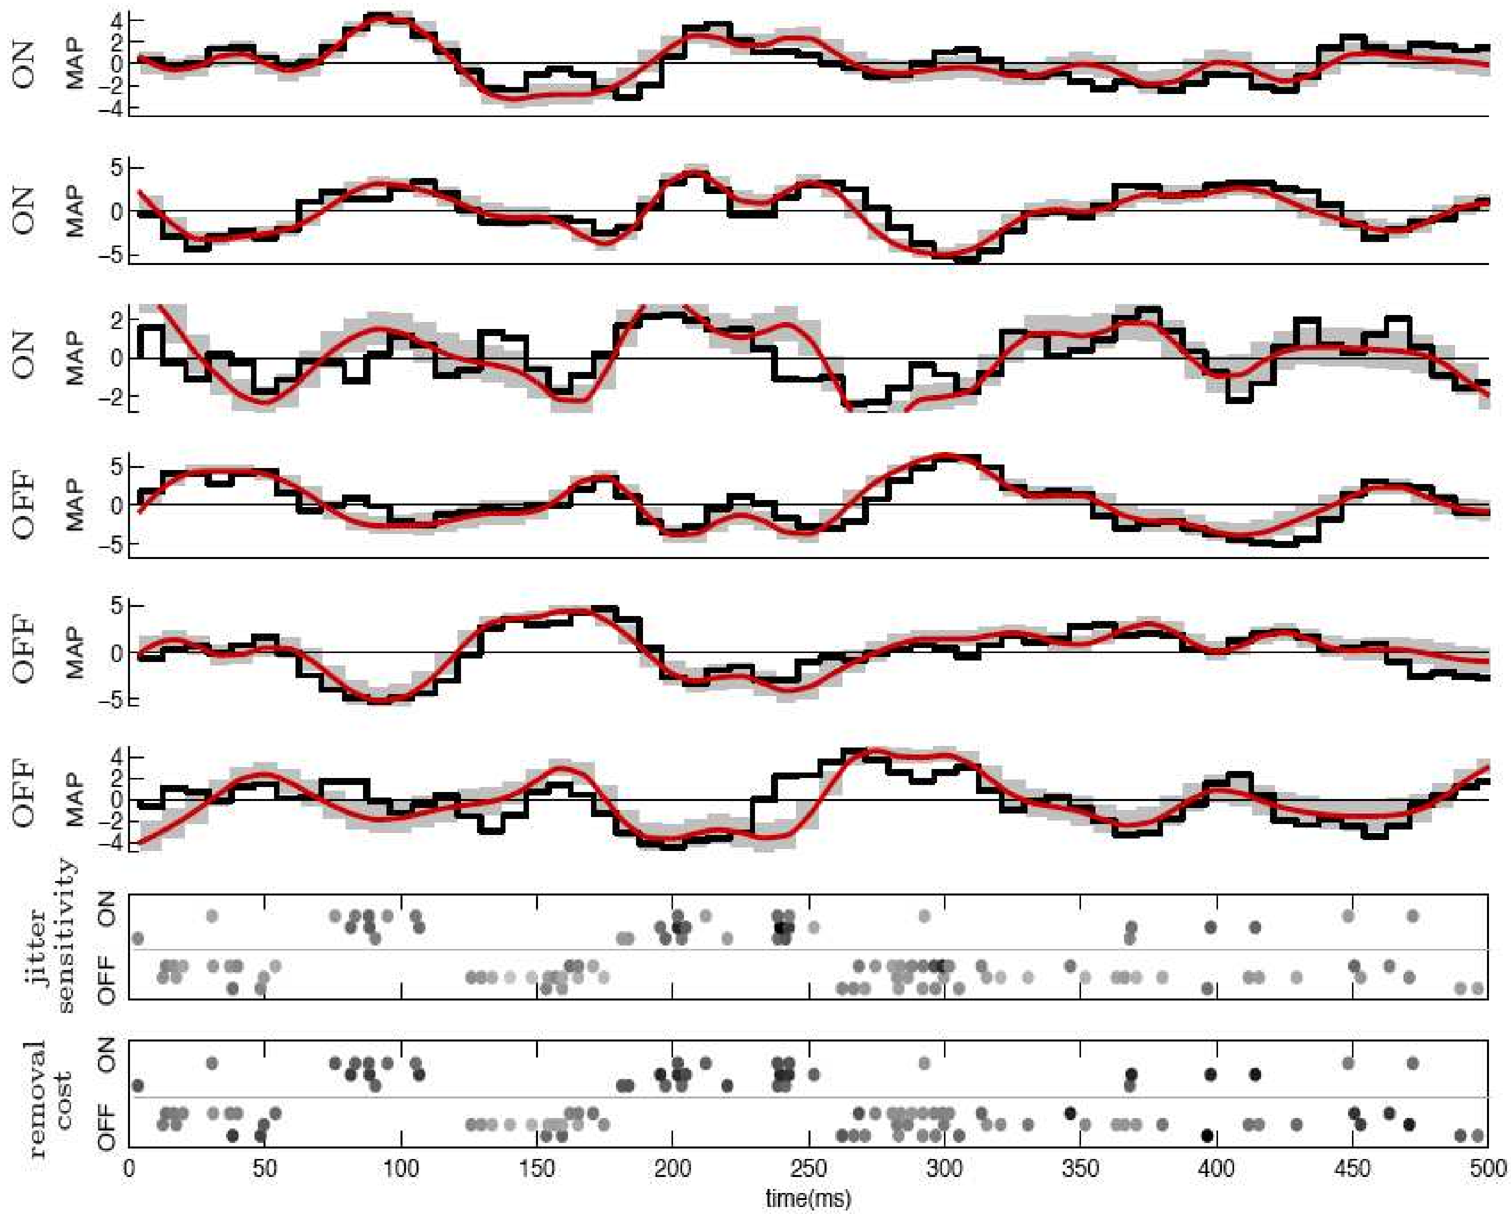

— further applications: decoding velocity signals (Lalor et al., 2009), tracking images perturbed by eye jitter (Pfau et al., 2009)

Next steps: reconsidering the model

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

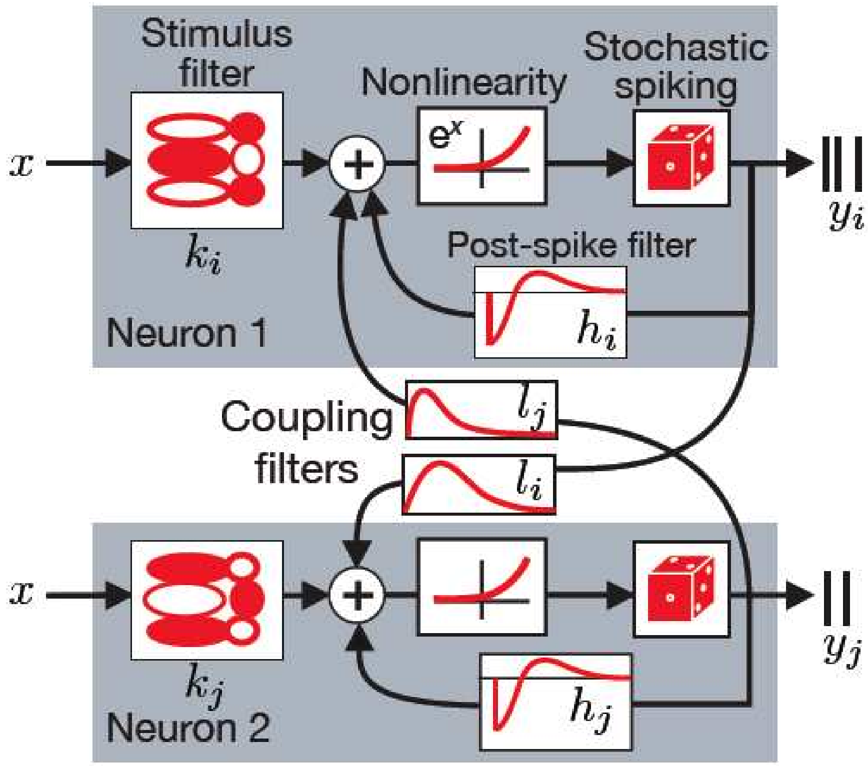

Pros:

