Coding and computation by neural ensembles in the retina

Liam Paninski

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

May 20, 2008

Support: NIH CRCNS award, Sloan Research Fellowship, NSF CAREER award.

The neural code

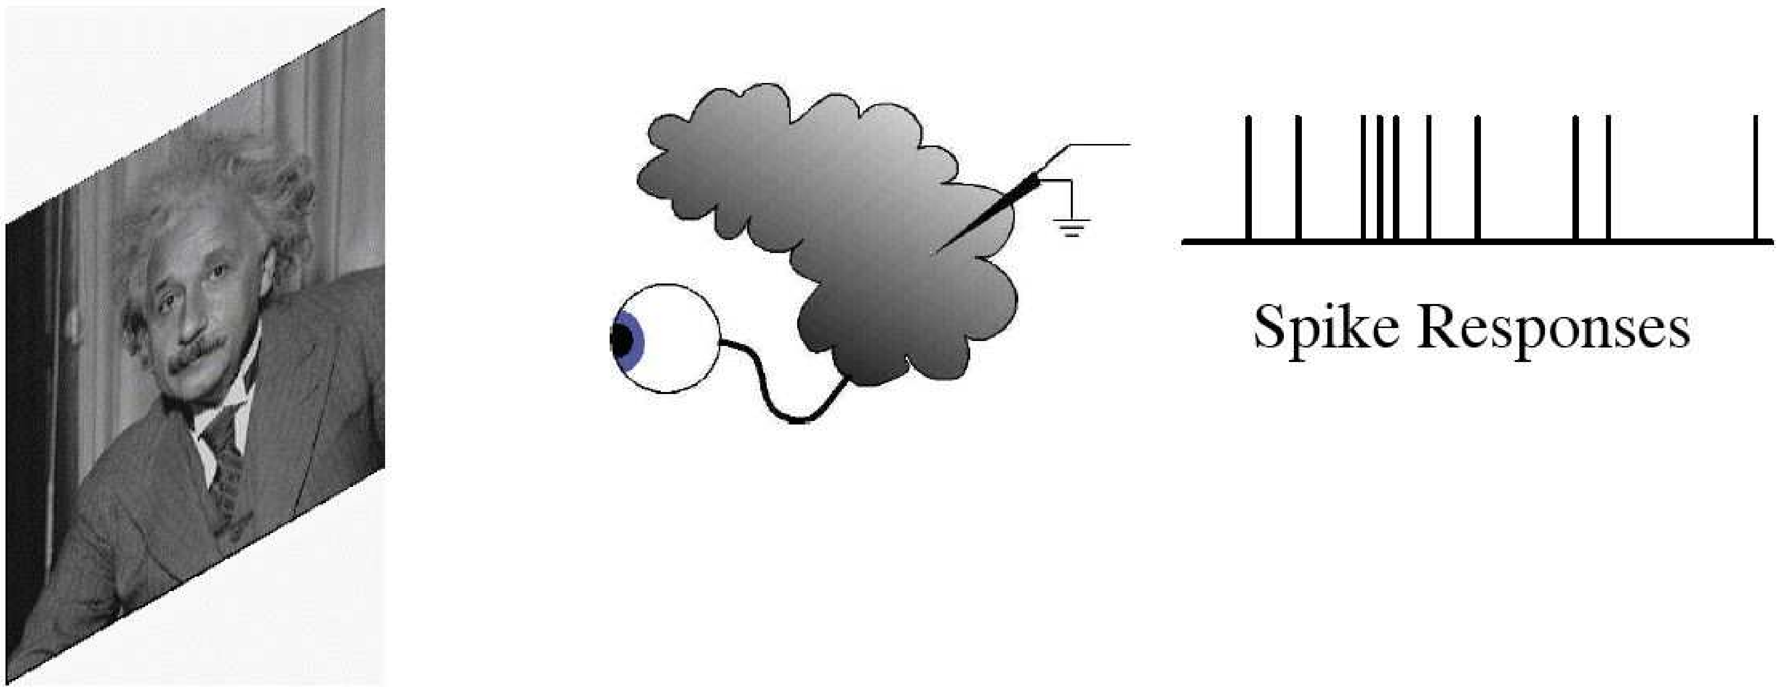

Input-output relationship between

- External observables x (sensory stimuli, motor responses...)
- Neural variables y (spike trains, population activity...)

