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

October 30, 2009

— with J. Pillow (UT Austin), E. Simoncelli (NYU), E.J. Chichilnisky, J. Gauthier, J. Shlens (Salk), E. Lalor (TC Dublin), S. Koyama (CMU), Y. Ahmadian, J. Kulkarni, H. Liu, T. Machado, D. Pfau, X. Pitkow, M. Vidne (Columbia).

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

Retinal ganglion neuronal data

Preparation: dissociated macaque retina

— extracellularly-recorded responses of populations of RGCs

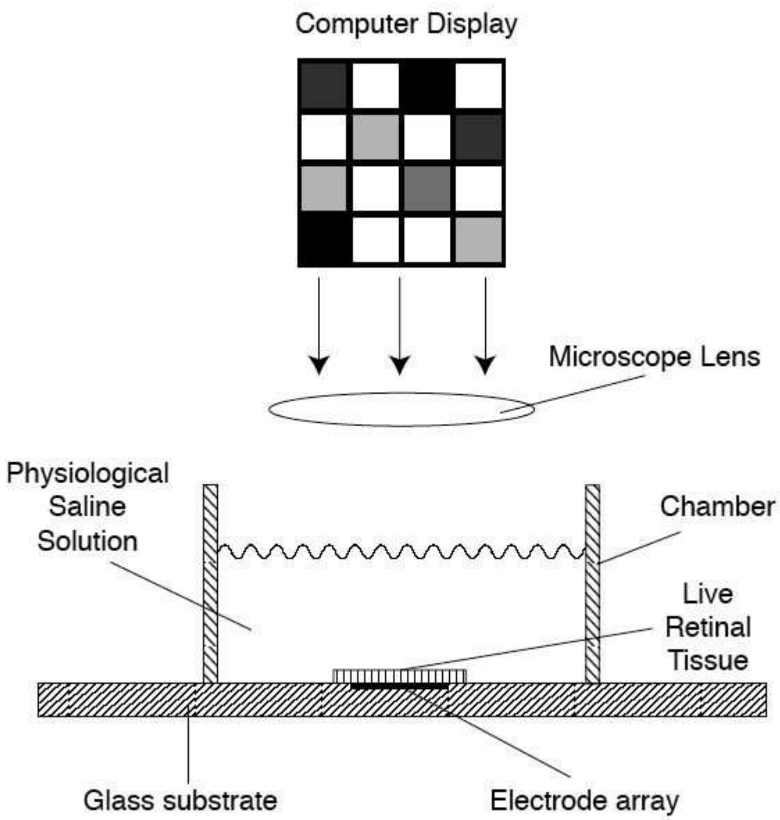

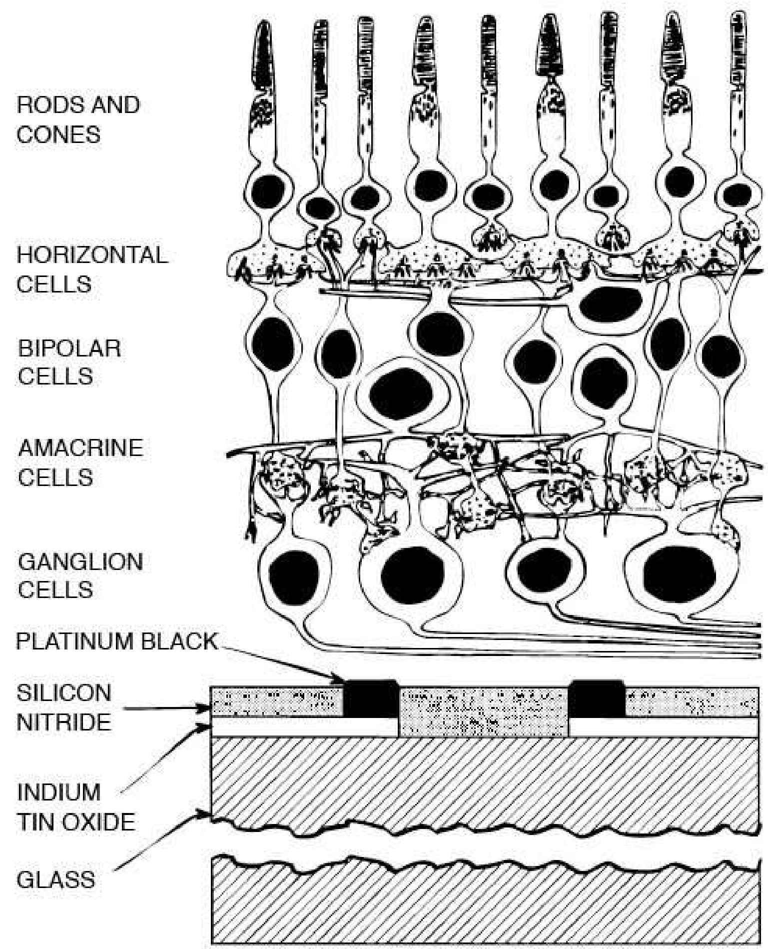

Stimulus: random spatiotemporal visual stimuli (Pillow et al., 2008)