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

Cons:

- Instantaneous coupling filters
- No explicit Common Input

Considering common input effects

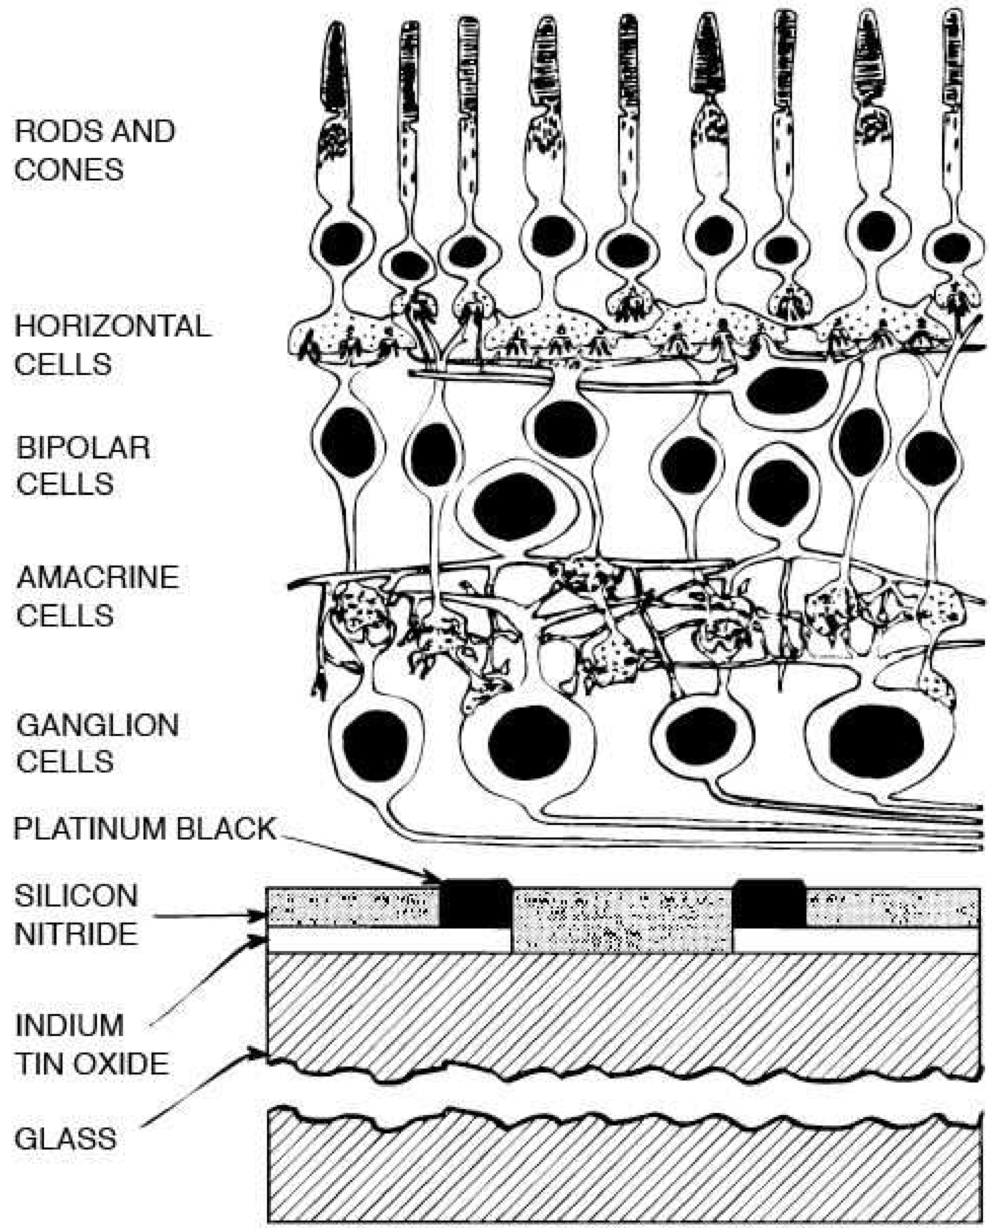

— universal problem in network analysis: can't observe all neurons!

Intracellular findings:

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

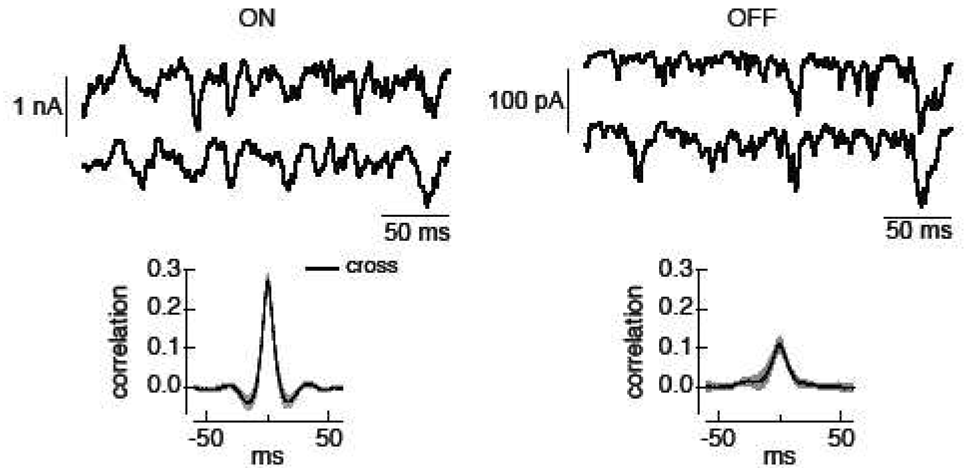

Khuc Trong & Rieke Nature Neuro 2008

- ON RGCs are weakly electrically coupled

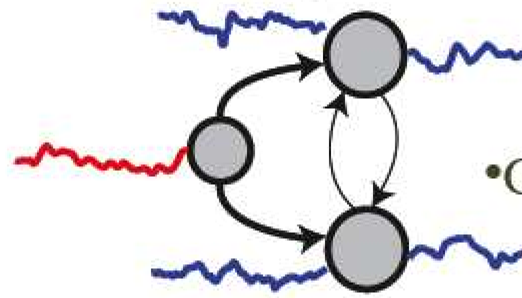

•ON RGCs

- No electrical coupling seen between OFF RGCs

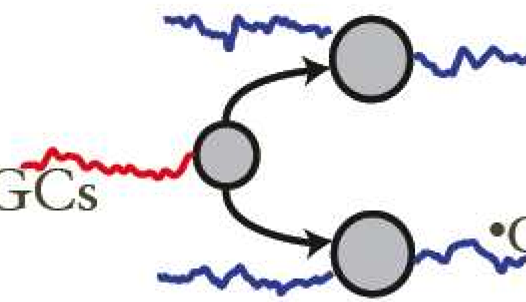

•OFF RGCs

Extension: including common input effects

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

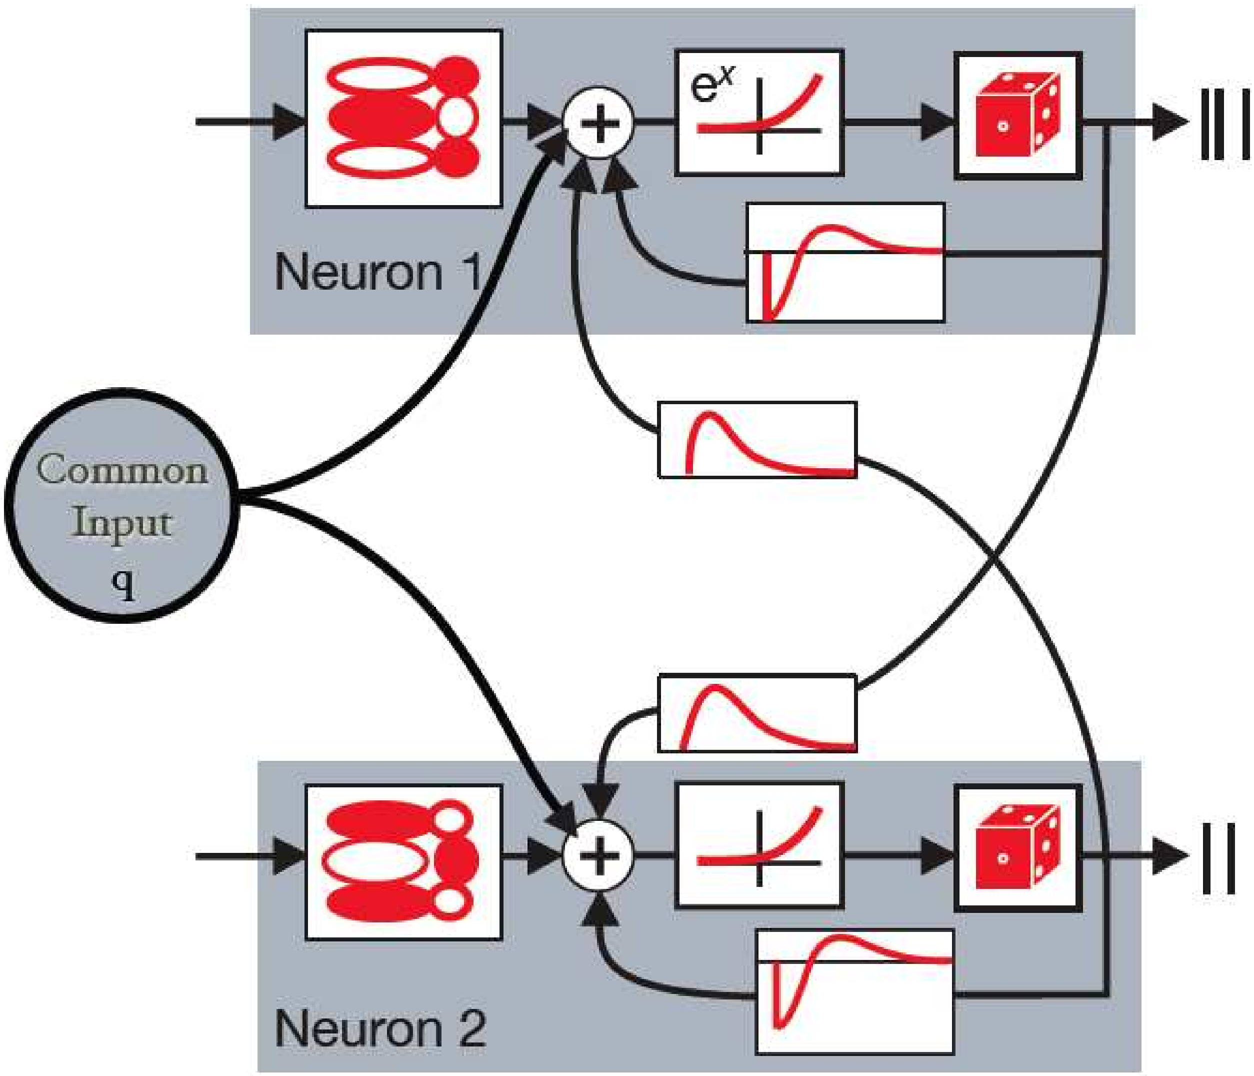

Direct state-space optimization methods

To fit parameters, optimize approximate marginal likelihood:

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

— Q is a very high-dimensional latent (unobserved) "common input" term. Taken to be a Gaussian process here with autocorrelation time ≈ 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