Encoding problem: $p(y|x)$; decoding problem: $p(x|y)$

Retinal ganglion neuronal data

Preparation: dissociated macaque retina

— extracellularly-recorded responses of populations of RGCs

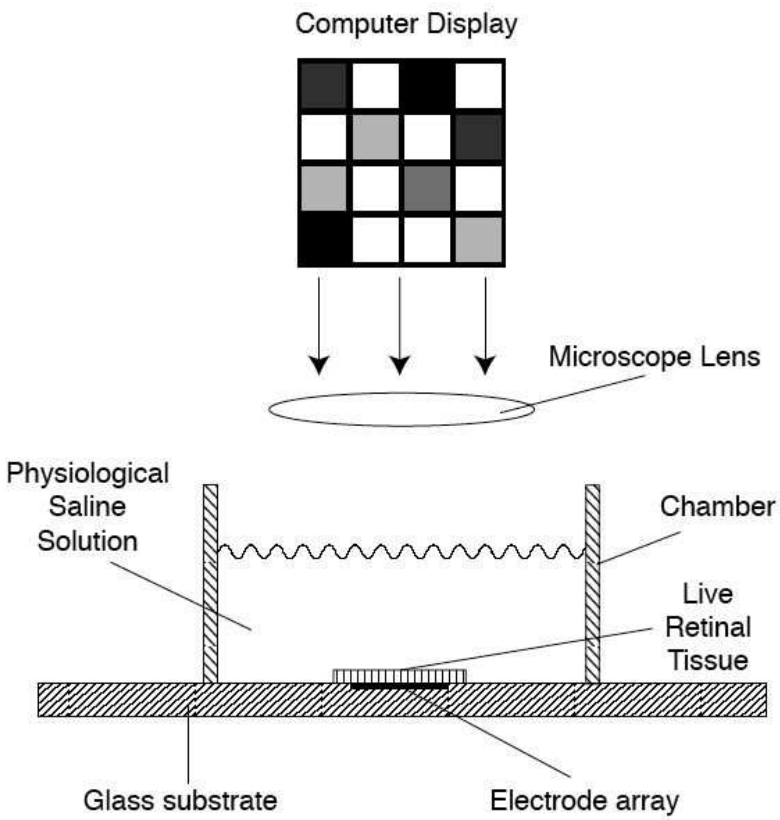

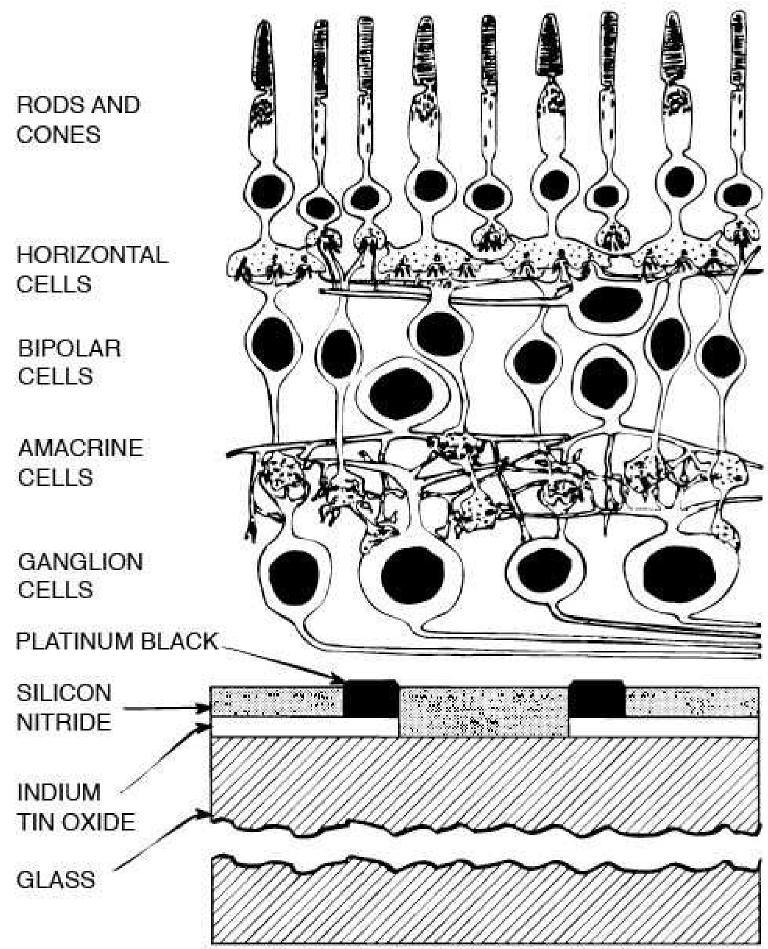

Stimulus: random spatiotemporal visual stimuli (Pillow et al., 2008)