Receptive fields tile visual space

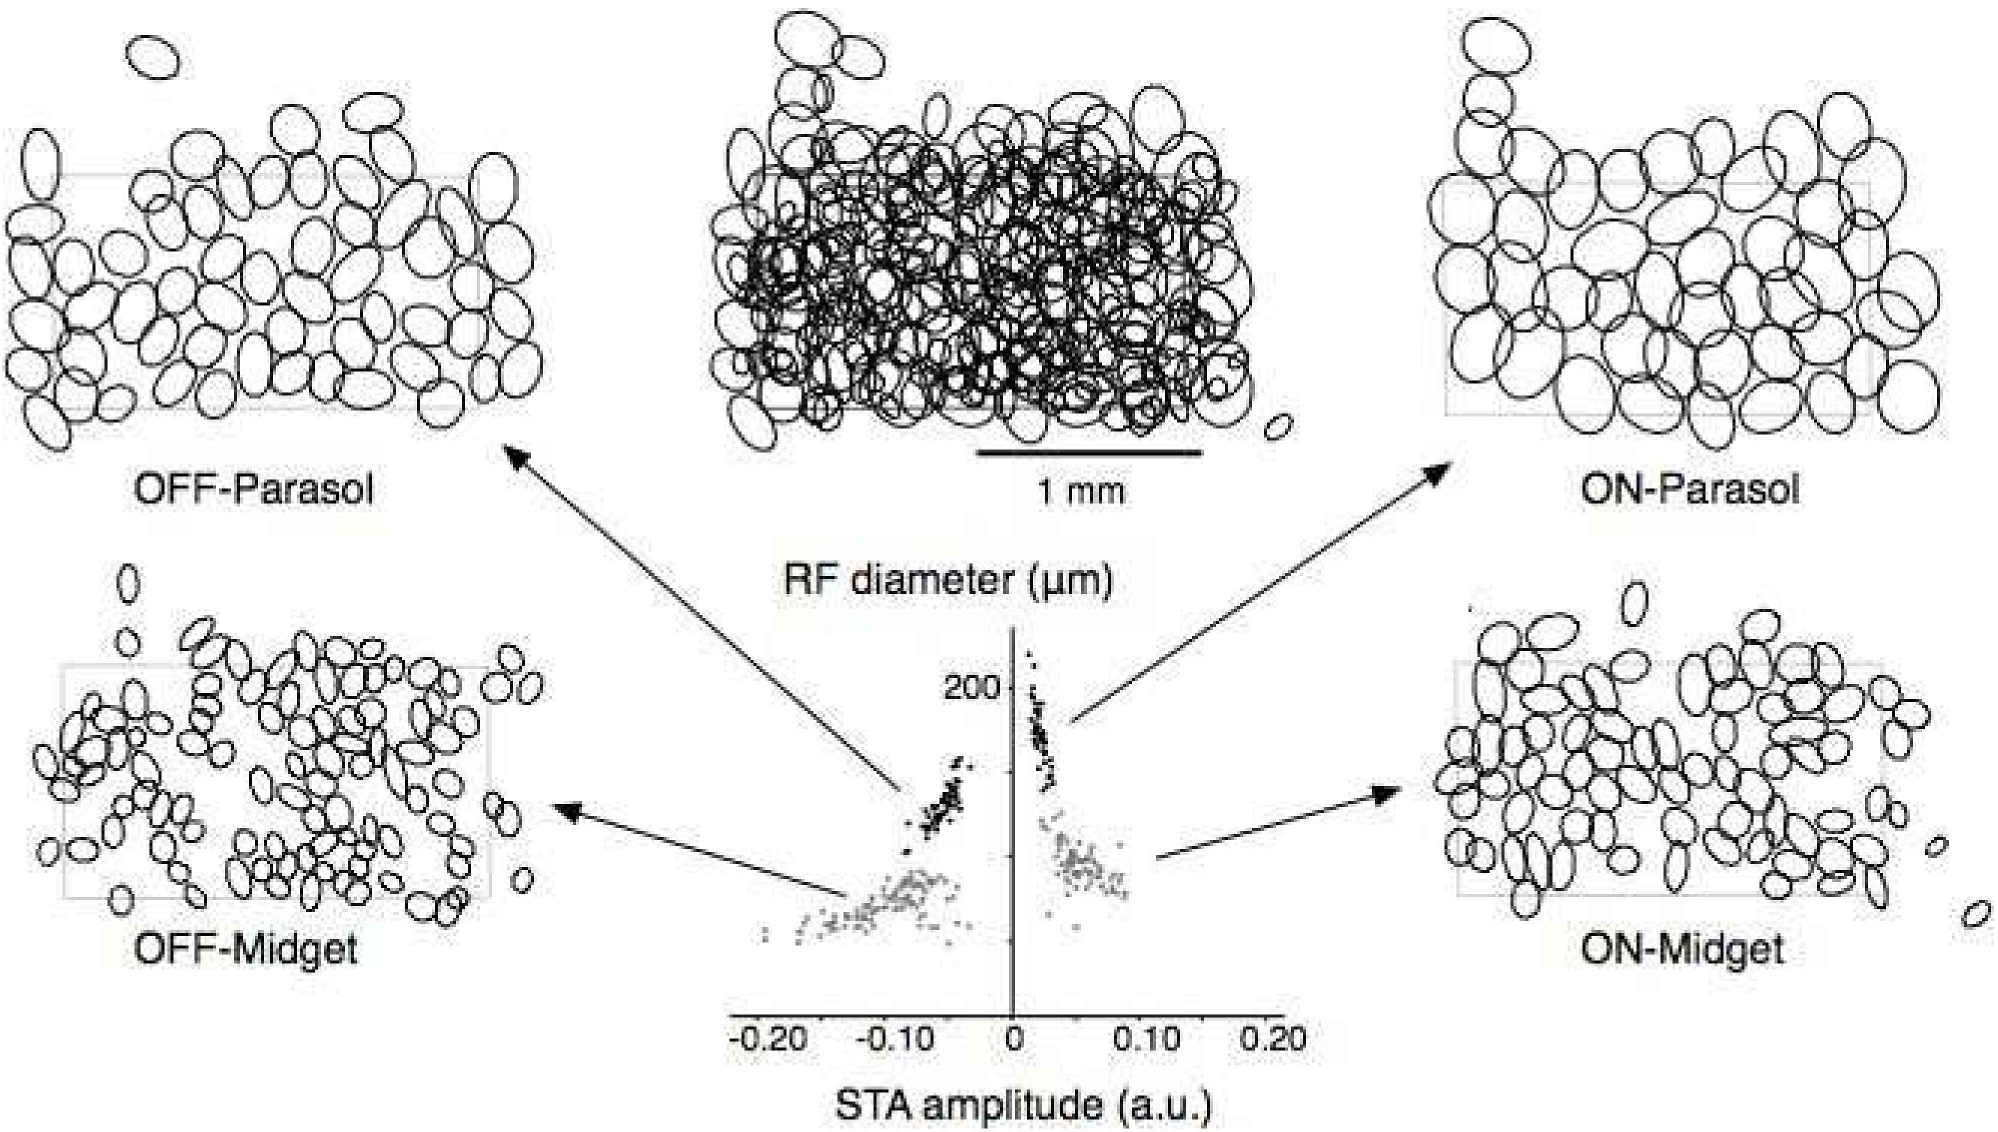

Multineuronal point-process model

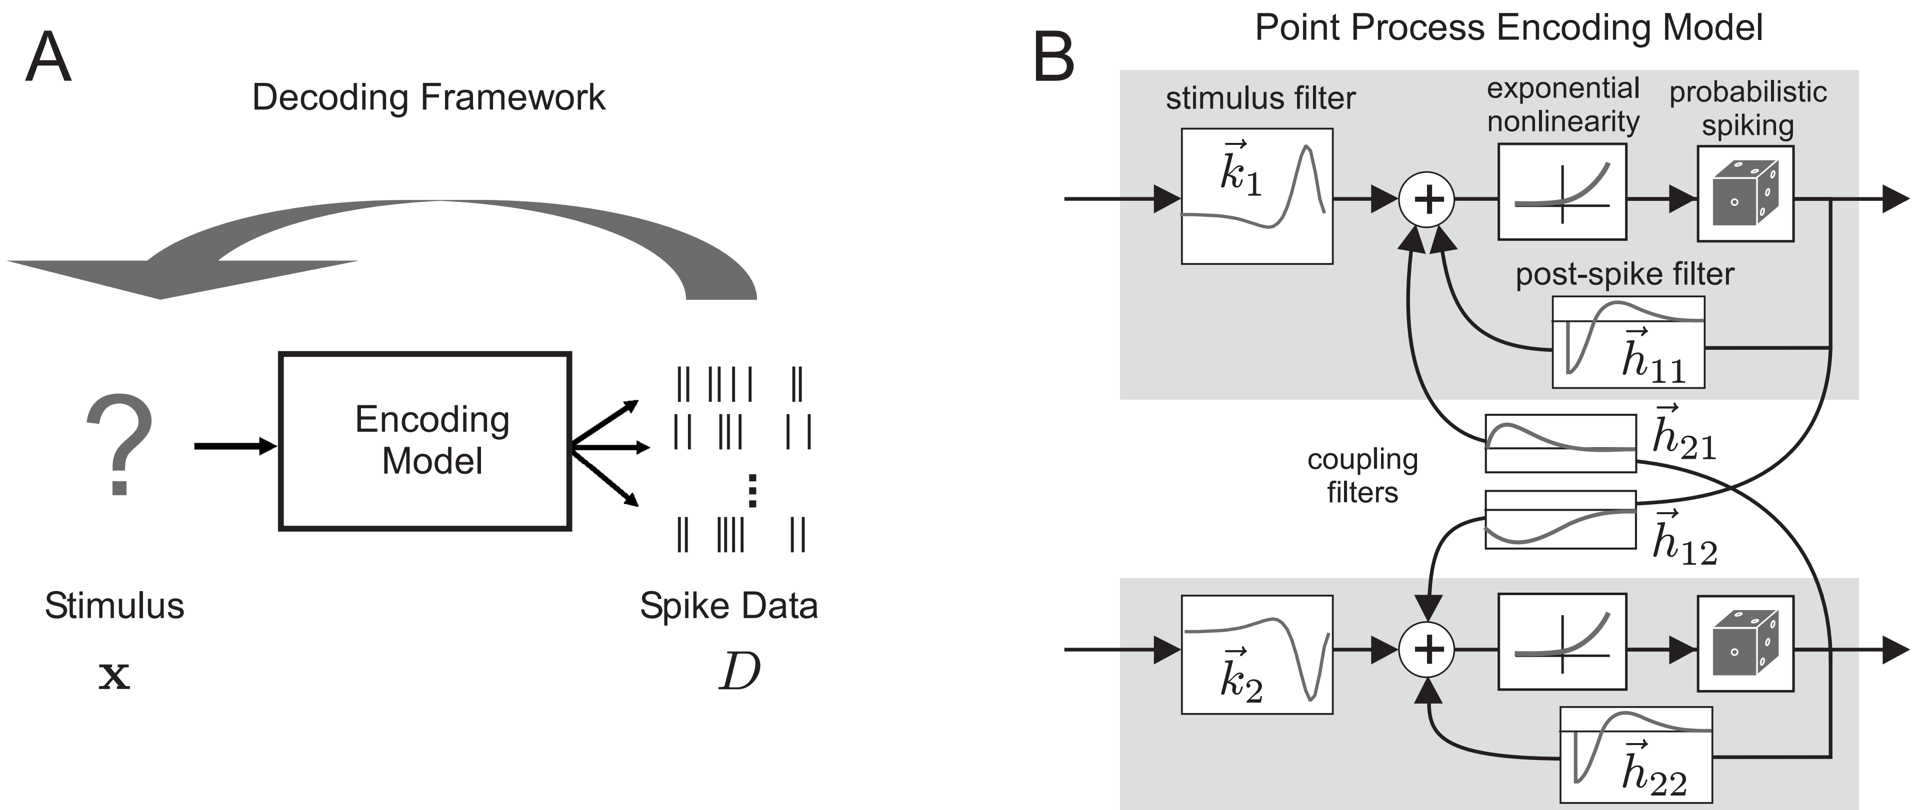

$$\lambda_i(t) = f\left(b_i + \vec{k}_i \cdot \vec{x}(t) + \sum_{i',j} h_{i',j} n_{i'}(t-j)\right),$$

— GLM; fit by L_1-penalized maximum likelihood (concave optimization)
(Paninski, 2004; Truccolo et al., 2005; Pillow et al., 2008)