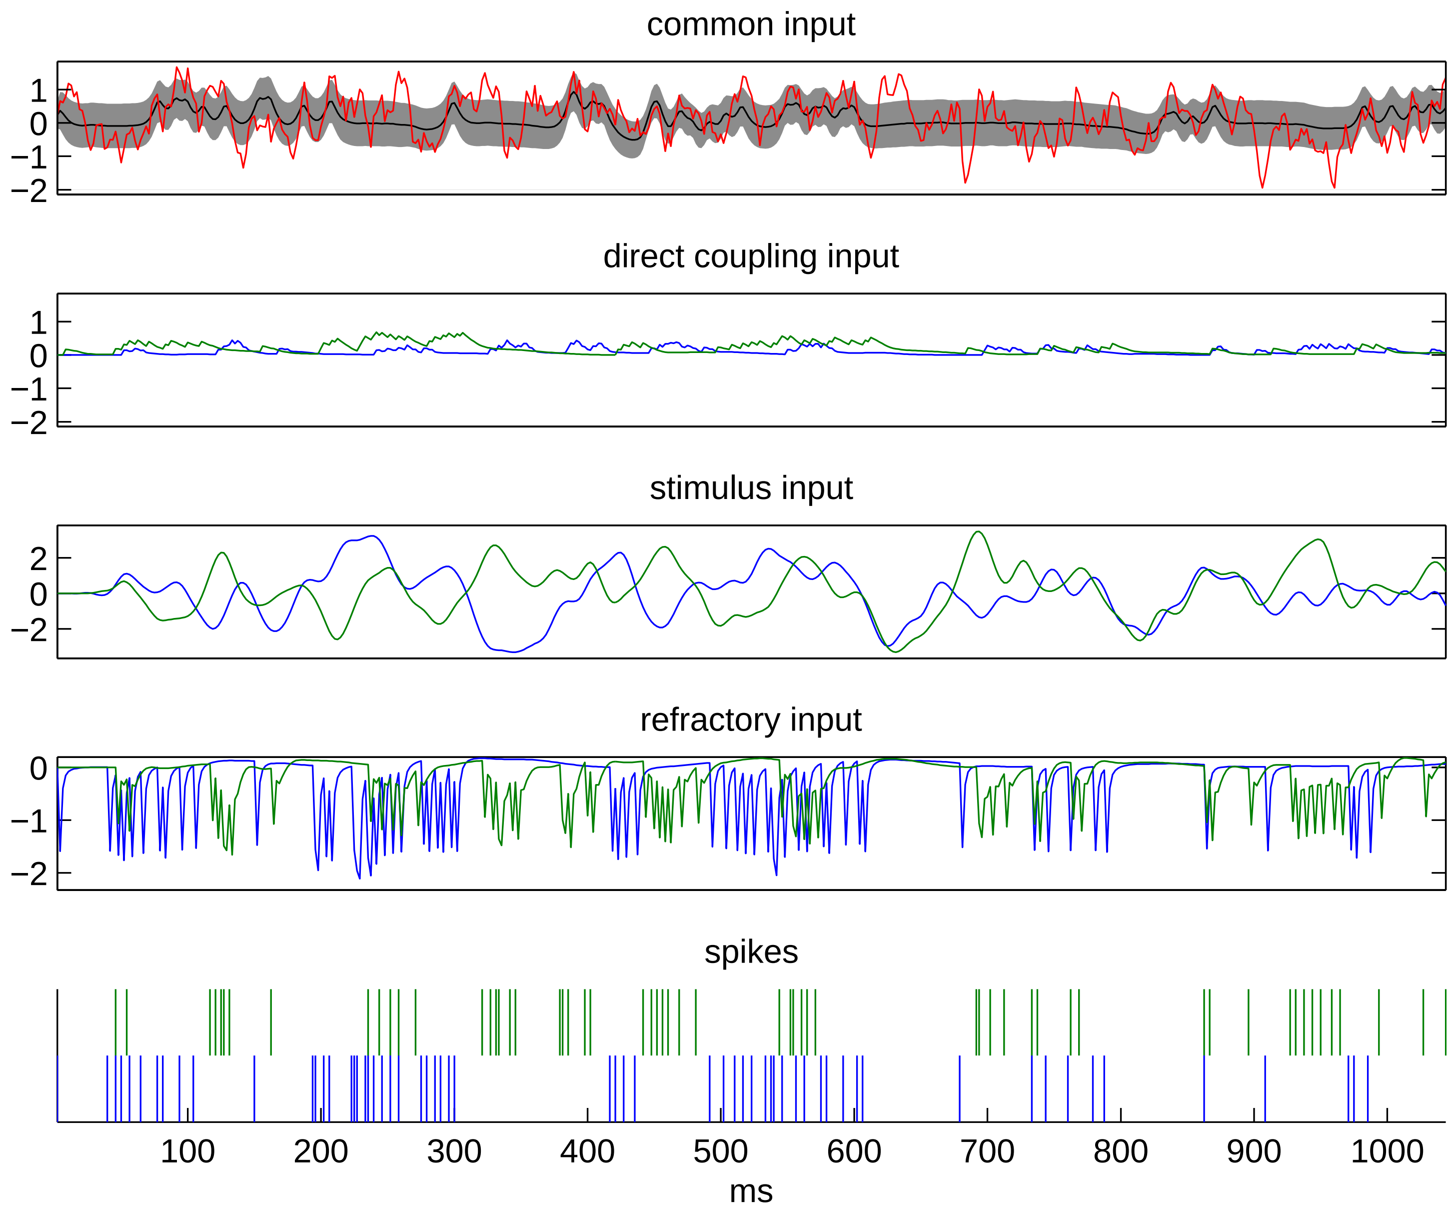

— note that inferred direct coupling effects are now relatively small.