Receptive fields tile visual space

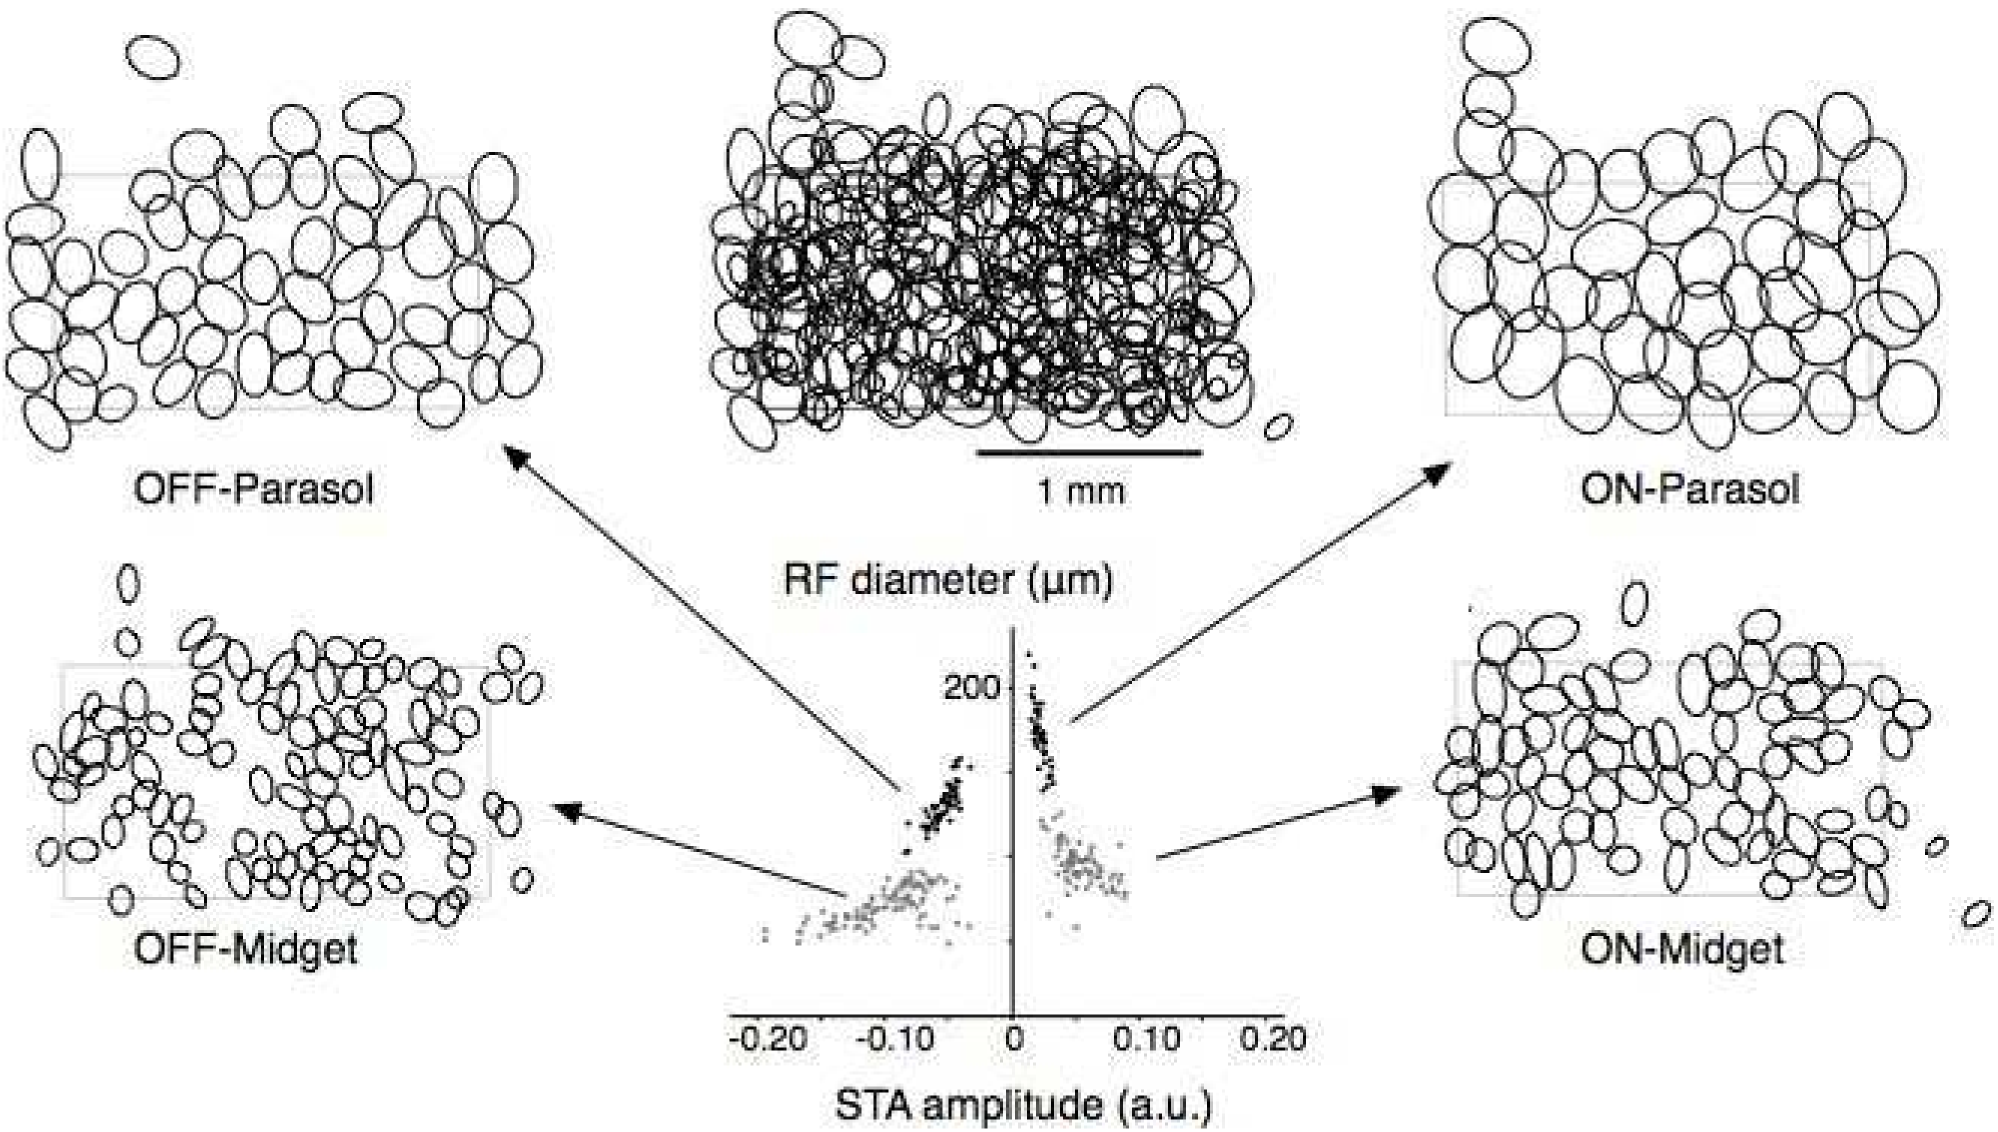

Multineuronal point-process model

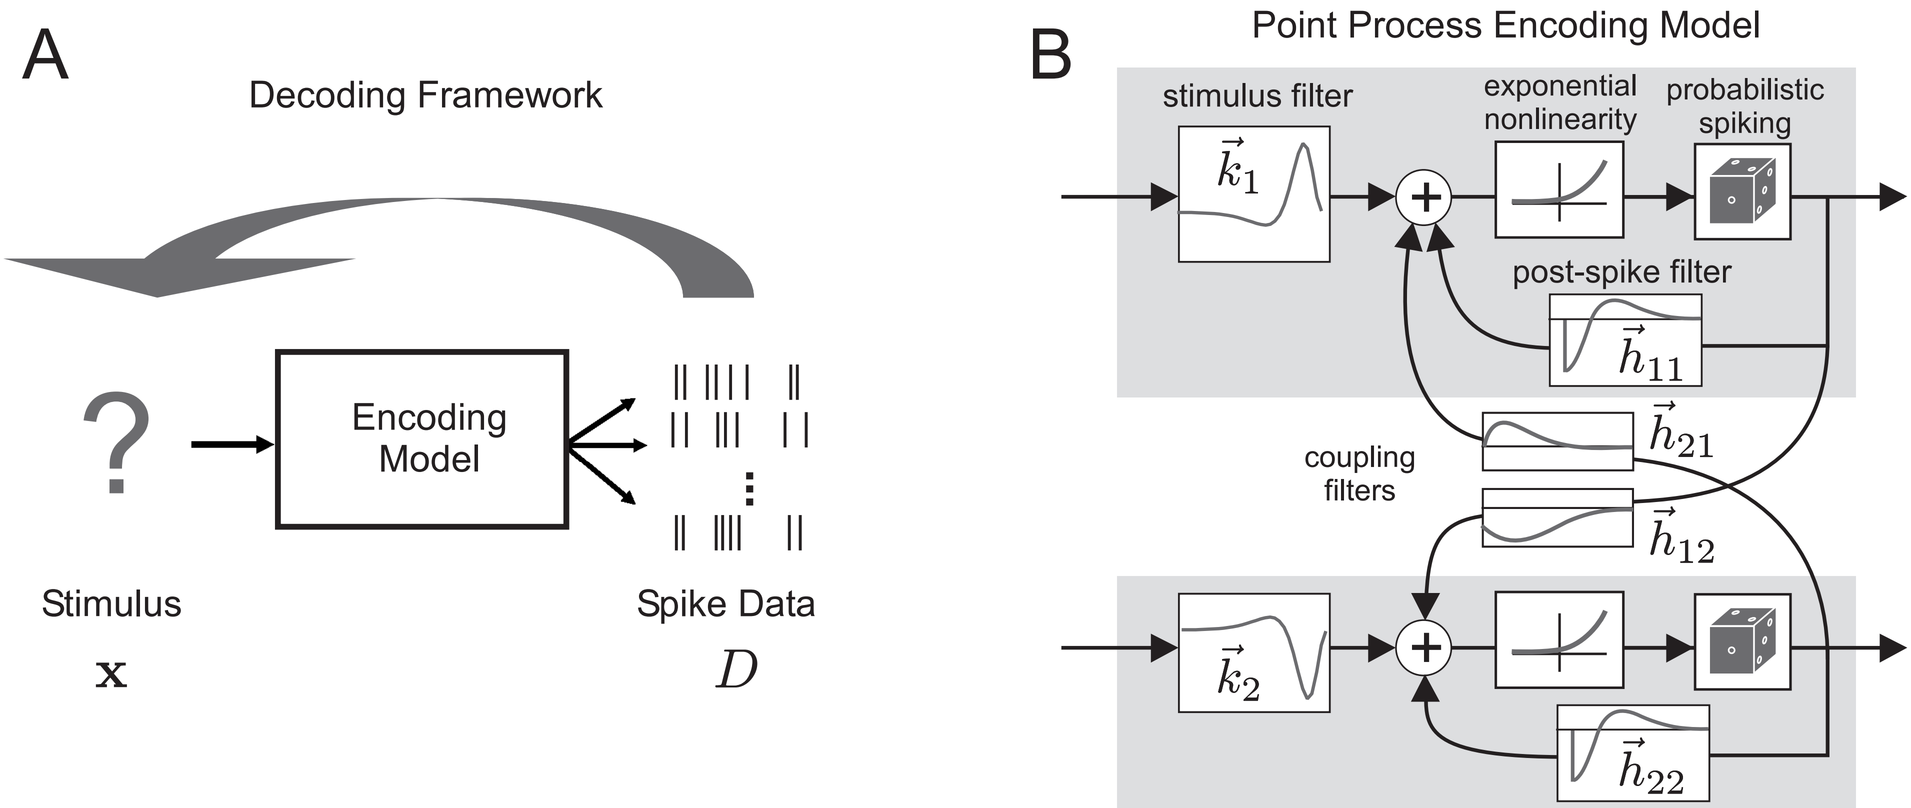

$$\lambda_i(t) = f\left(b + \vec{k}_i \cdot \vec{x}(t) + \sum_{i',j} h_{i',j} n_{i'}(t-j)\right),$$

— Fit by maximum likelihood (concave optimization) (Paninski, 2004)