Model captures spatiotemporal cross-corrs

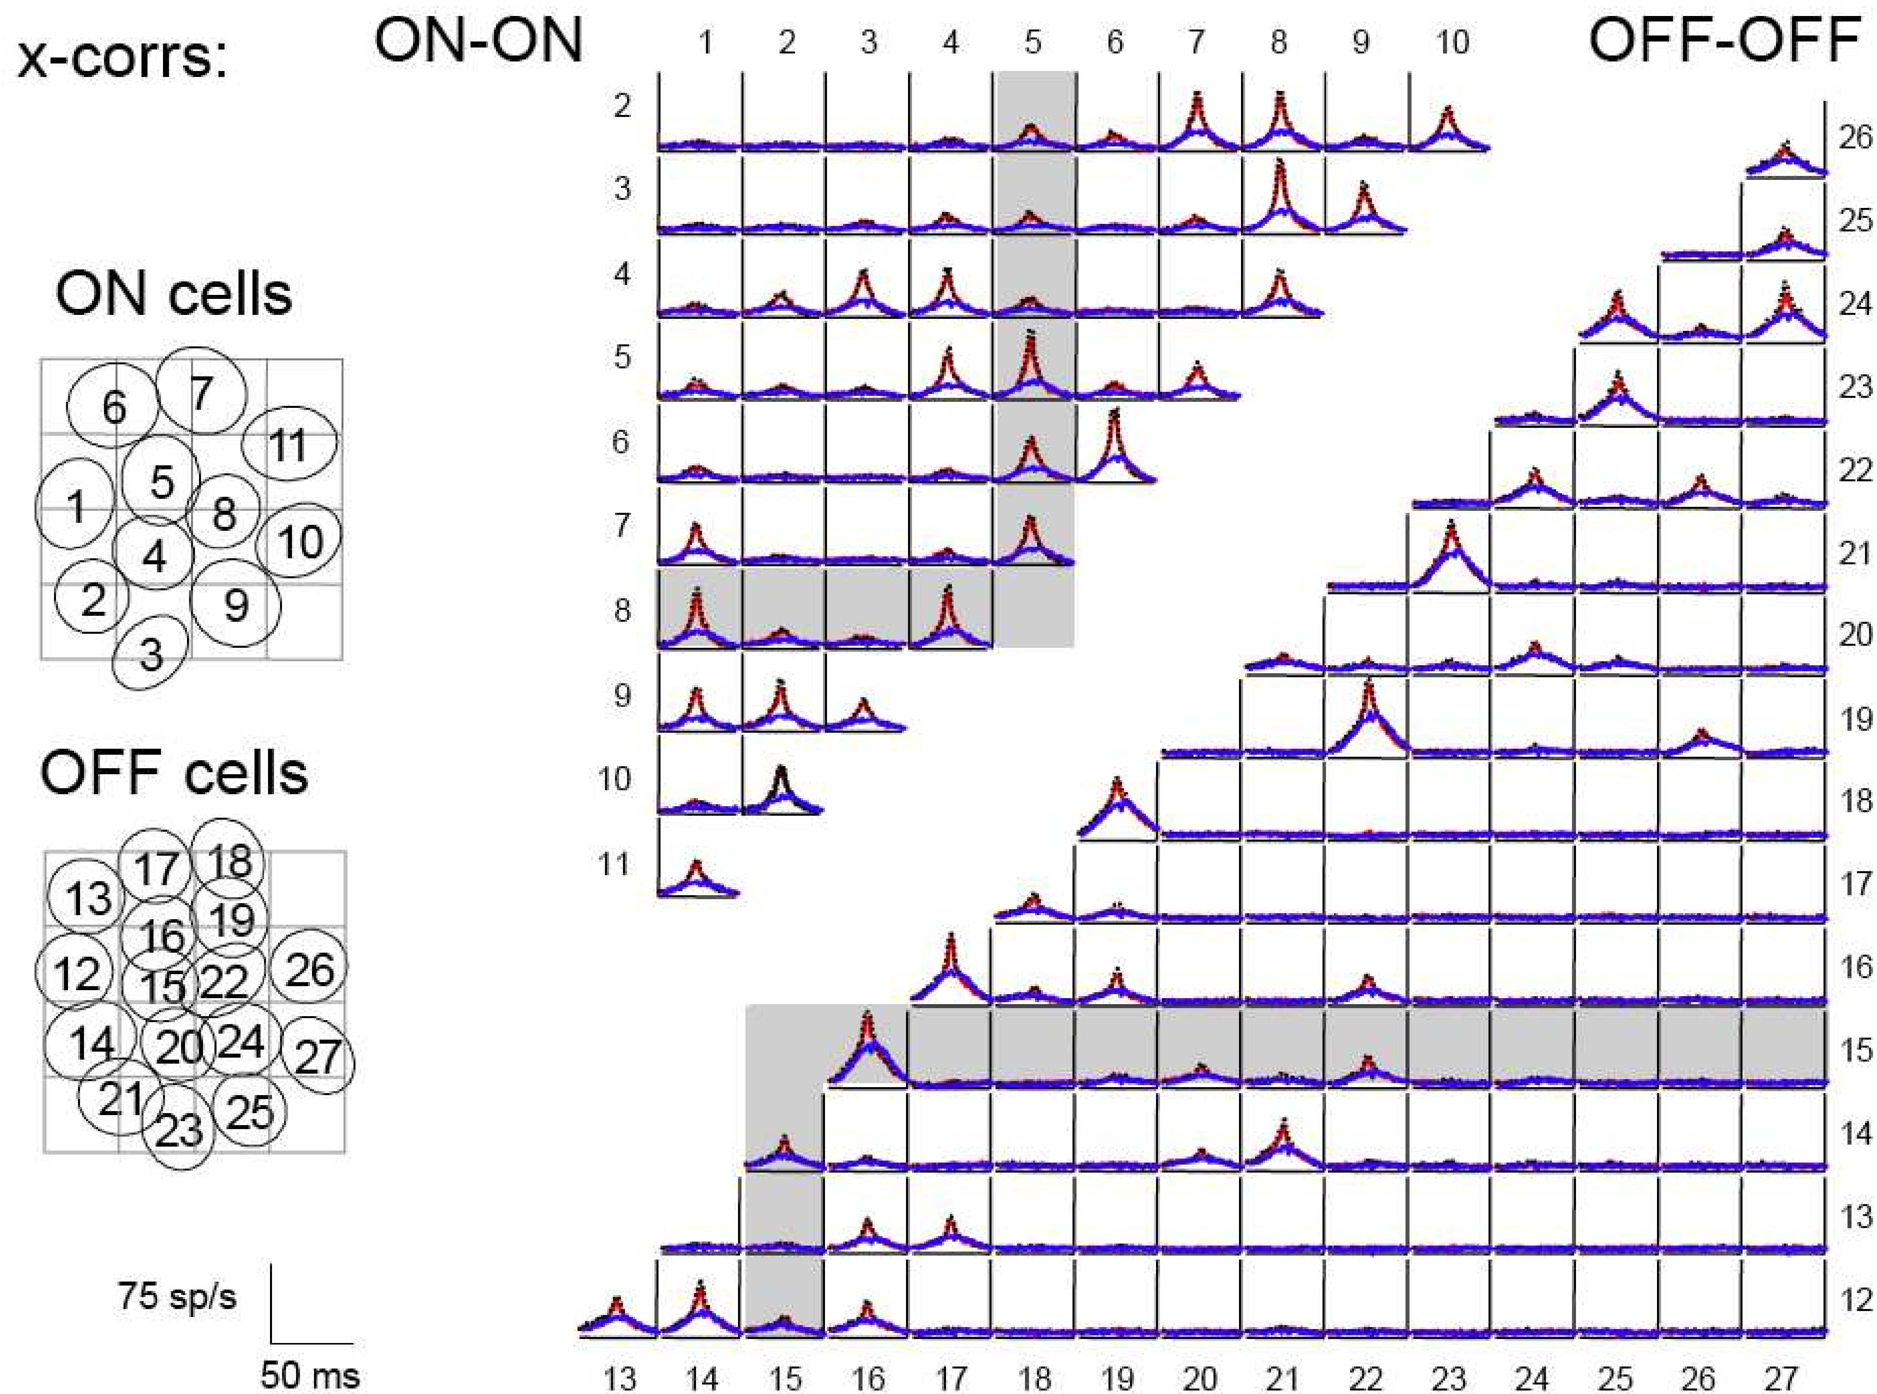

Optimal Bayesian decoding

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

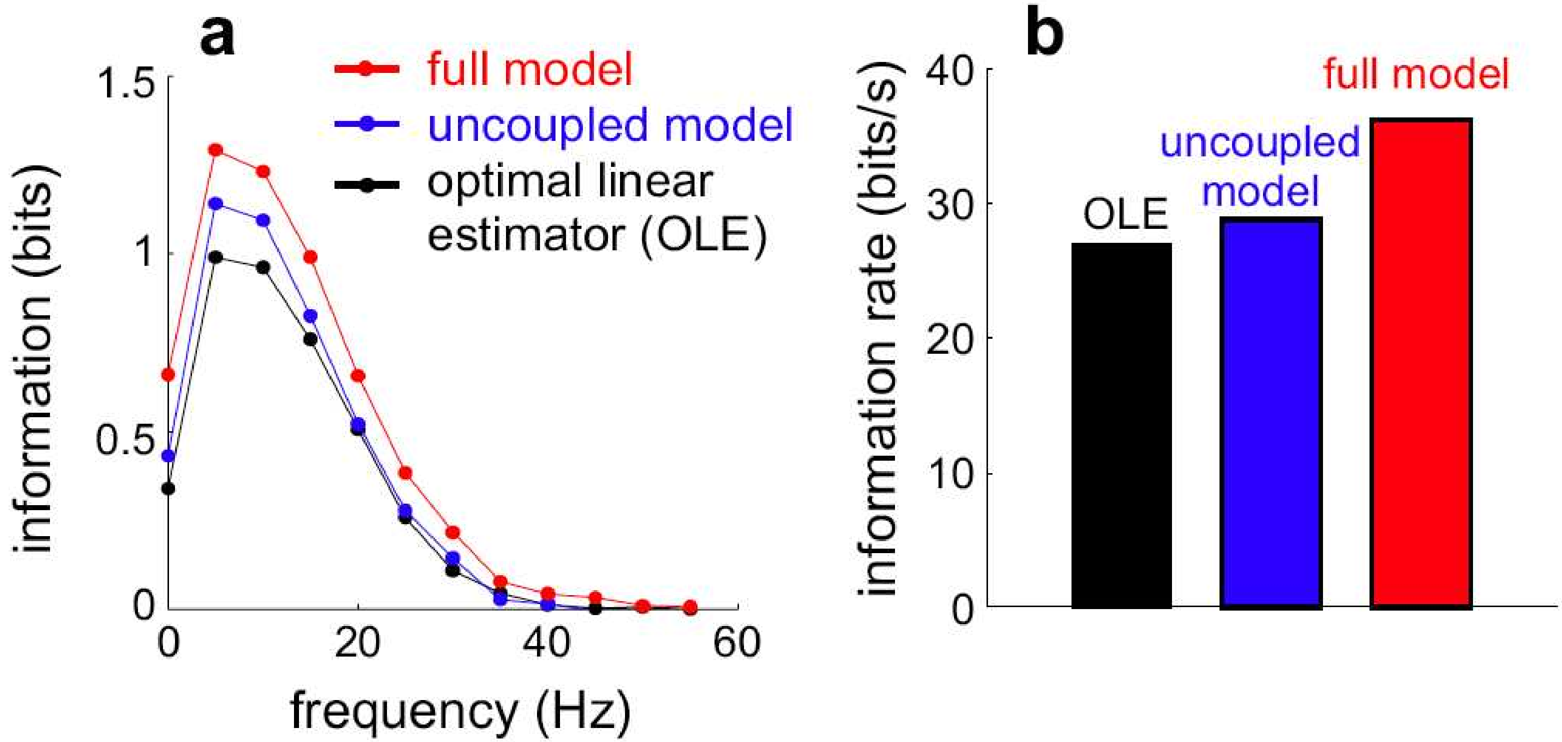

— Computational points:

- $\log P(spikes|\vec{x})$ 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., 2009).

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

Application: how important is timing?

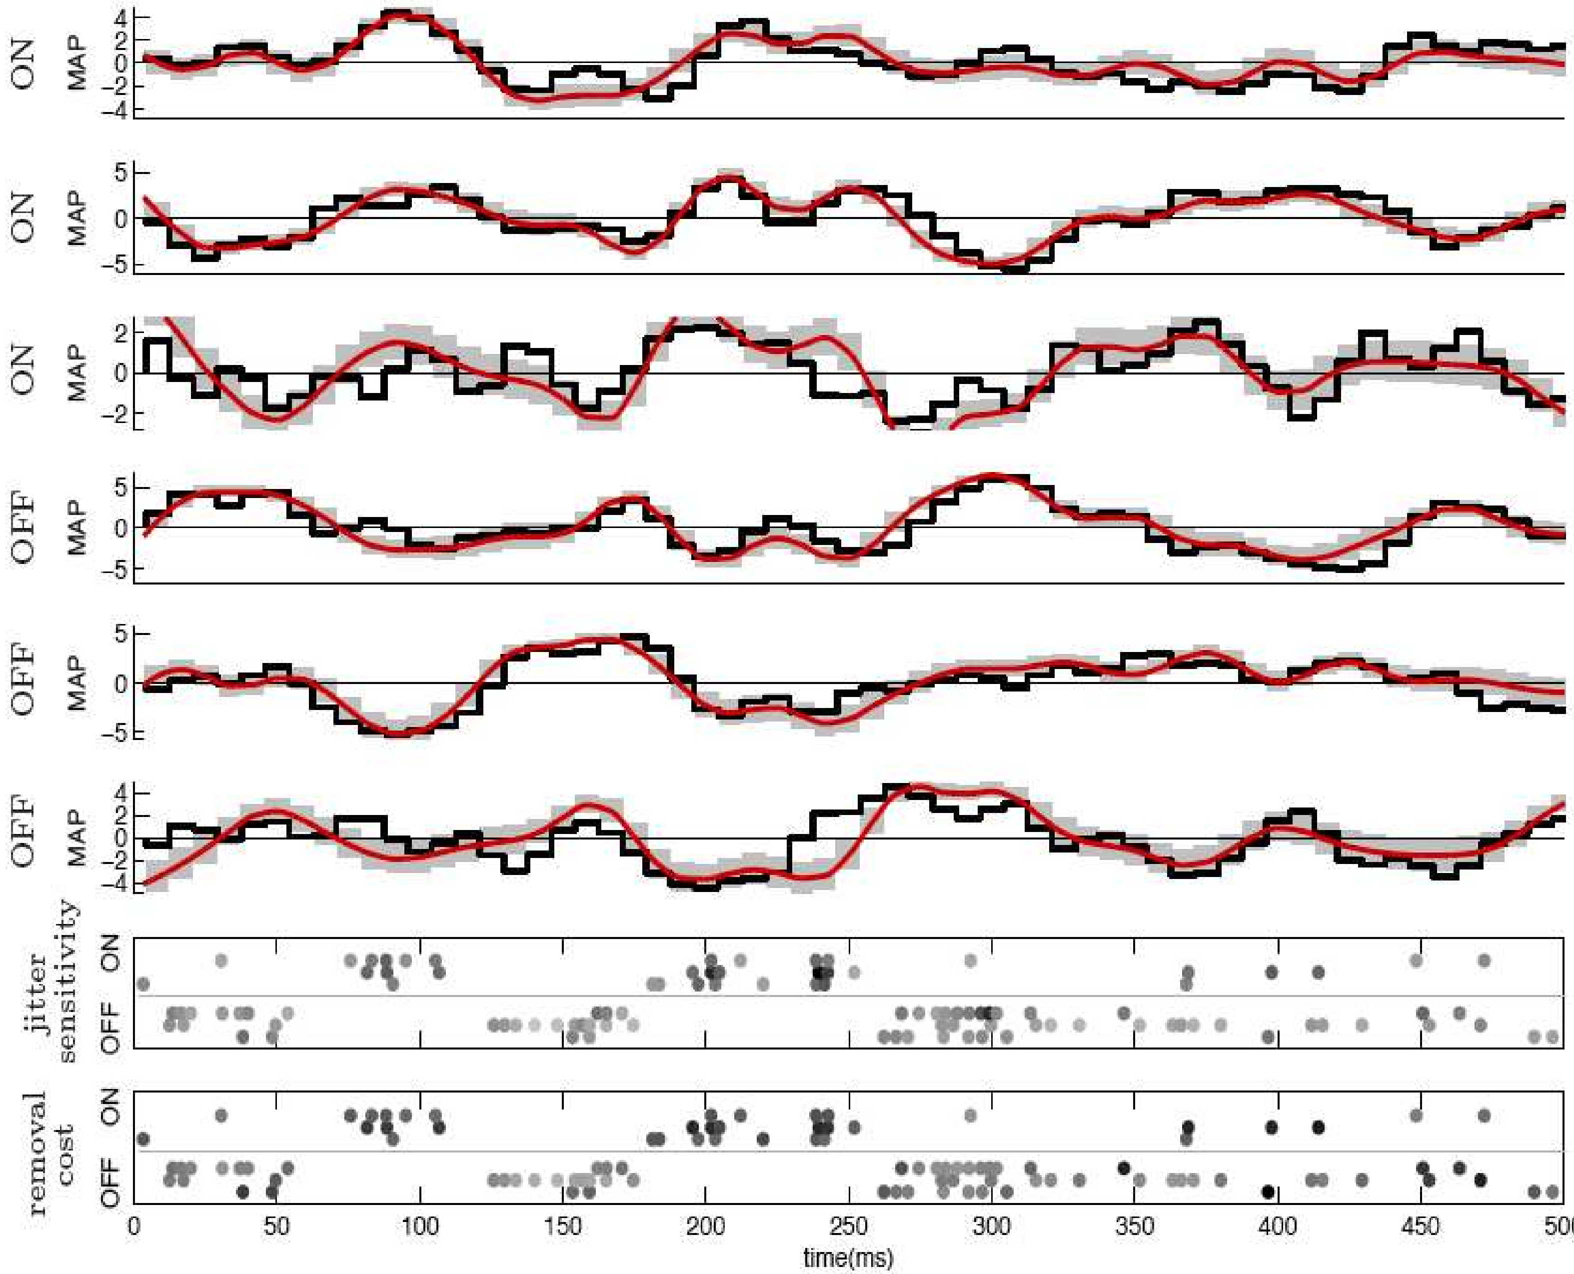

— Fast decoding methods let us look more closely (Ahmadian et al., 2009)

Spike sensitivity is strongly context-dependent

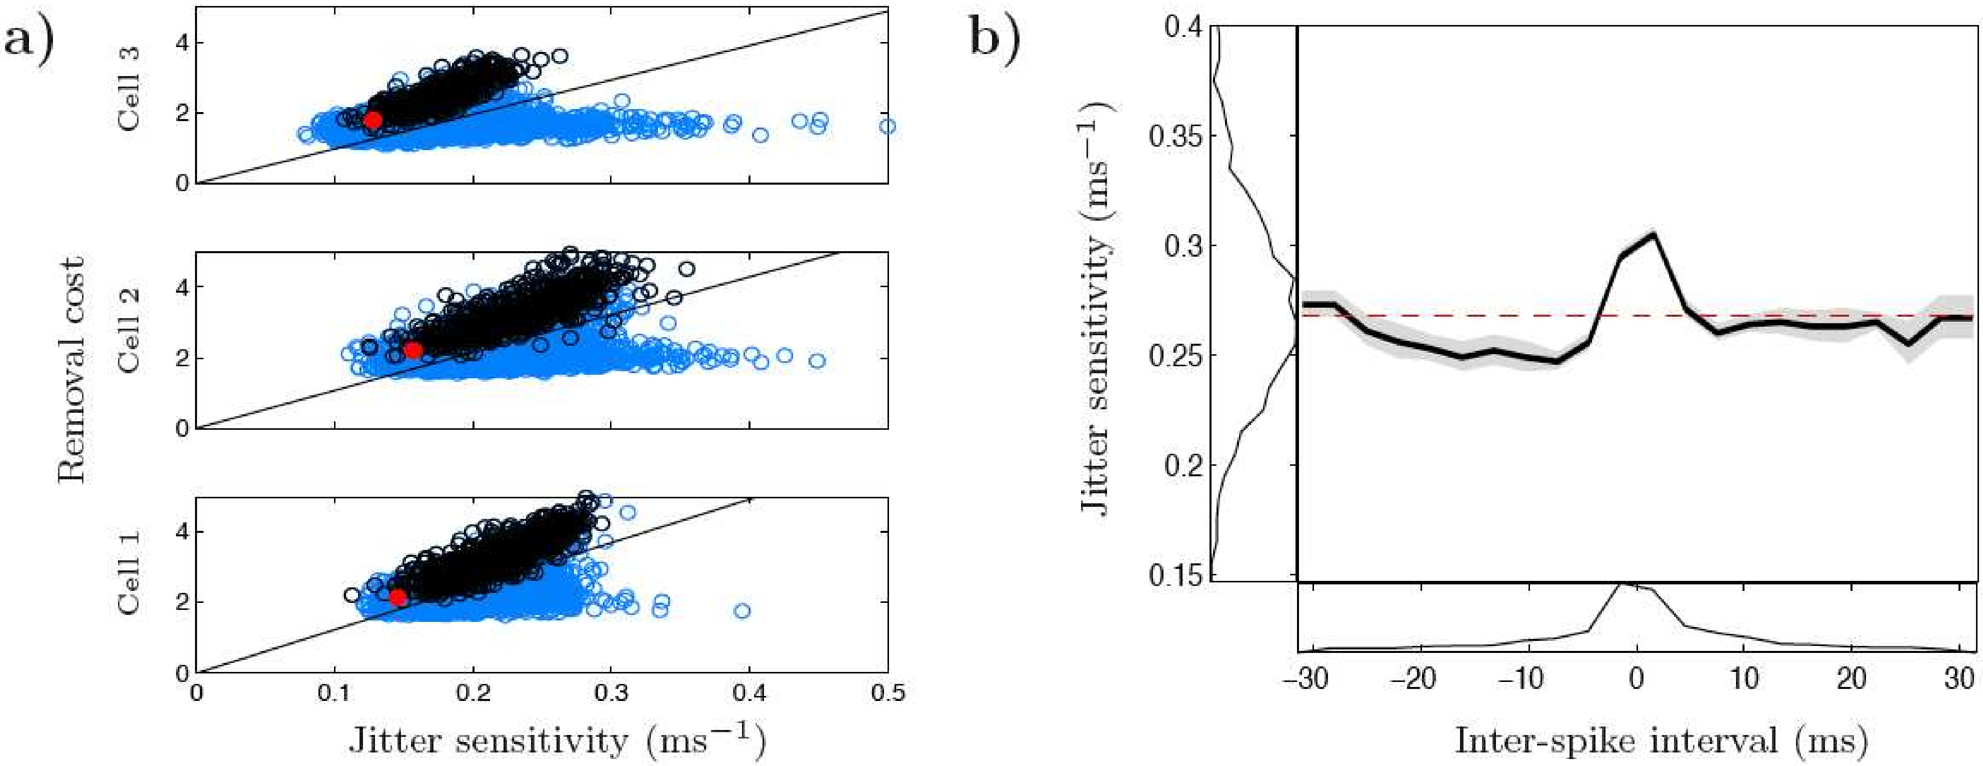

— Reflects nonlinearity of decoder $\hat{x}(r)$: linear decoder is context-independent

— Cost of spike addition/deletion $\approx$ cost of jittering by 10 ms (Victor, 2000): natural time scale of spike train.