Common-input-only model captures x-corrs

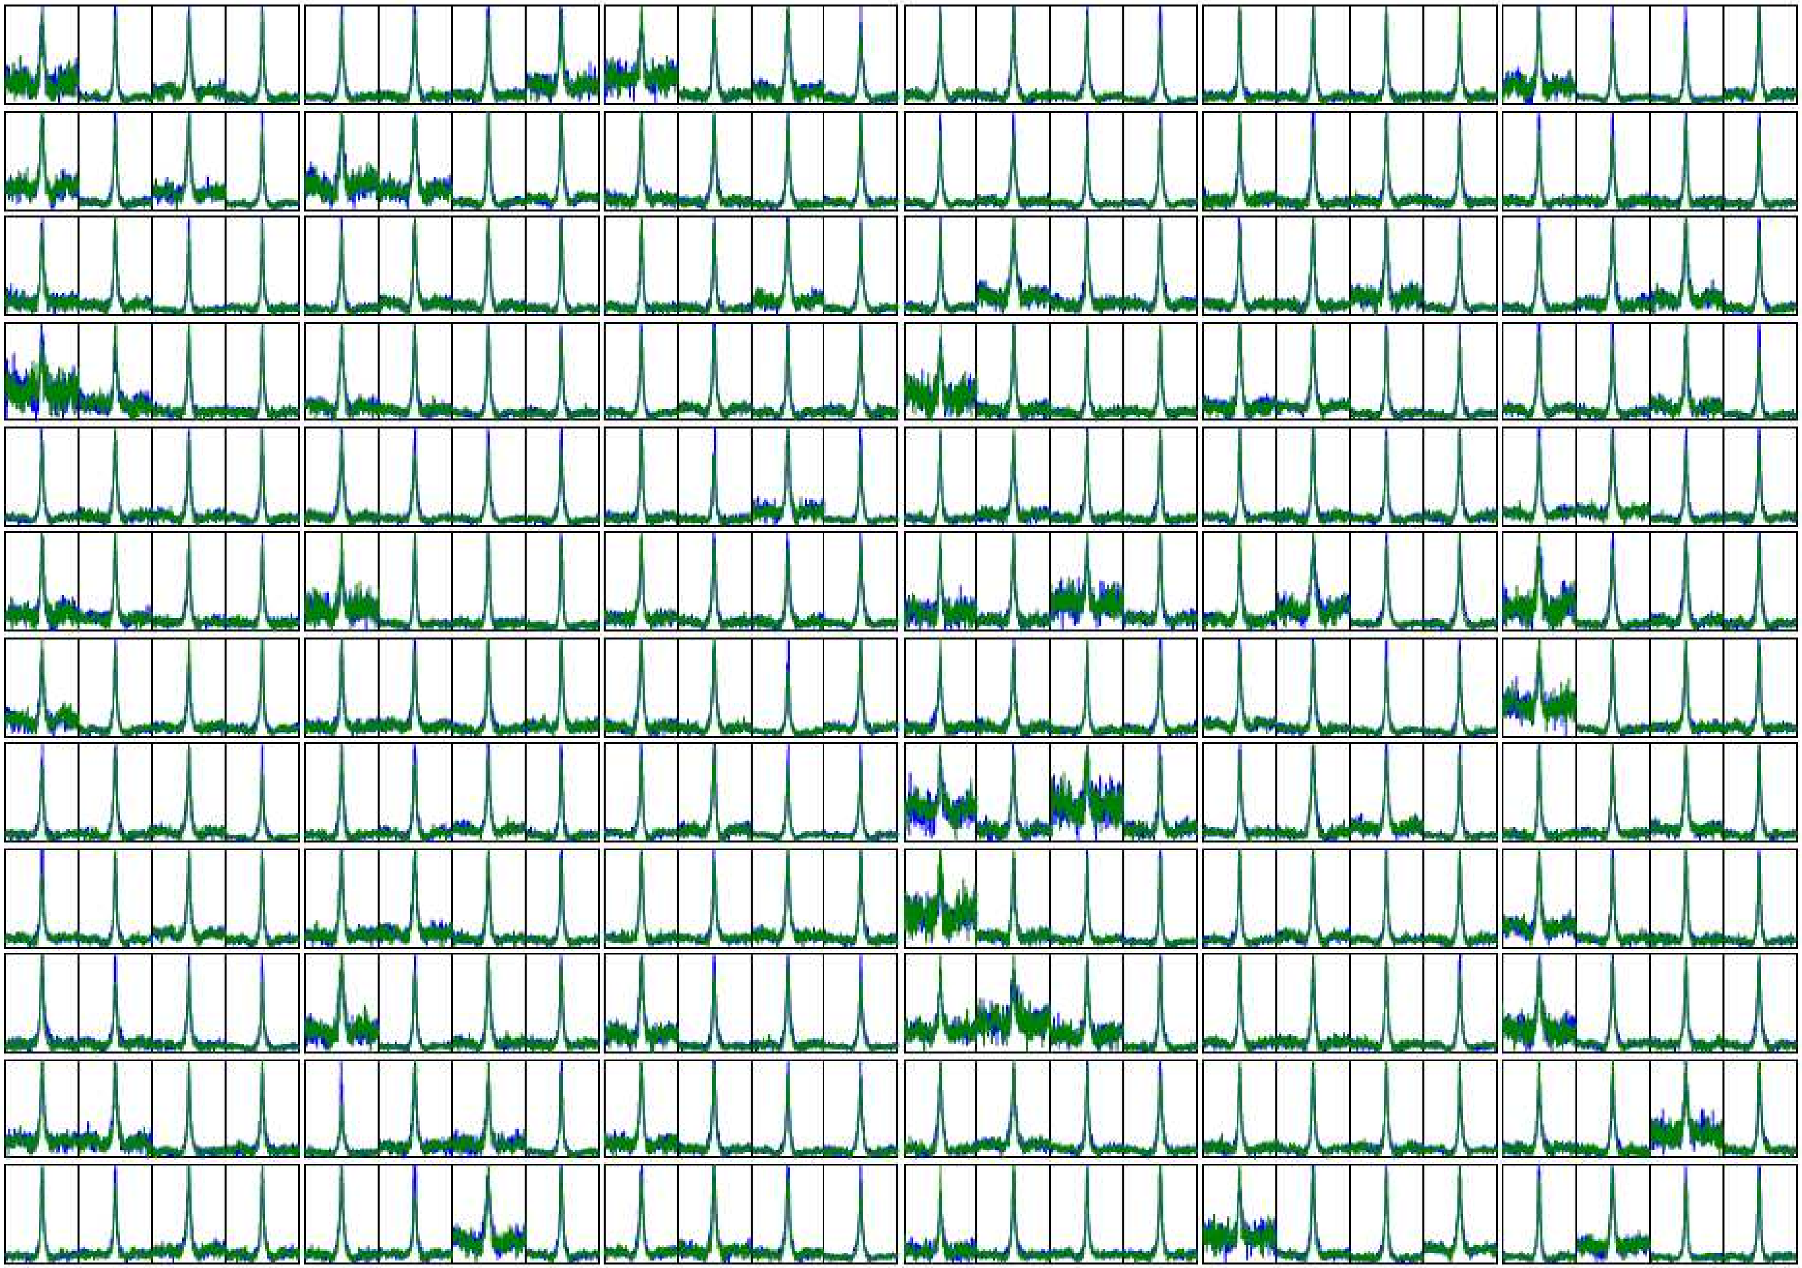

— single and triple-cell activities captured well, too (Vidne et al., 2009)