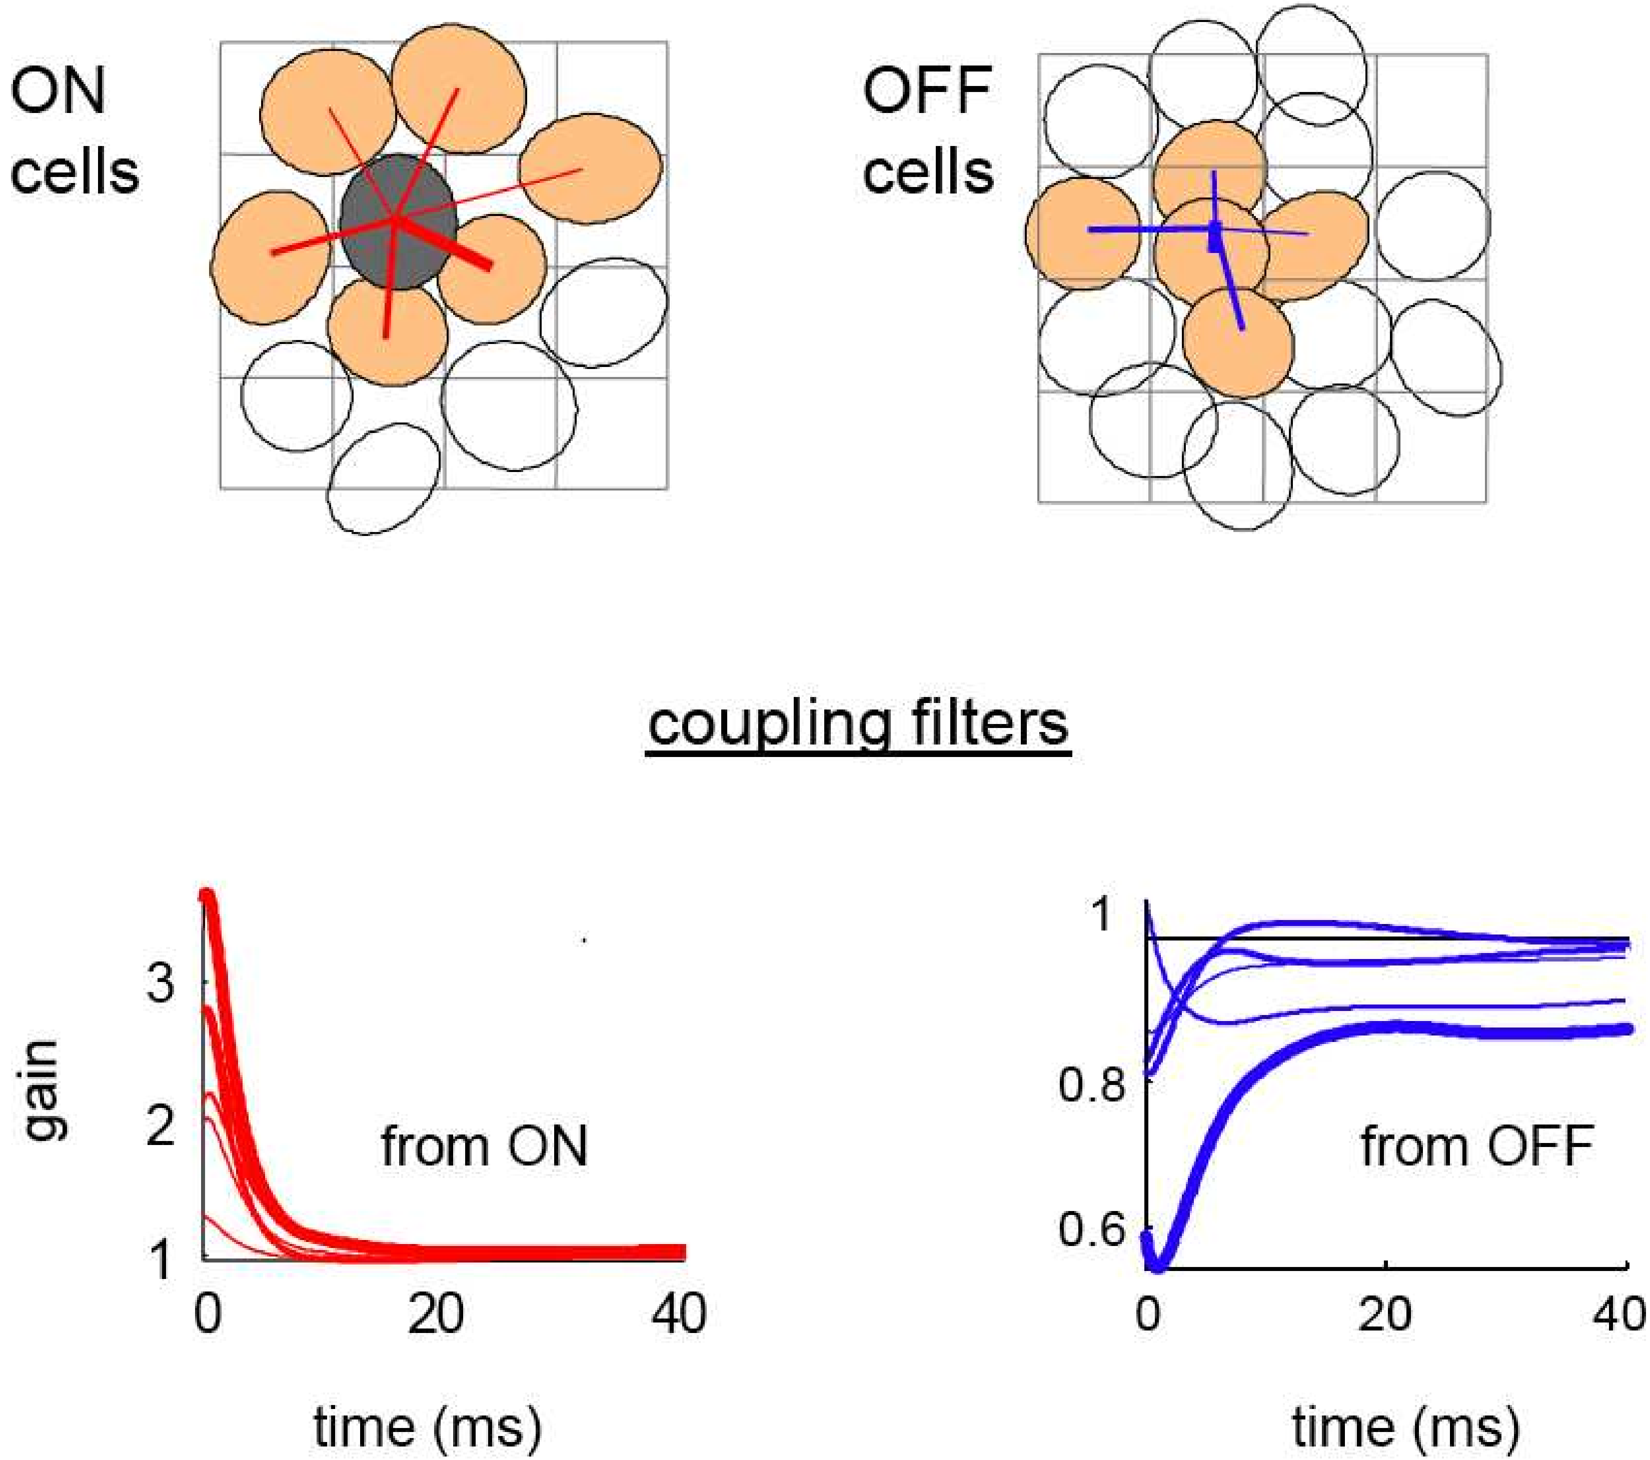
ON
cells
OFF
cells
coupling filters
gain
3
2
1
0
20
40
from ON
time (ms)
1
0.8
0.6
0
20
40
from OFF
time (ms)