Application: image stabilization

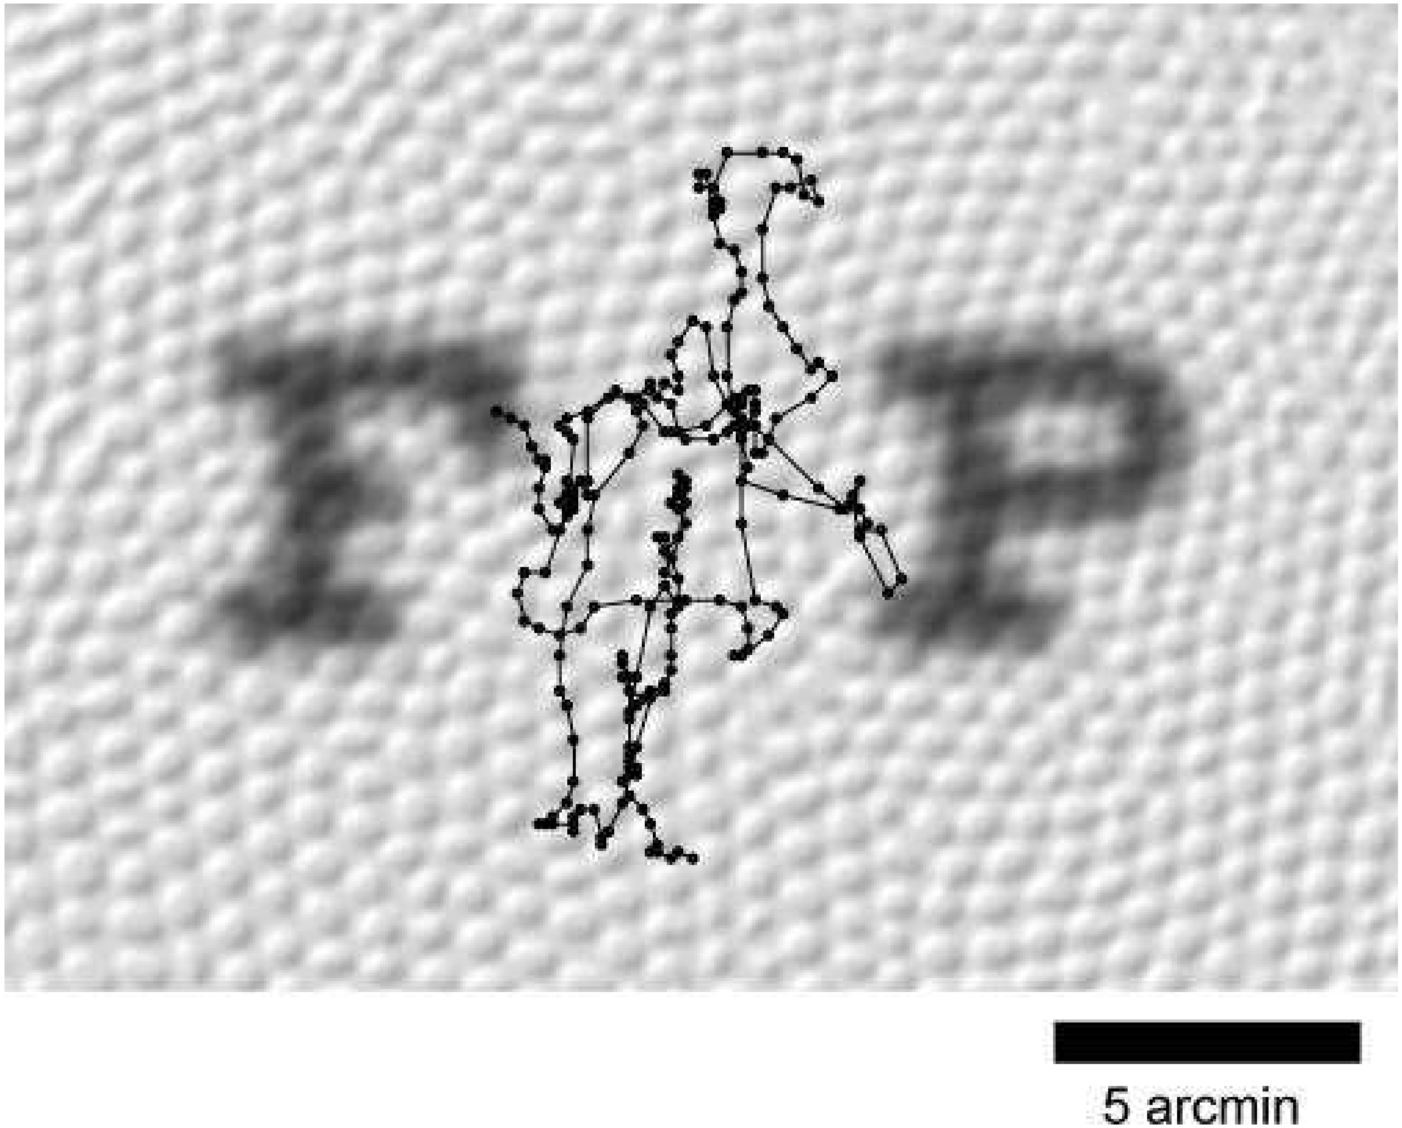

From (Pitkow et al., 2007): neighboring letters on the 20/20 line of the Snellen eye chart. Trace shows 500 ms of eye movement.

Bayesian methods for image stabilization

Have to marginalize out random eye movements:

$$p(I|spikes) \propto p(I)p(spikes|I) = p(I)\int p(spikes|e,I)p(e)de;$$

e denotes eye path; integration by particle-filter methods.

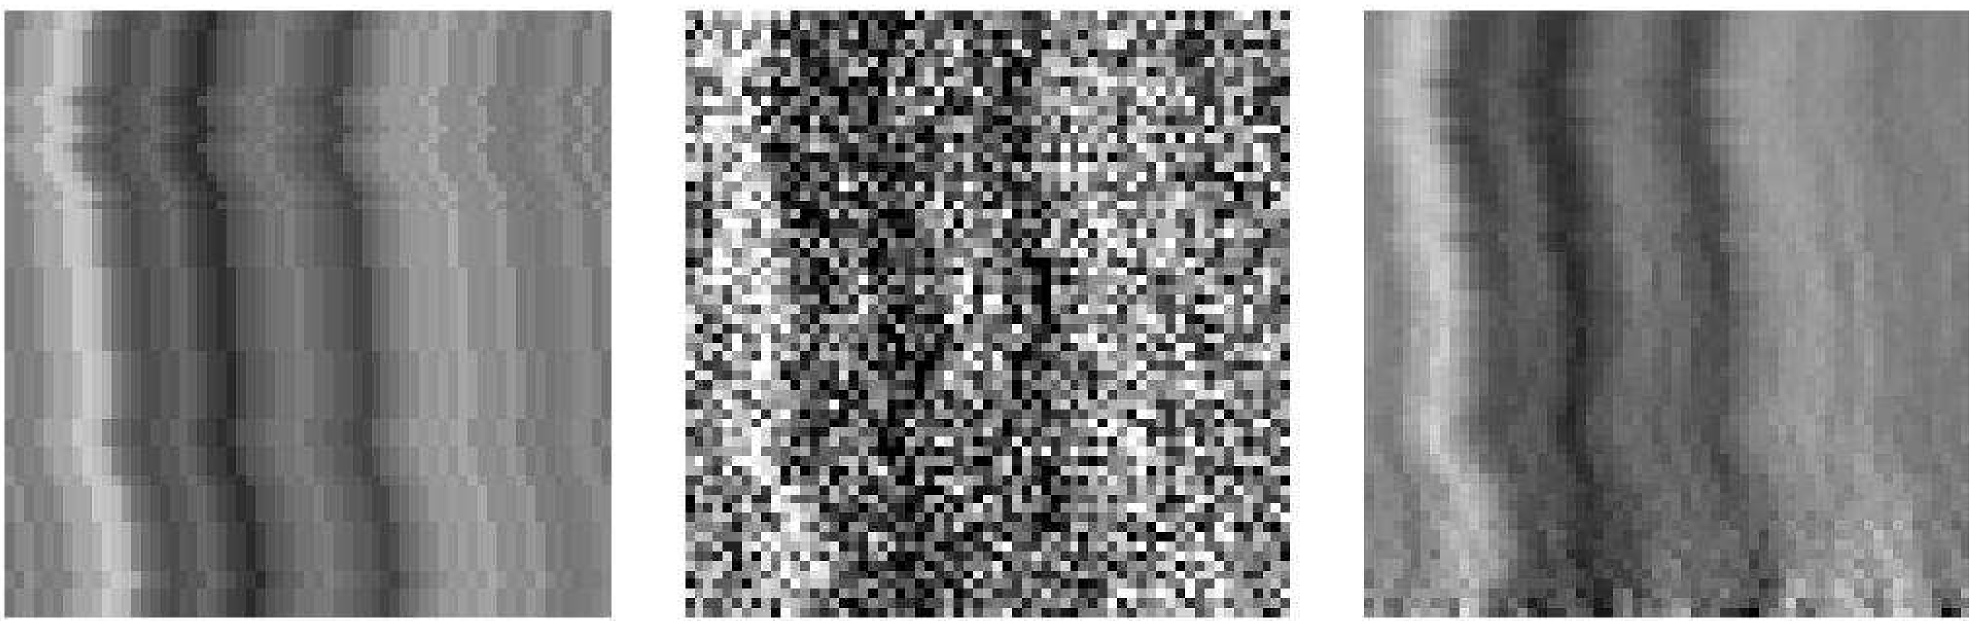

true image w/ translations; observed noisy retinal responses; estimated image.

Reconsidering the model

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

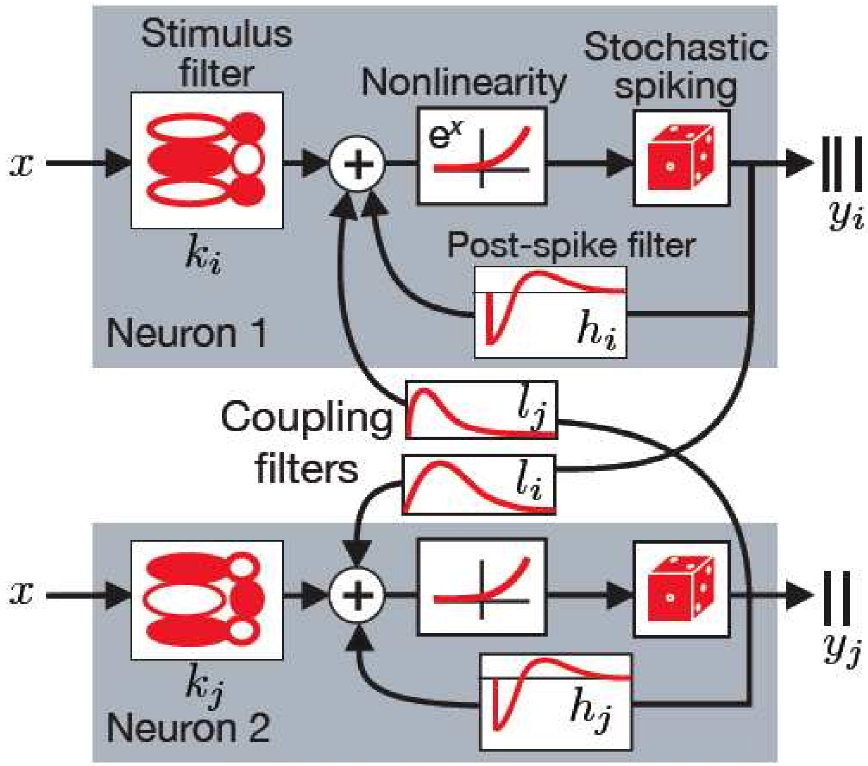

Pros:

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

Cons:

- Instantaneous coupling filters
- No explicit Common Input

Considering common input effects

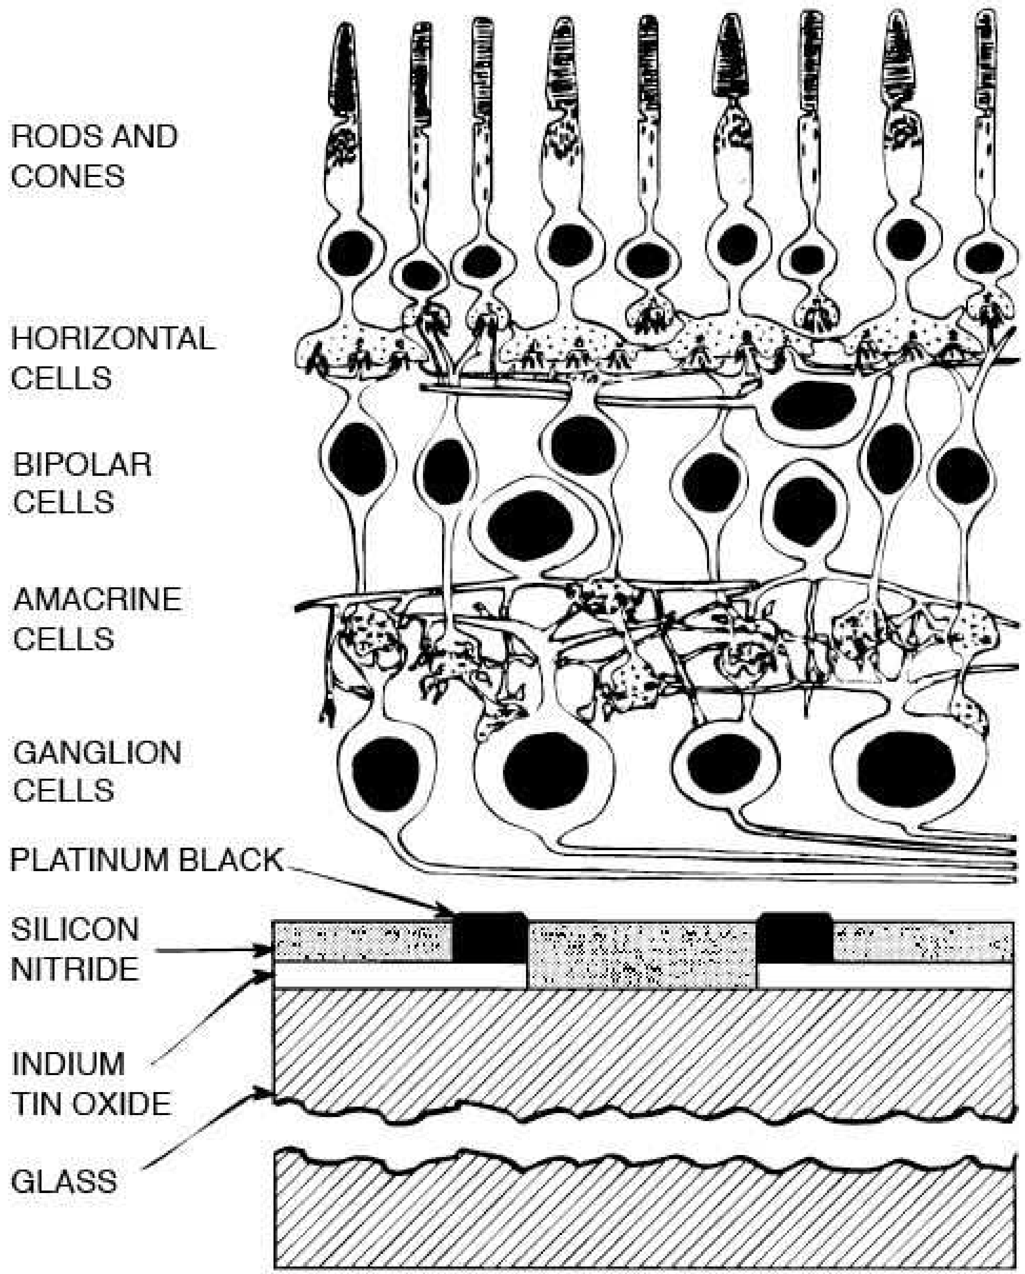

Intracellular findings:

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

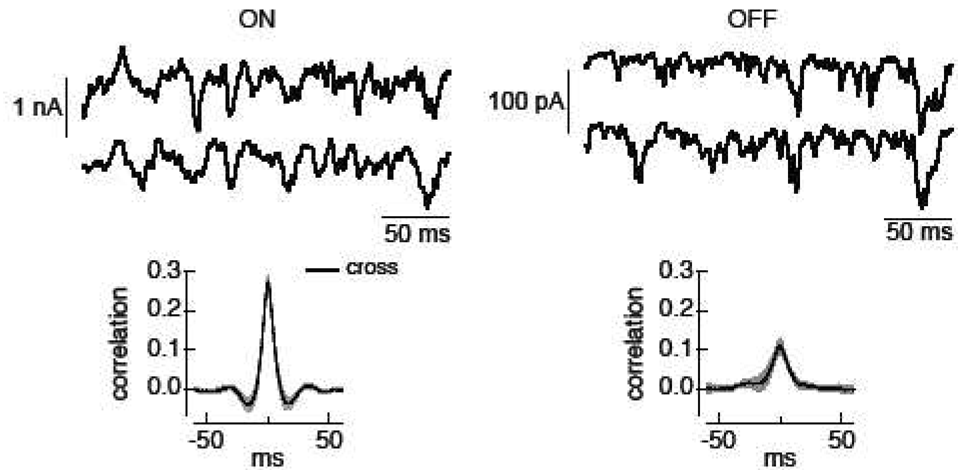

Khuc Trong & Rieke Nature Neuro 2008

- ON RGCs are weakly electrically coupled

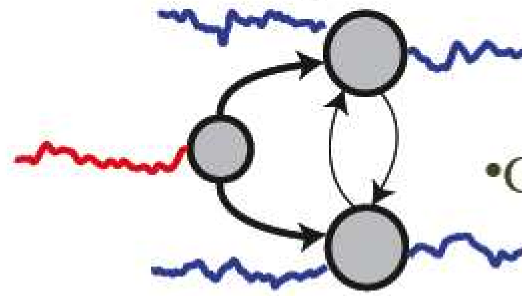

•ON RGCs

- No electrical coupling seen between OFF RGCs

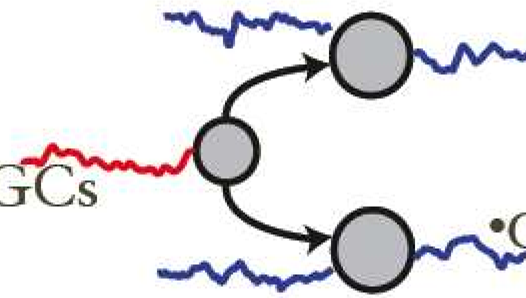

•OFF RGCs

Extension: including common input effects

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

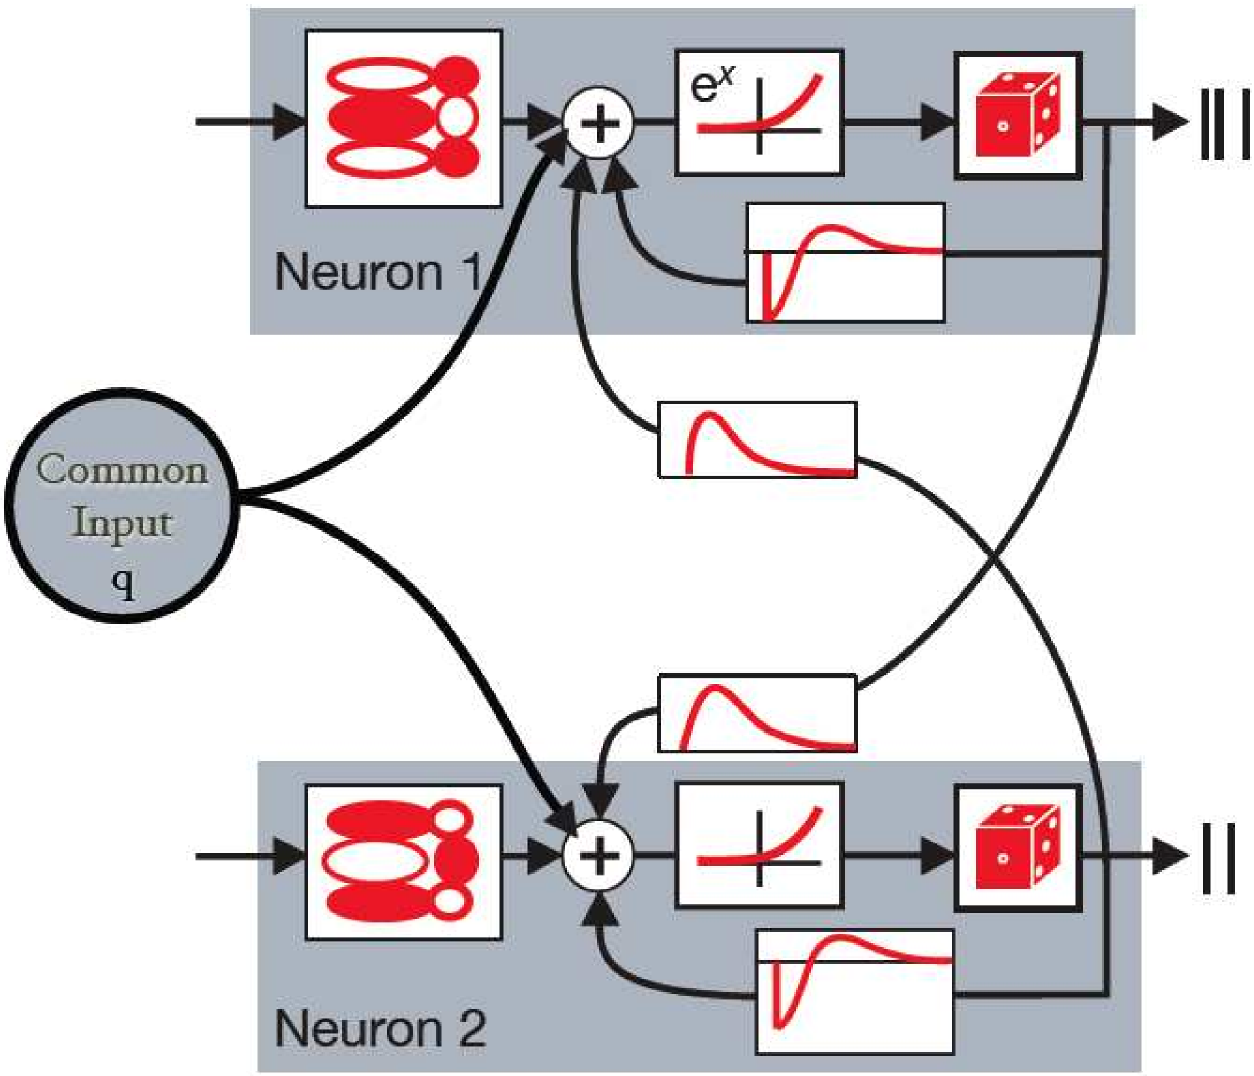

Direct state-space optimization methods

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 $\approx$ 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., 2009).