Decoding the stimulus and hidden input

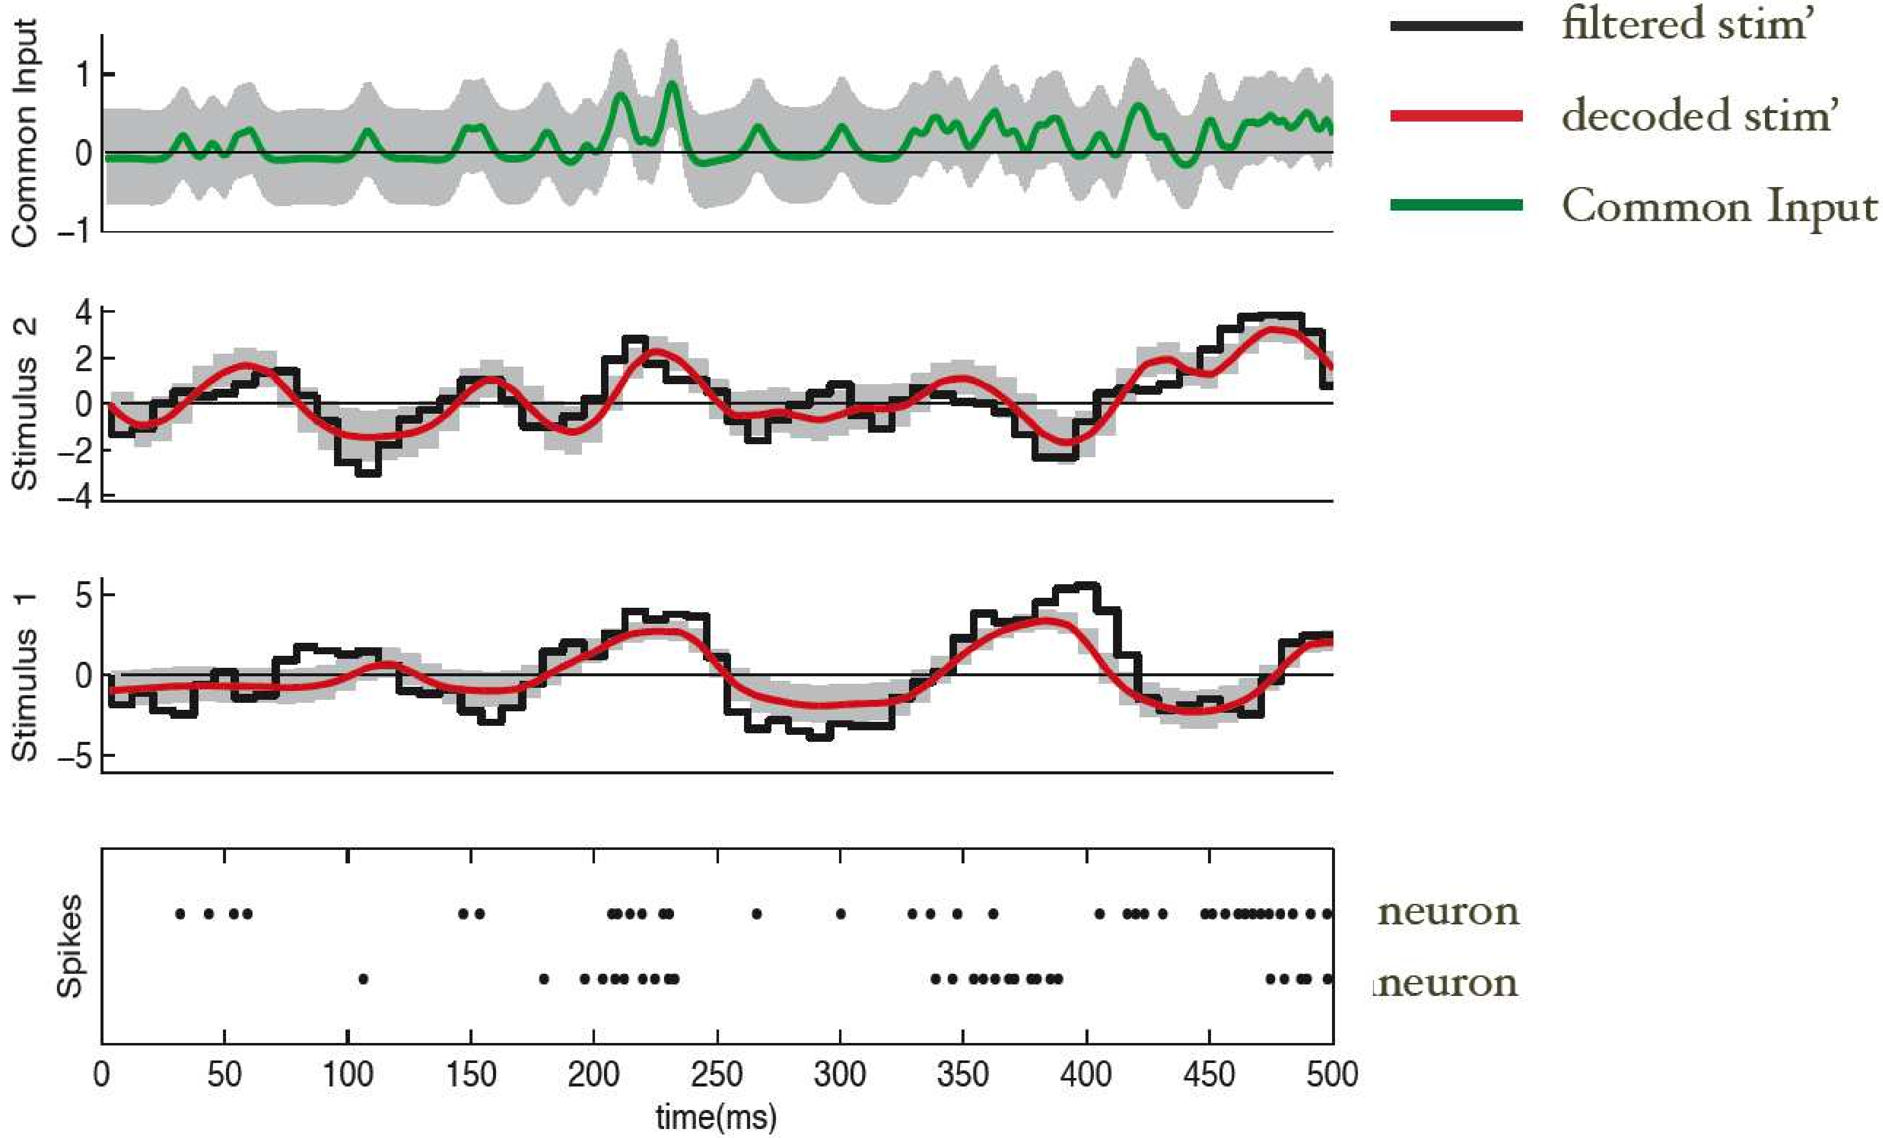

$$\arg\max_{\vec{x}} p(\vec{x}|y,\theta) = \arg\max_{\vec{x}} \int p(\vec{x},Q|y,\theta)dQ \approx \arg\max_{\vec{x},Q} p(\vec{x},Q|y,\theta)$$