Network vs. stimulus drive

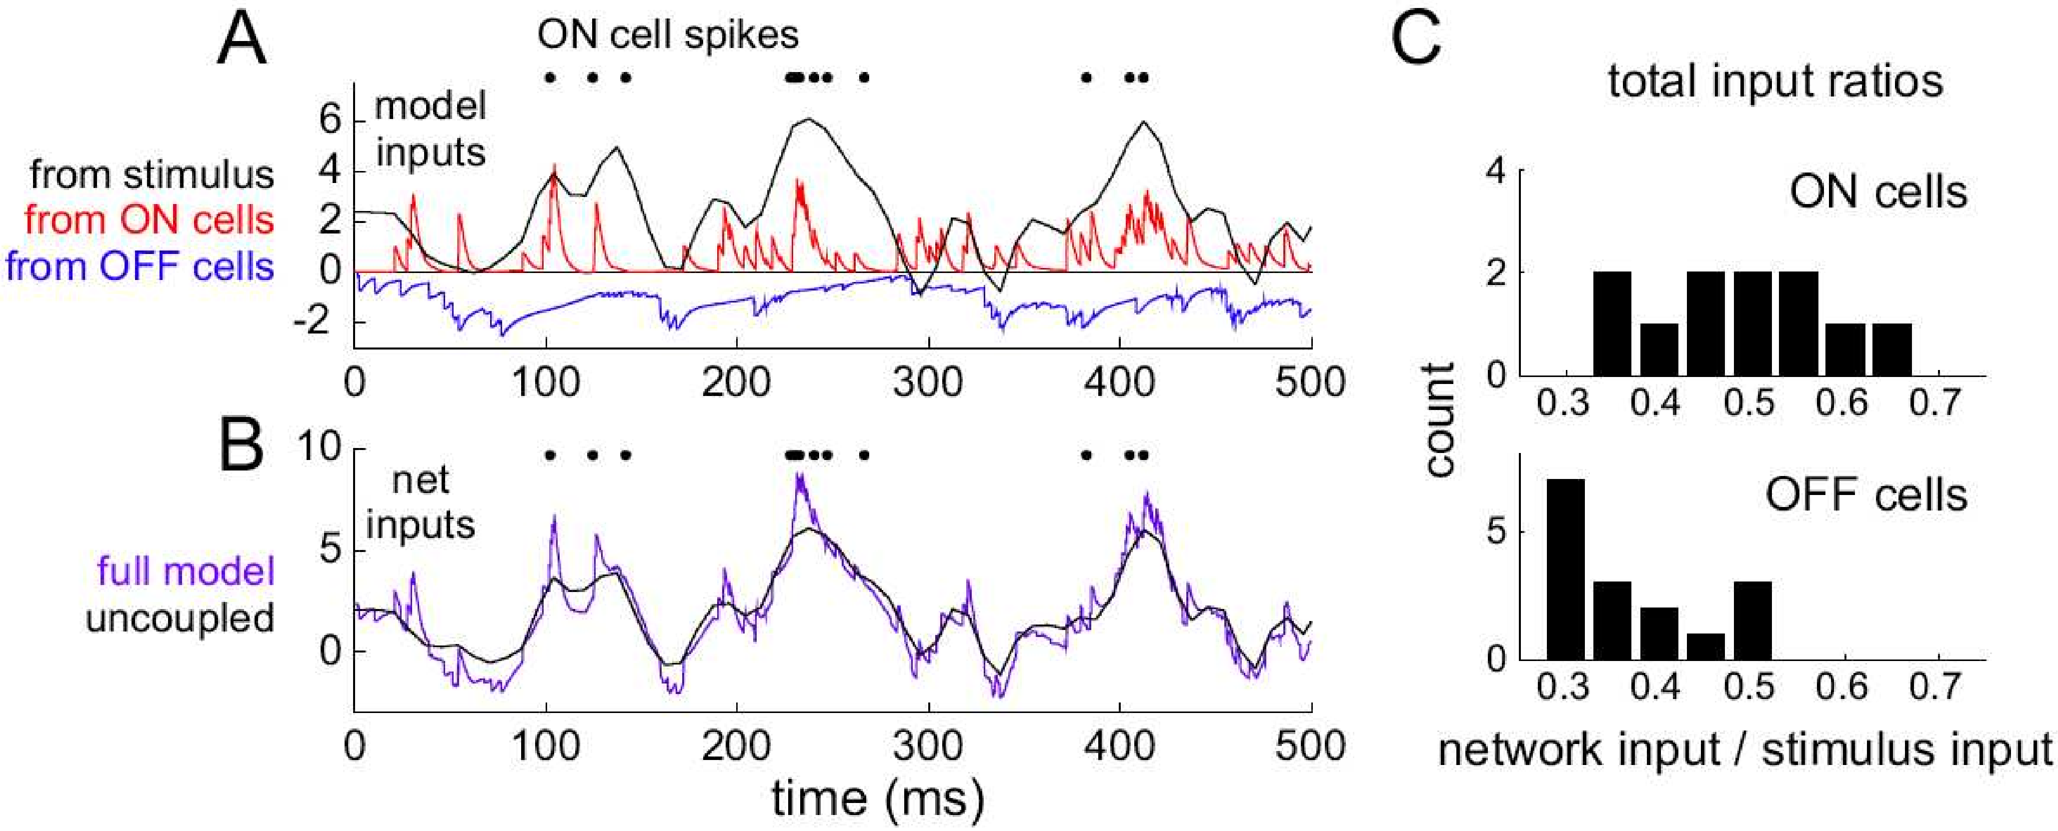

— Network effects are ≈ 50% as strong as stimulus effects

Spike Train Prediction

• improved prediction, but not of mean rate!