Inferred common input effects are strong

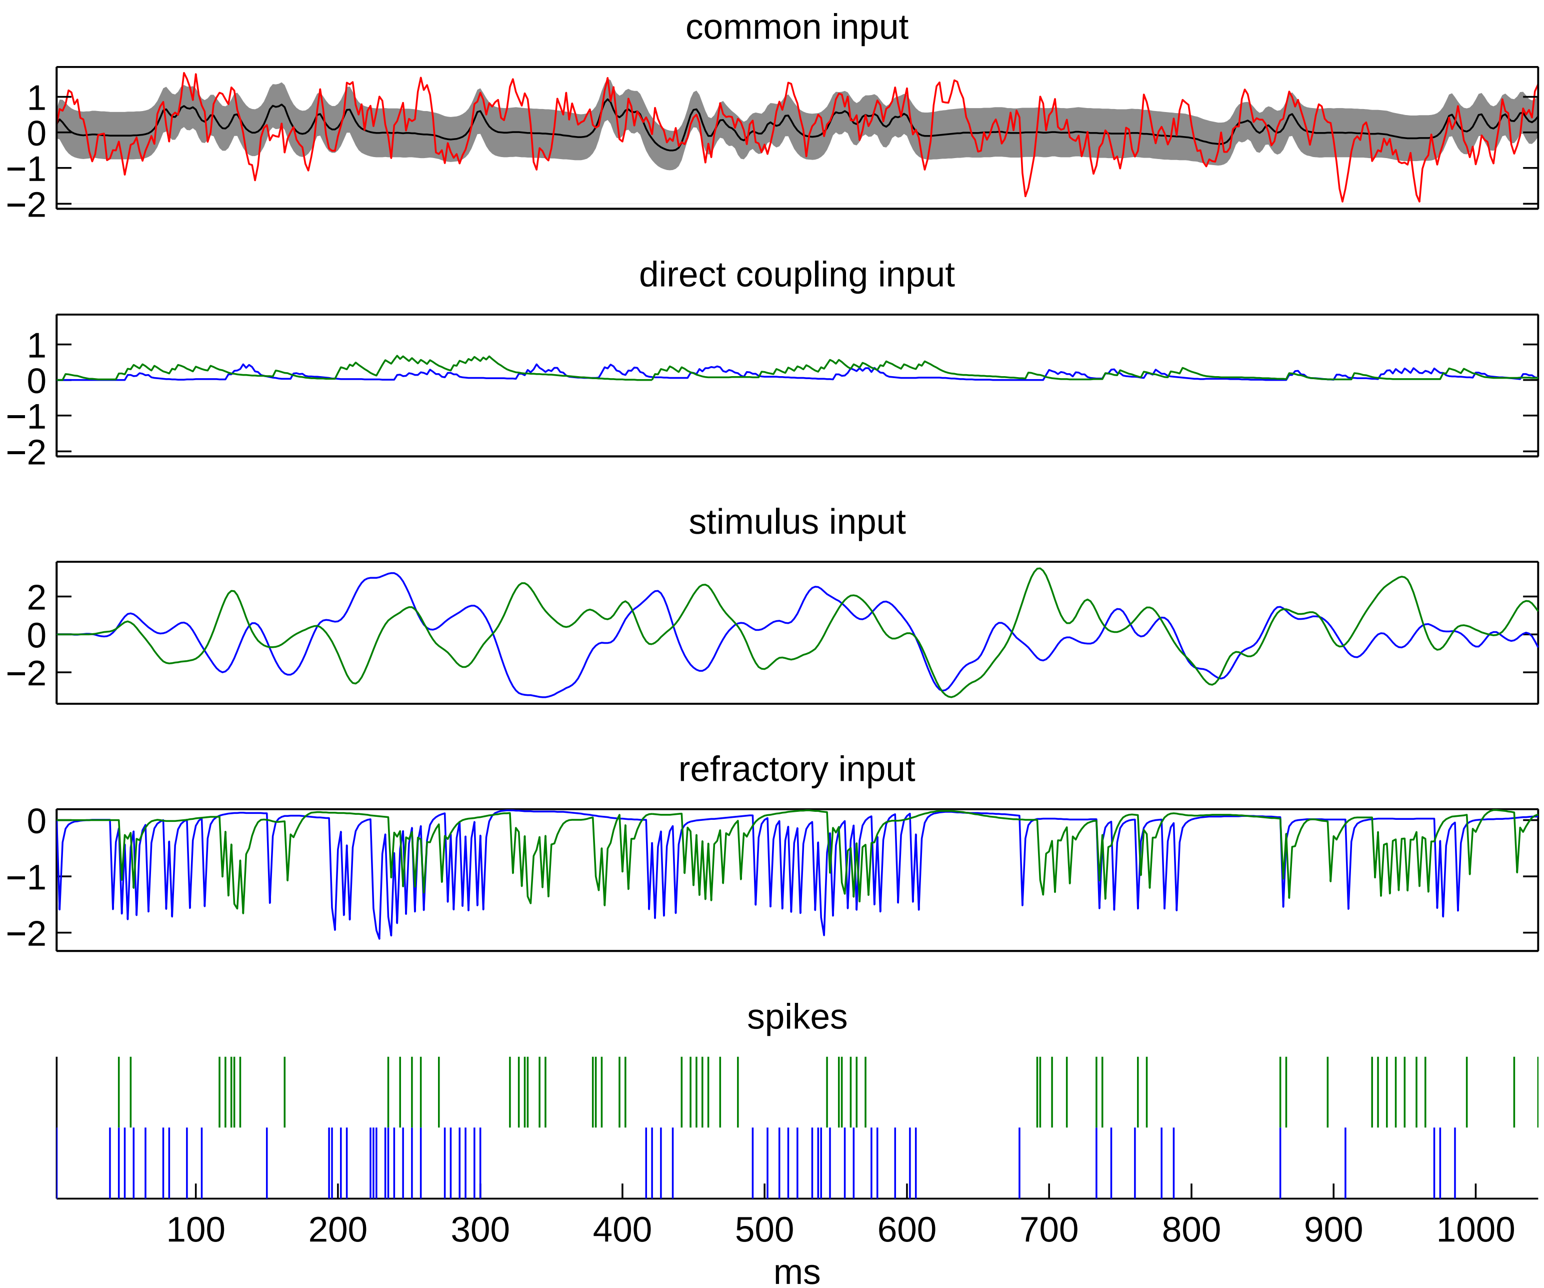

Common-input-only model captures x-corrs

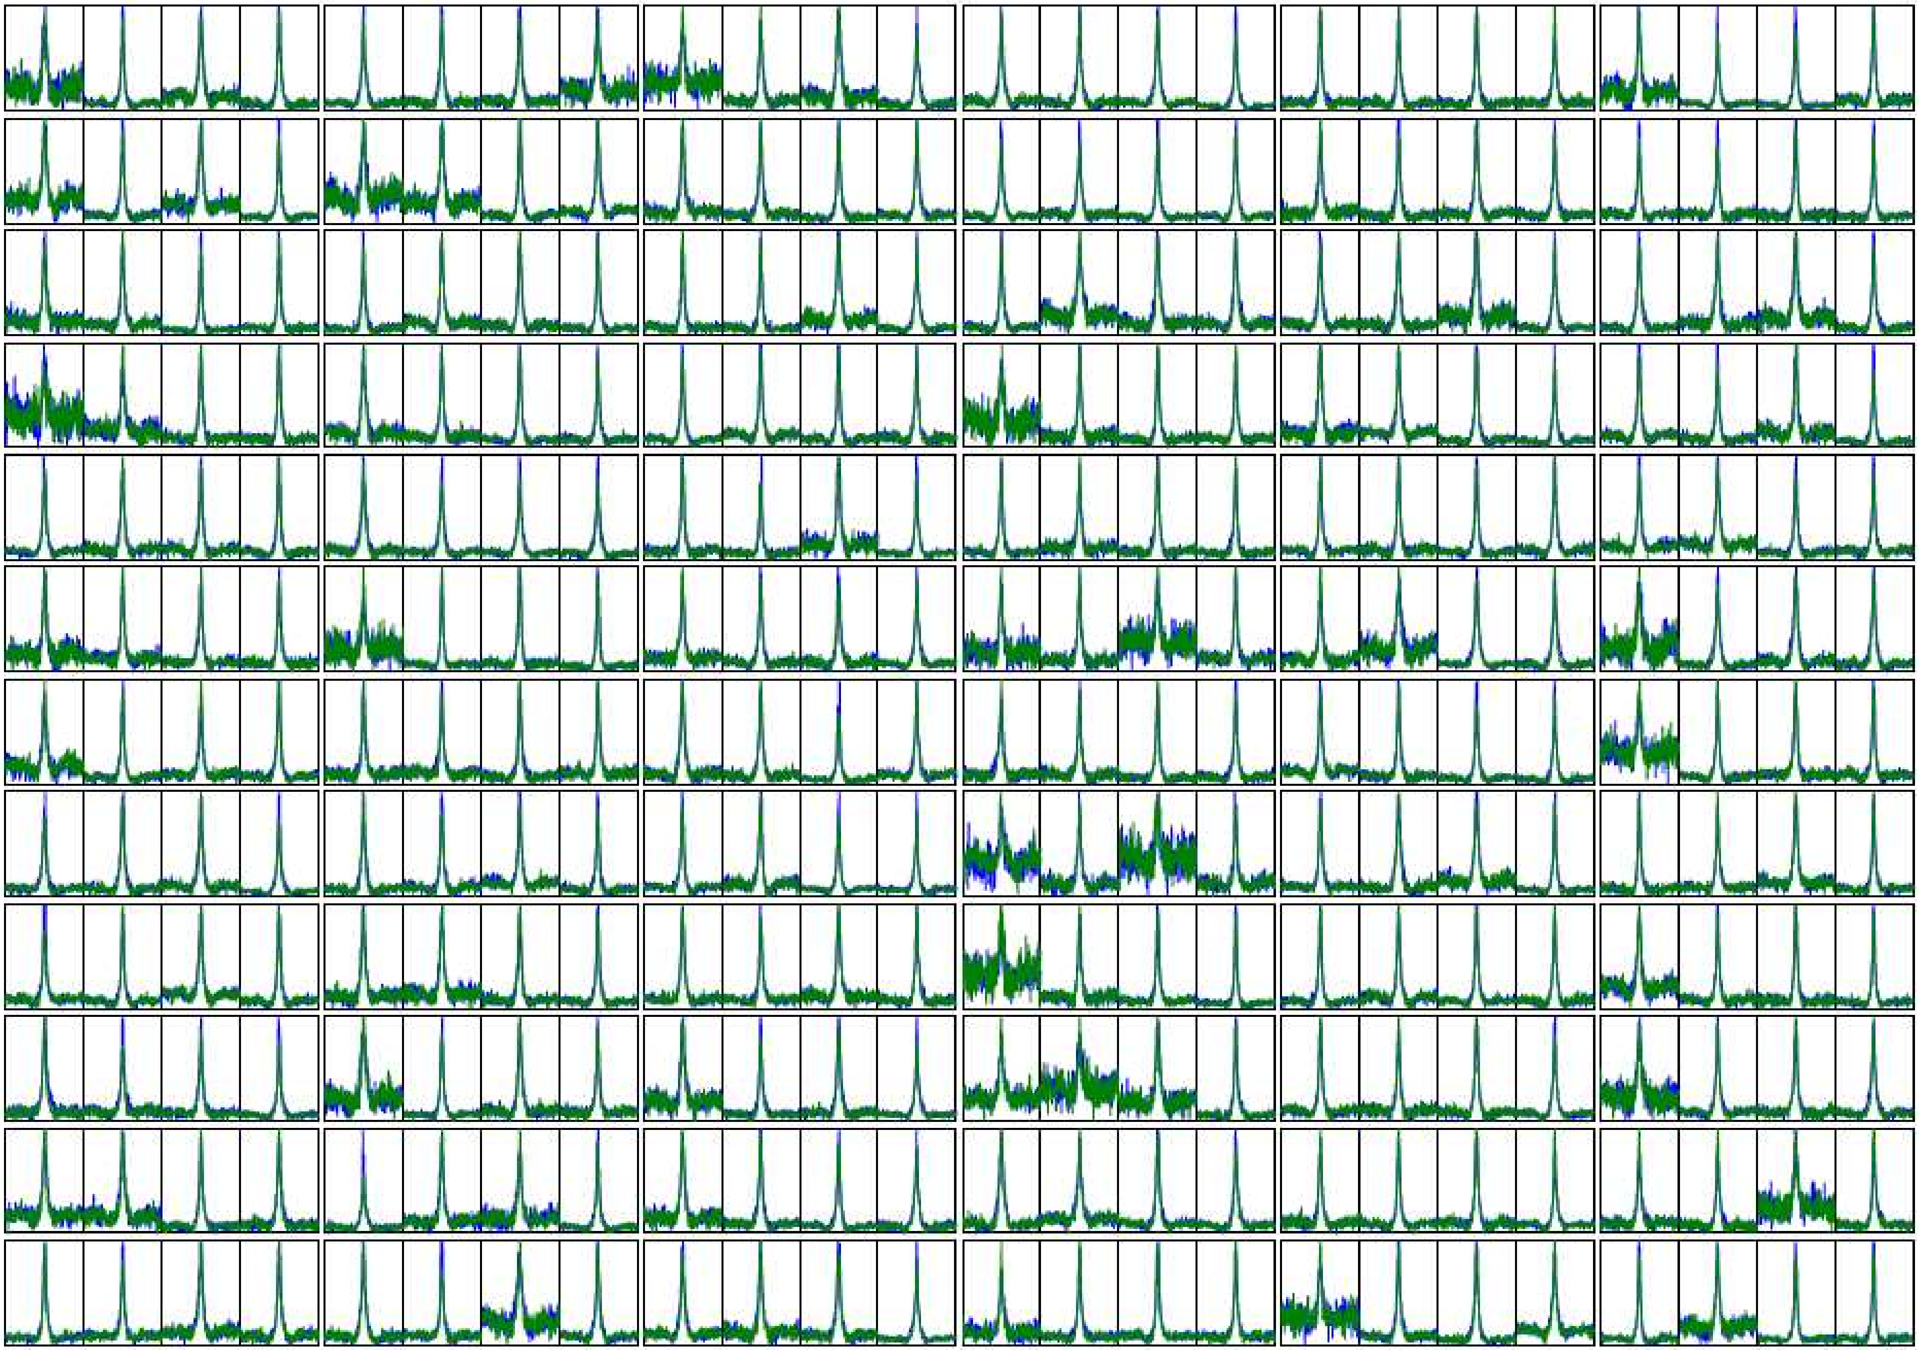

Decoding the stimulus and hidden input

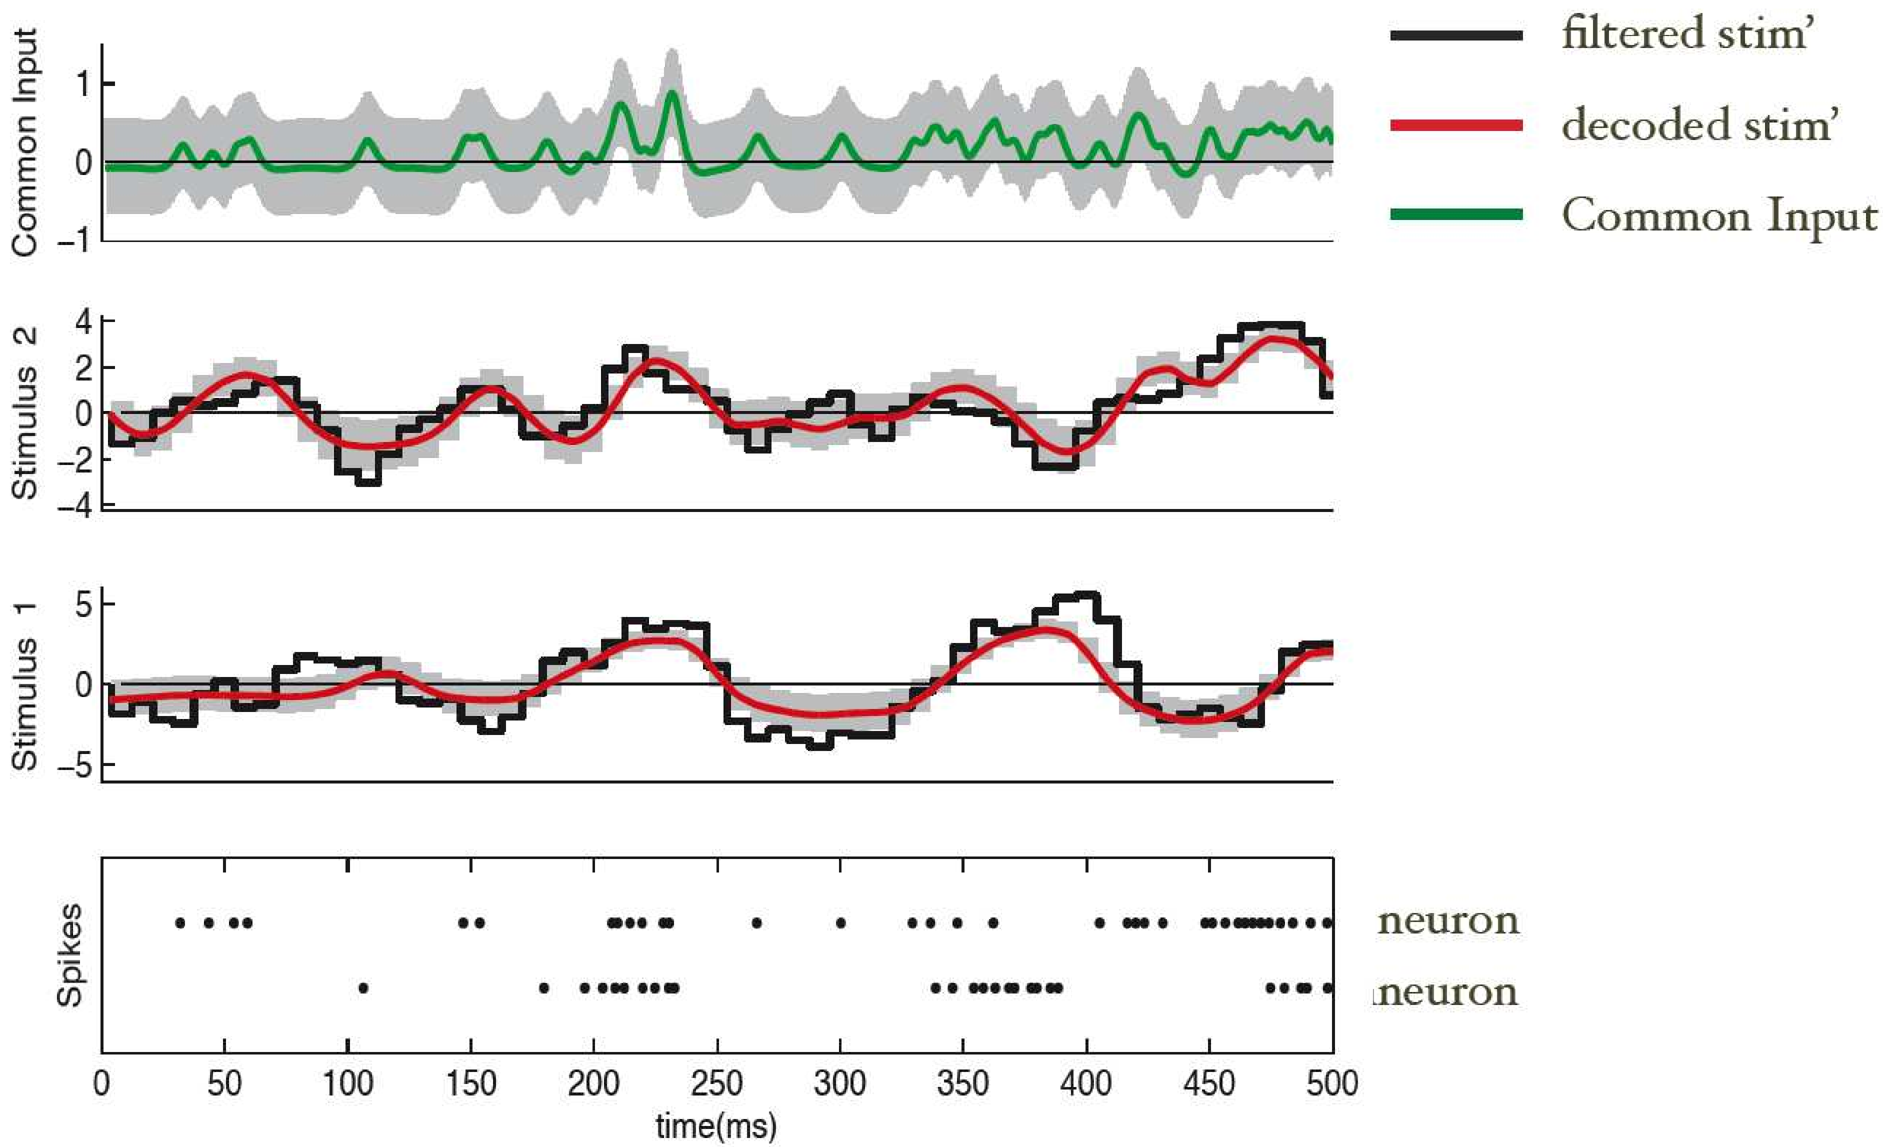

$$\arg\max_{\vec{x}} p(\vec{x}|y,\theta) = \arg\max_{\vec{x}} \int p(\vec{x},Q|y,\theta)dQ \approx \arg\max_{\vec{x},Q} p(\vec{x},Q|y,\theta)$$