Models lead to similar decoding performance

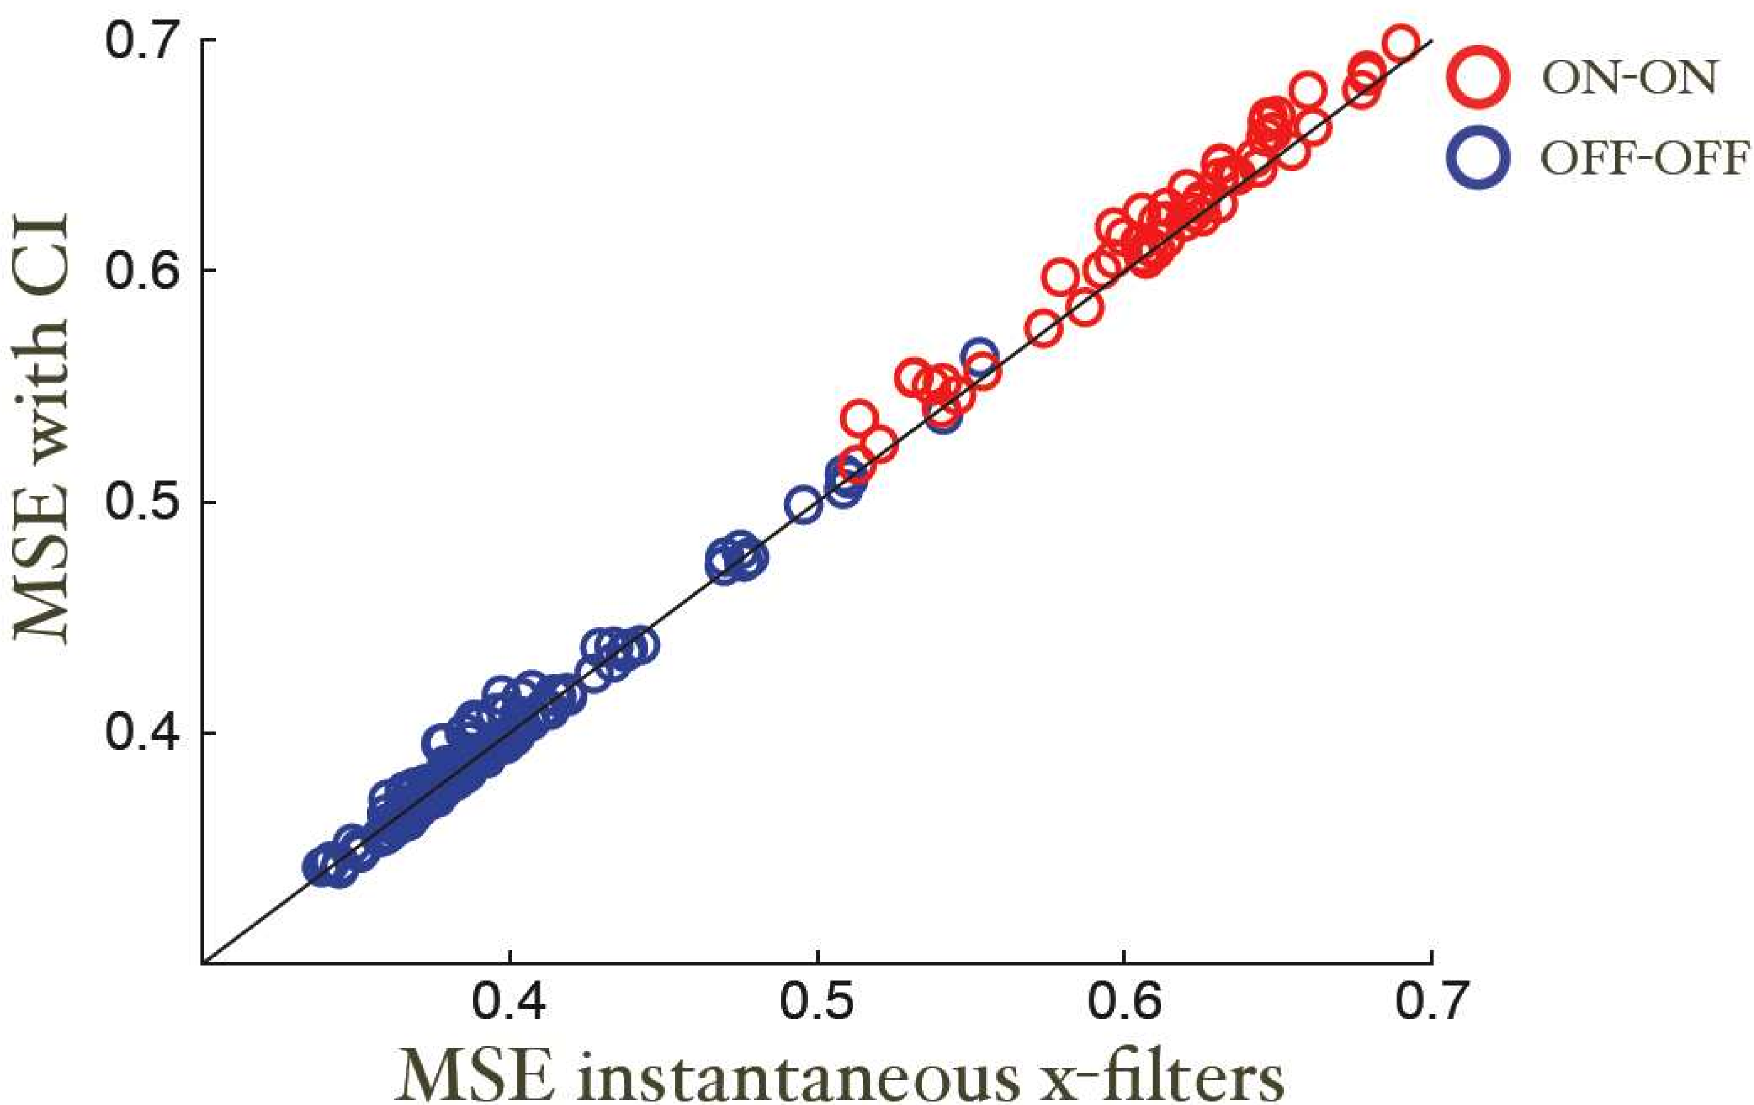

...but CI model is more robust to spike jitter and deletions (Vidne et al., 2009).

Next steps: inferring cones

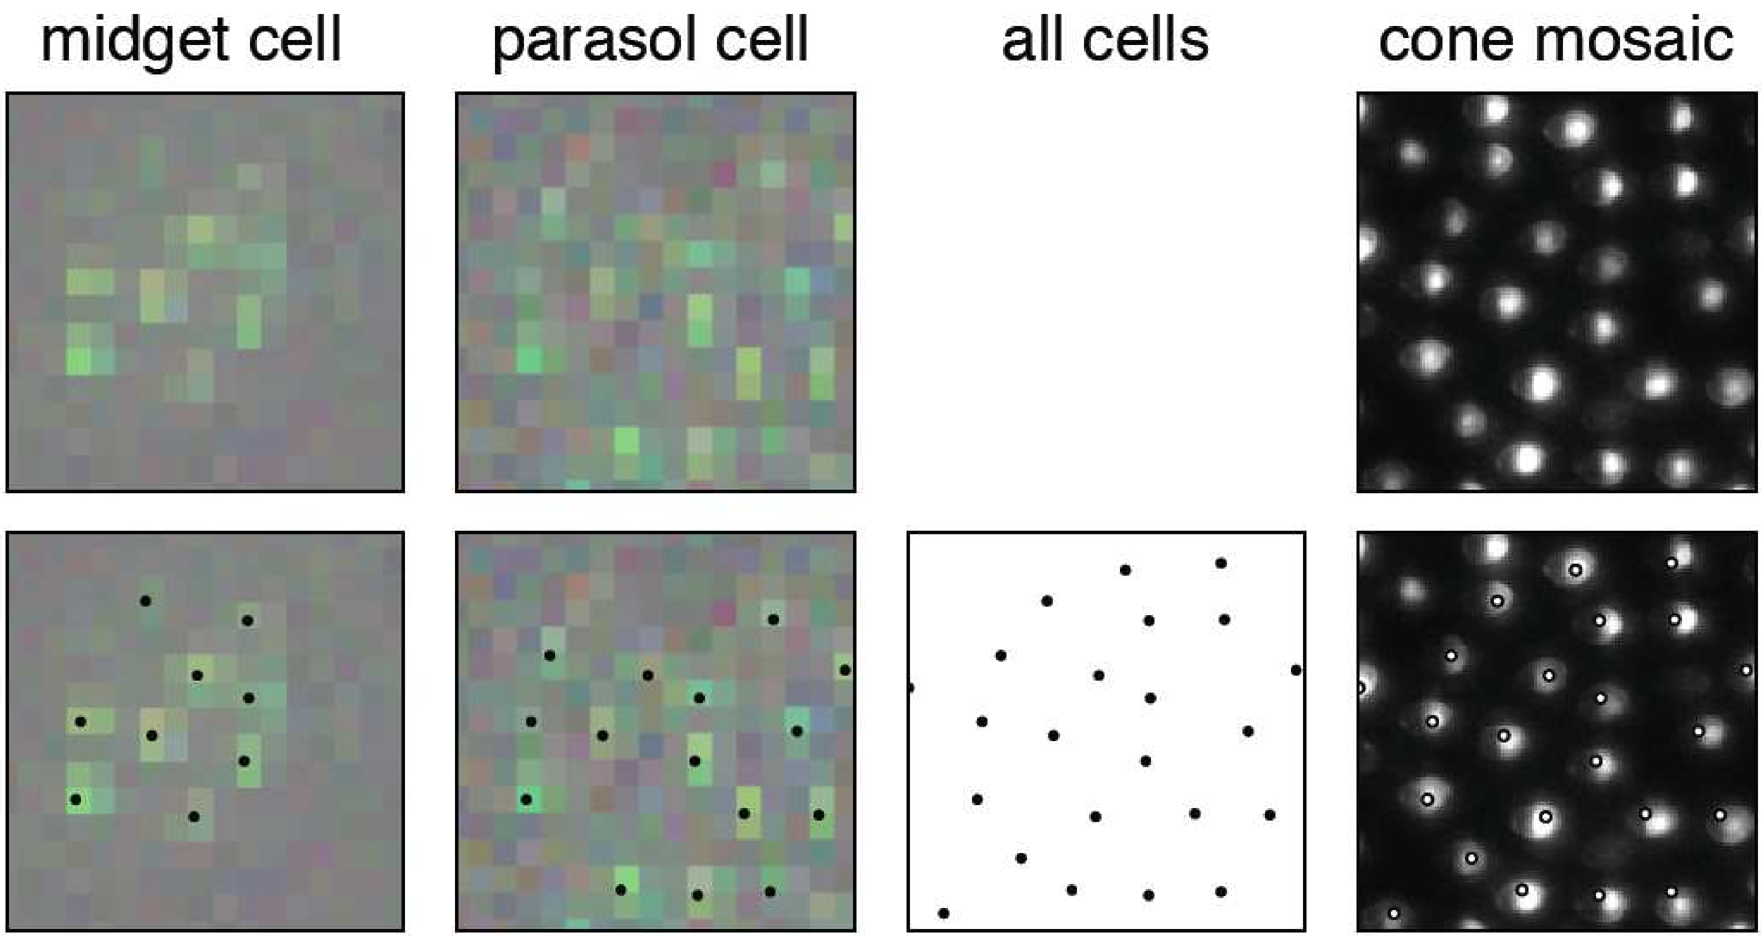

— cone locations and color identity can be inferred accurately with high spatial-resolution stimuli via maximum a posteriori estimates (Field et al., 2010).