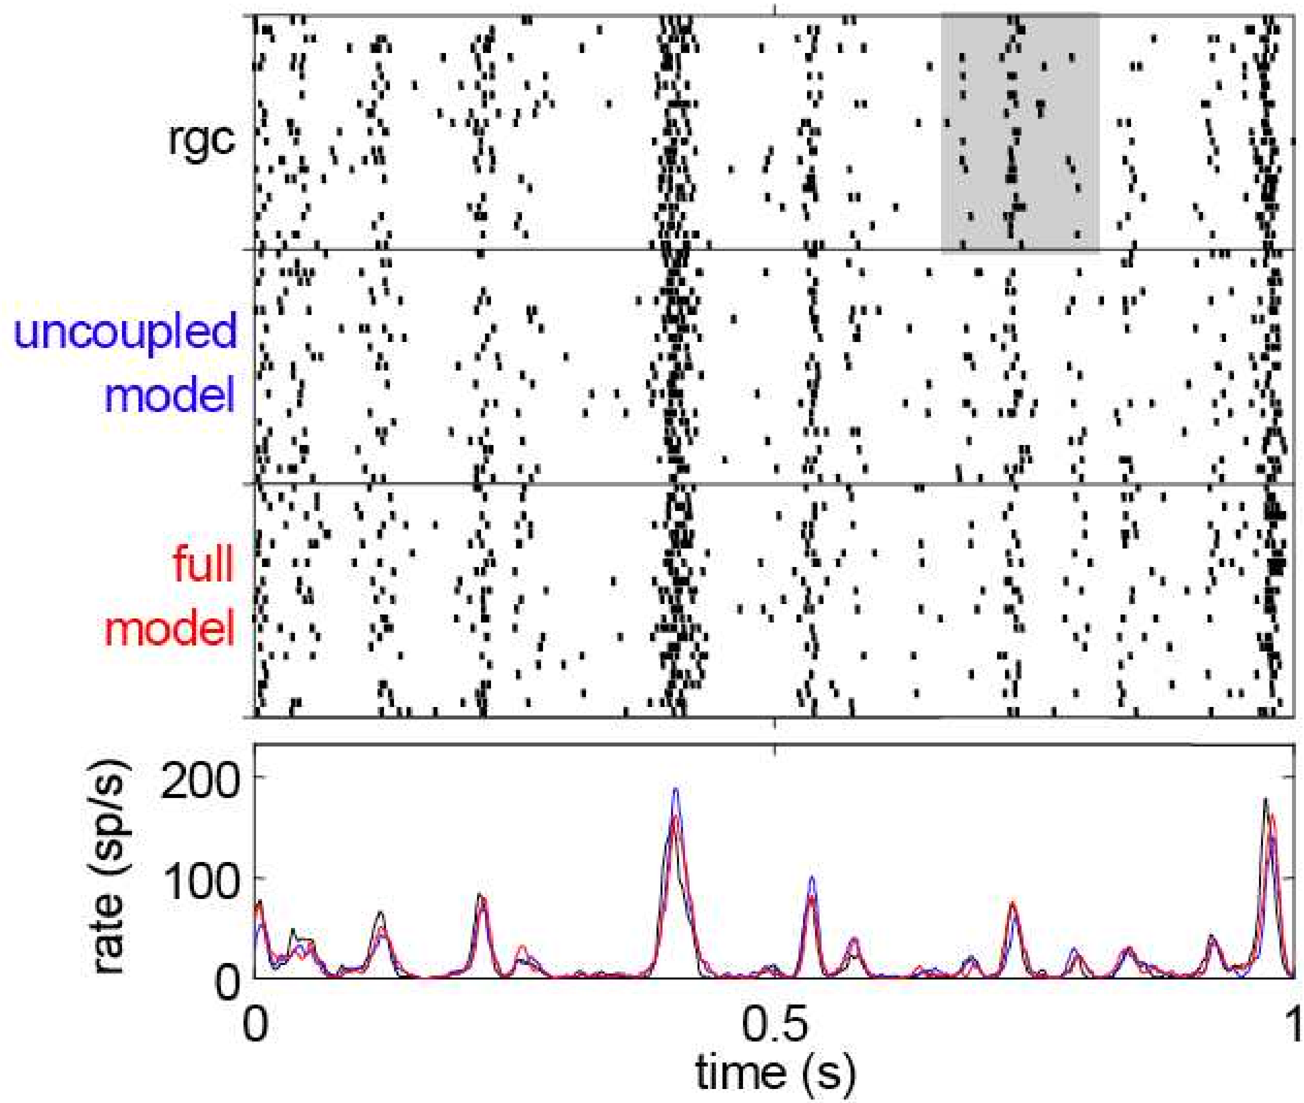

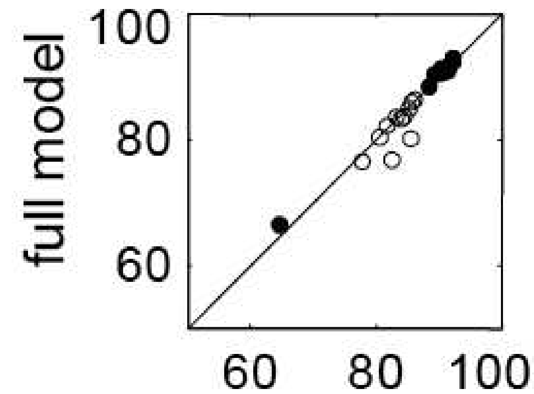

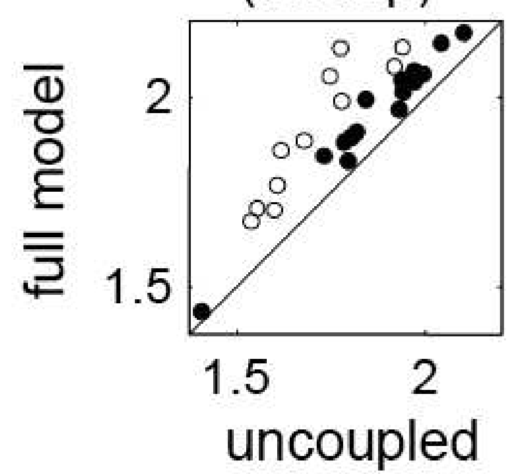

Network predictability analysis