Models lead to similar decoding performance

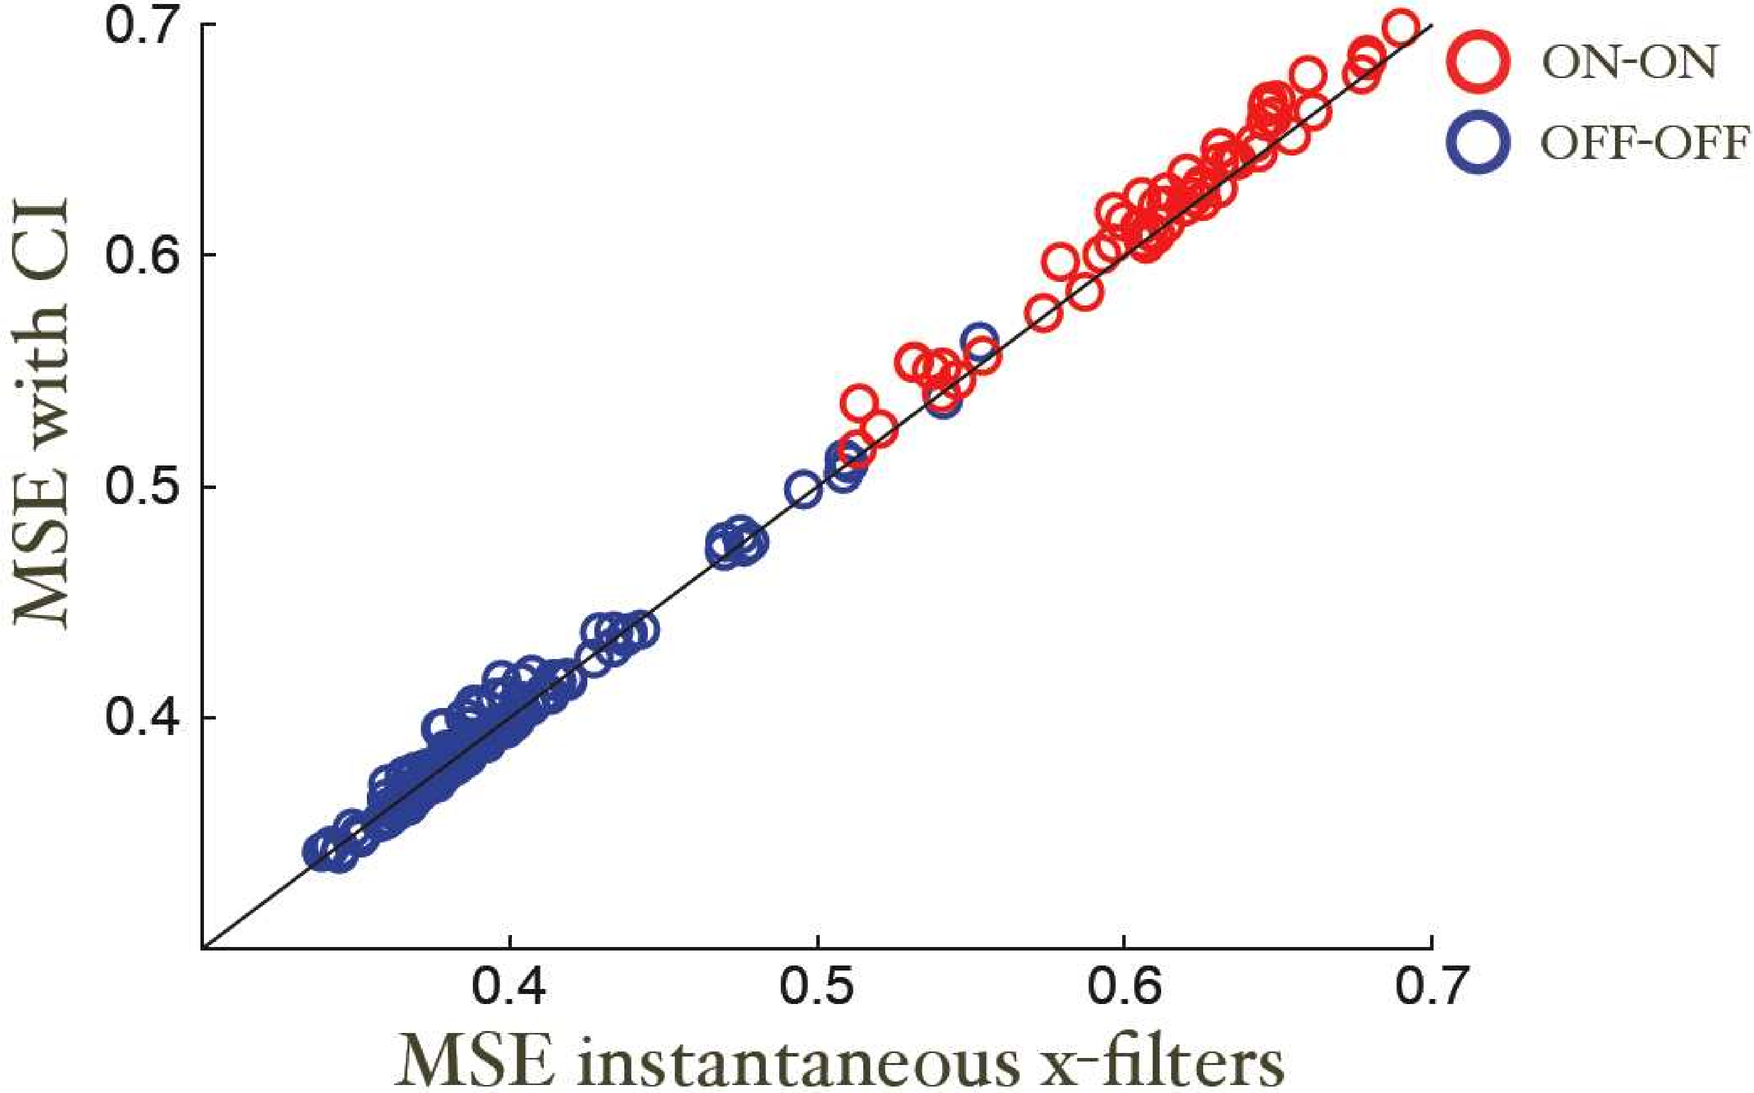

...but CI model is more robust to spike jitter and deletions.

Next steps: inferring cones

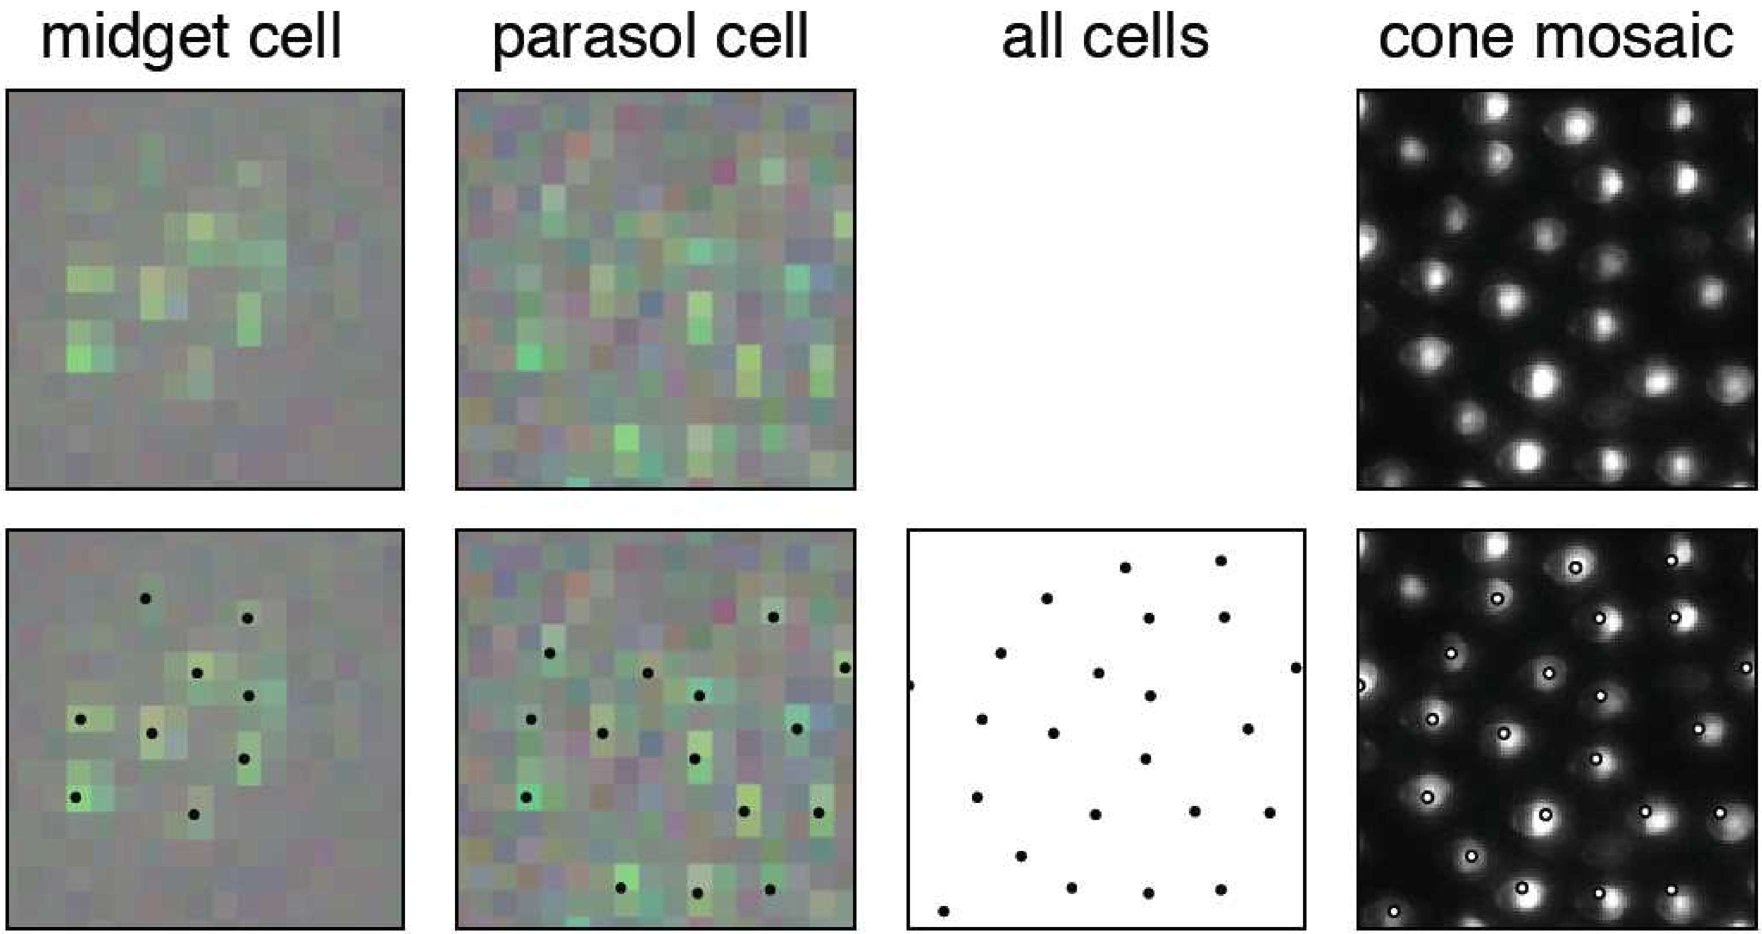

— cone locations and color identity can be inferred accurately via maximum a posteriori estimates.