Next steps: inferring circuitry?

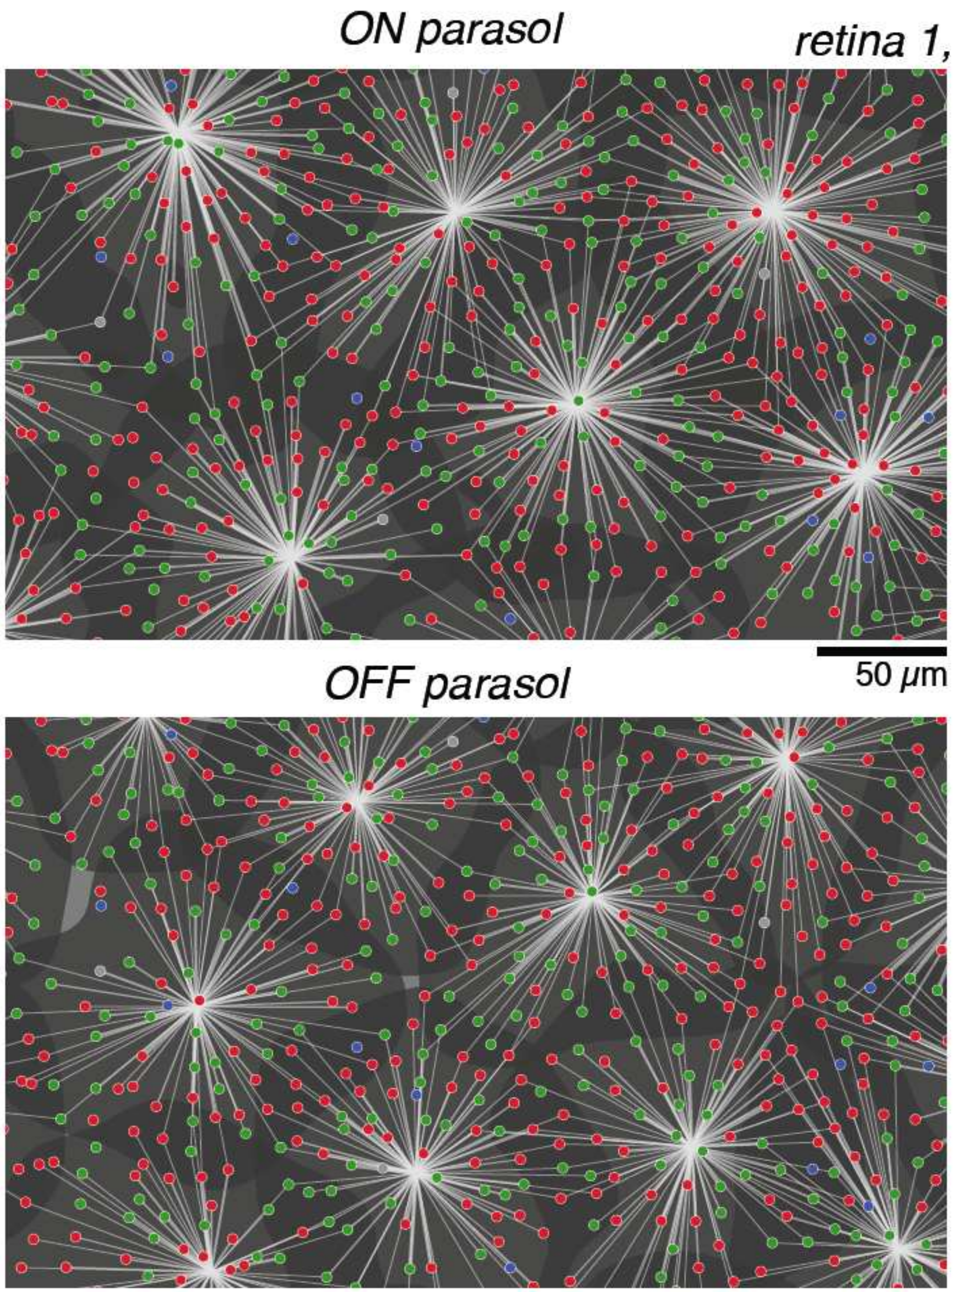

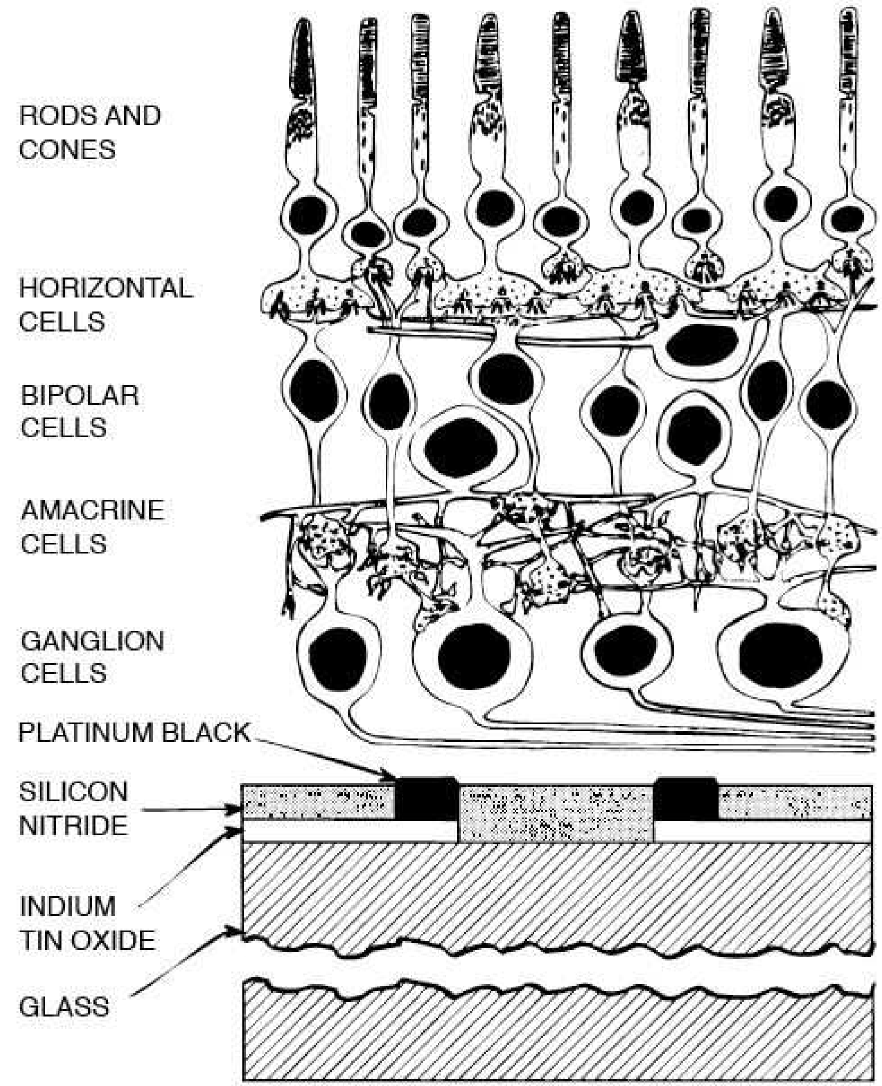

References

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

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

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.

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

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.

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). 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 input. *COSYNE*.