ON midget

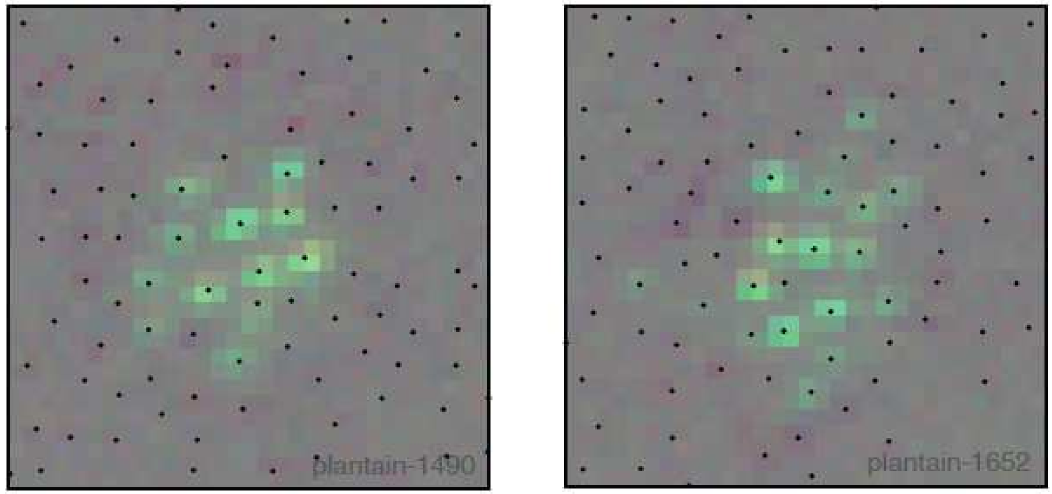

OFF midget

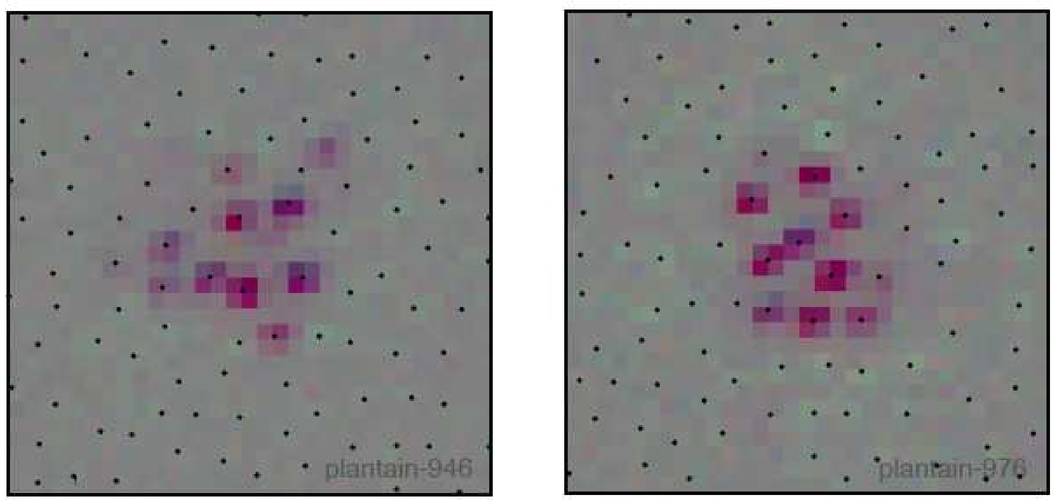

Next steps: inferring circuitry?

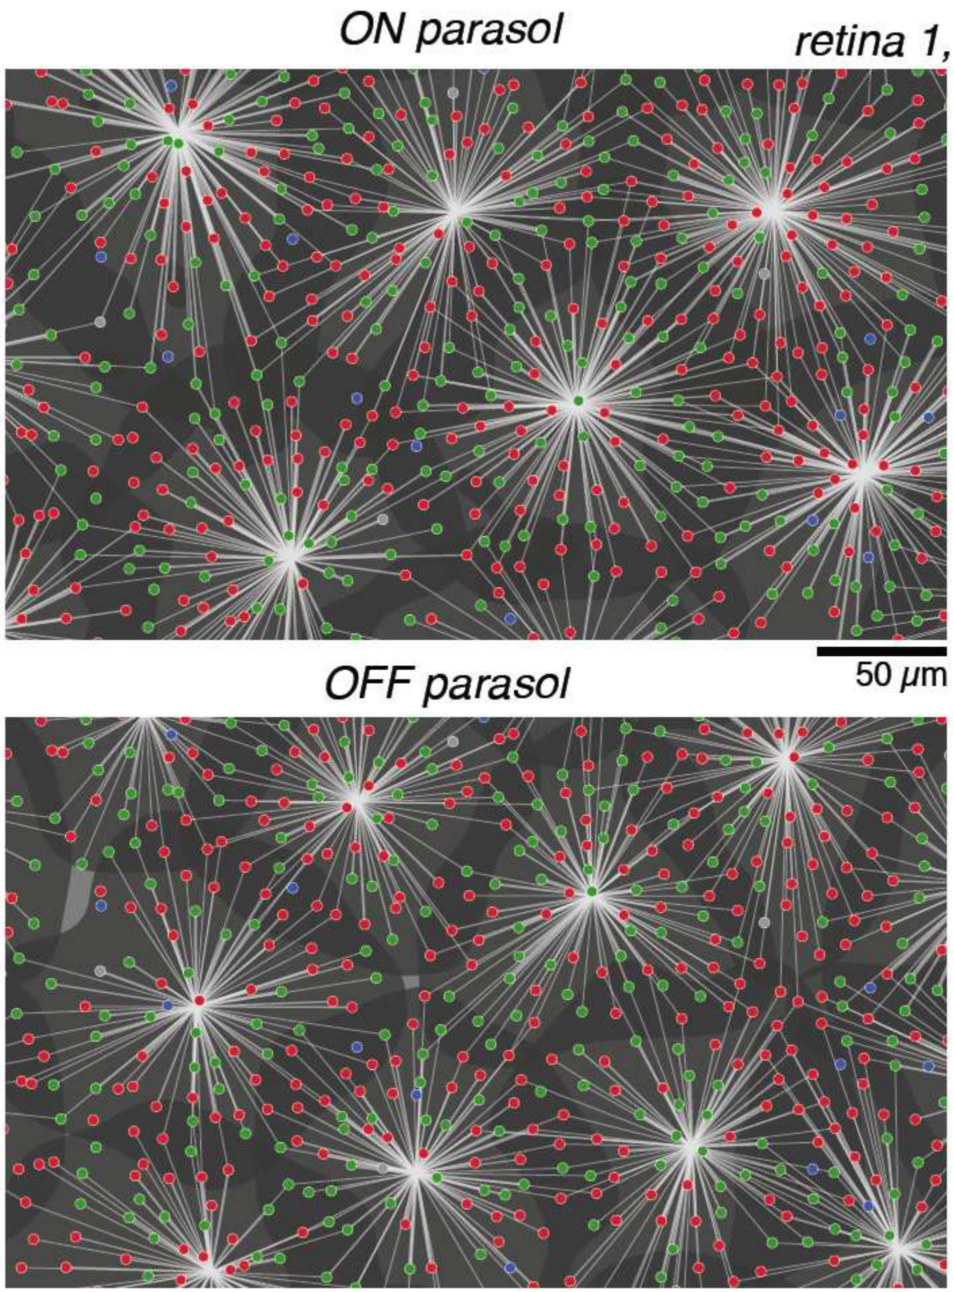

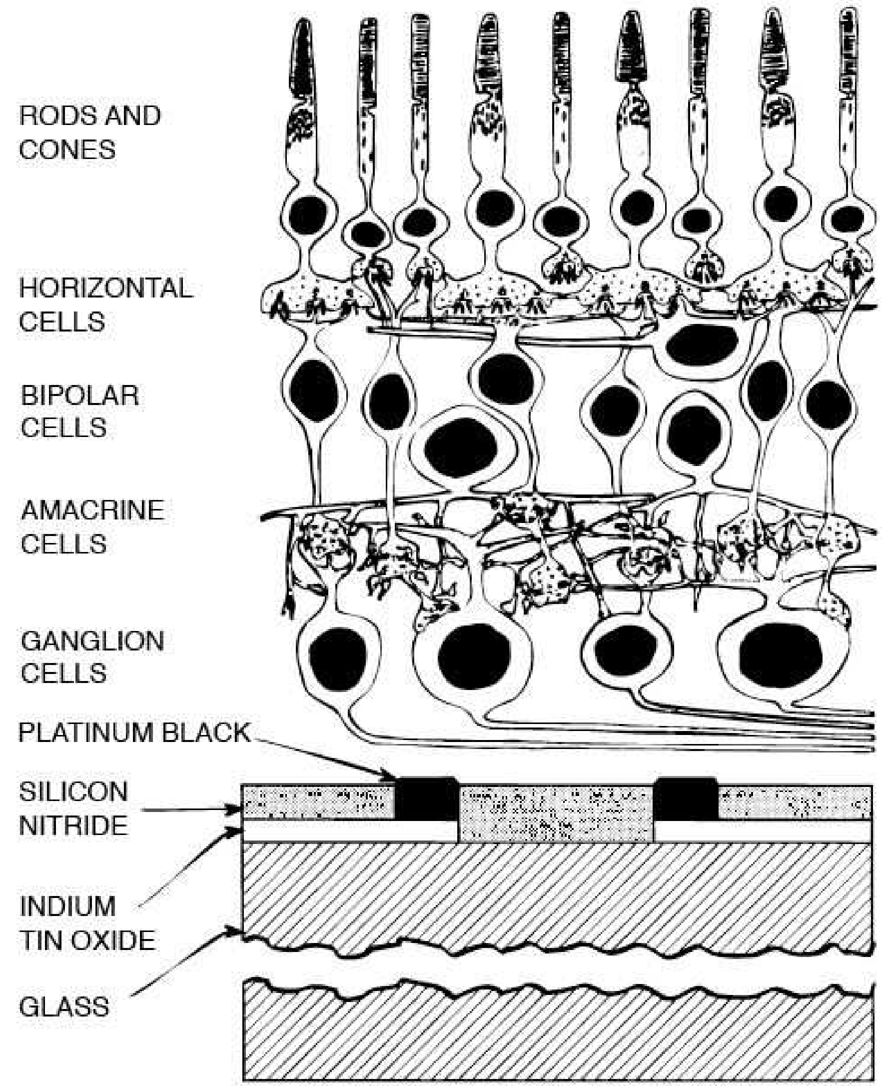

References

Ahmadian, Y., Pillow, J., and Paninski, L. (2009). Efficient Markov Chain Monte Carlo methods for decoding population spike trains. *Under review, Neural Computation*.

Frechette, E., Sher, A., Grivich, M., Petrusca, D., Litke, A., and Chichilnisky, E. (2005). Fidelity of the ensemble code for visual motion in the primate retina. *J Neurophysiol*, 94(1):119–135.

Khuc-Trong, P. and Rieke, F. (2008). Origin of correlated activity between parasol retinal ganglion cells. *Nature Neuroscience*, 11:1343–1351.

Lalor, E., Ahmadian, Y., and Paninski, L. (2008). Optimal decoding of stimulus velocity using a probabilistic model of ganglion cell populations in primate retina. *Journal of Vision*, Under review.

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. (2009). A new look at state-space models for neural data. *Journal of Computational Neuroscience*, In press.

Pillow, J., Ahmadian, Y., and Paninski, L. (2009). Model-based decoding, information estimation, and change-point detection in multi-neuron spike trains. *Under review, 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.

Pitkow, X., Sompolinsky, H., and Meister, M. (2007). A neural computation for visual acuity in the presence of eye movements. *PLOS Biology*, 5.

Truccolo, W., Eden, U., Fellows, M., Donoghue, J., and Brown, E. (2005). A point process framework for relating neural spiking activity to spiking history, neural ensemble and extrinsic covariate effects. *Journal of Neurophysiology*, 93:1074–1089.

Victor, J. (2000). How the brain uses time to represent and process visual information. *Brain Research*, 886:33–46.

Weiss, Y., Simoncelli, E., and Adelson, E. (2002). Motion illusions as optimal percepts. *Nature Neuroscience*, 5:598–604.