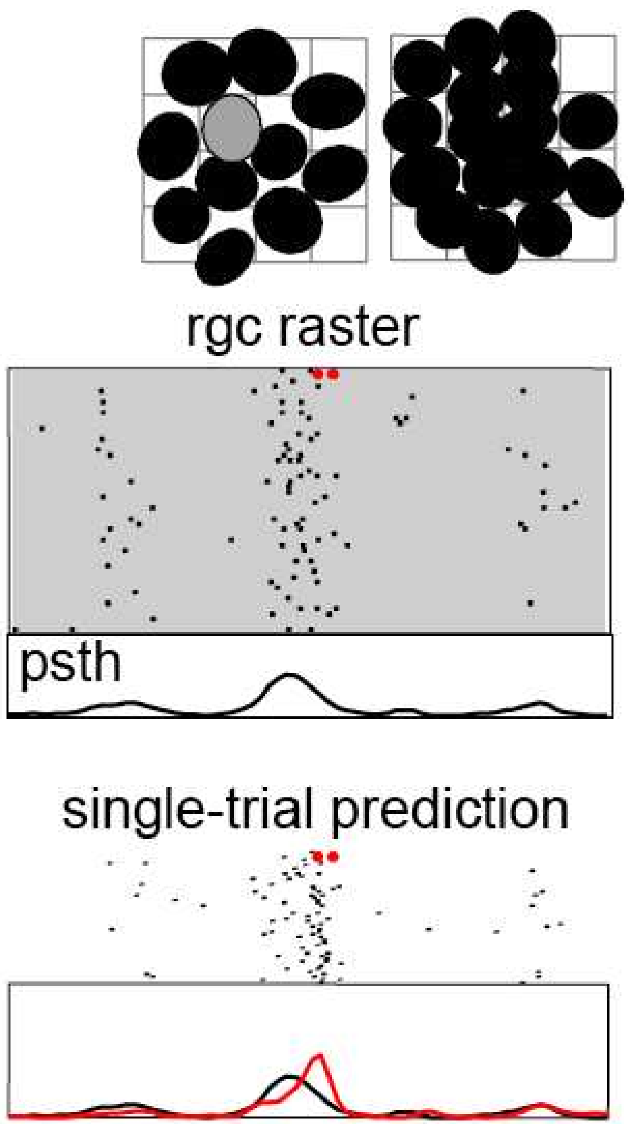

- fix all other neurons for a single trial
- draw single-trial predictions of this cell's spike train

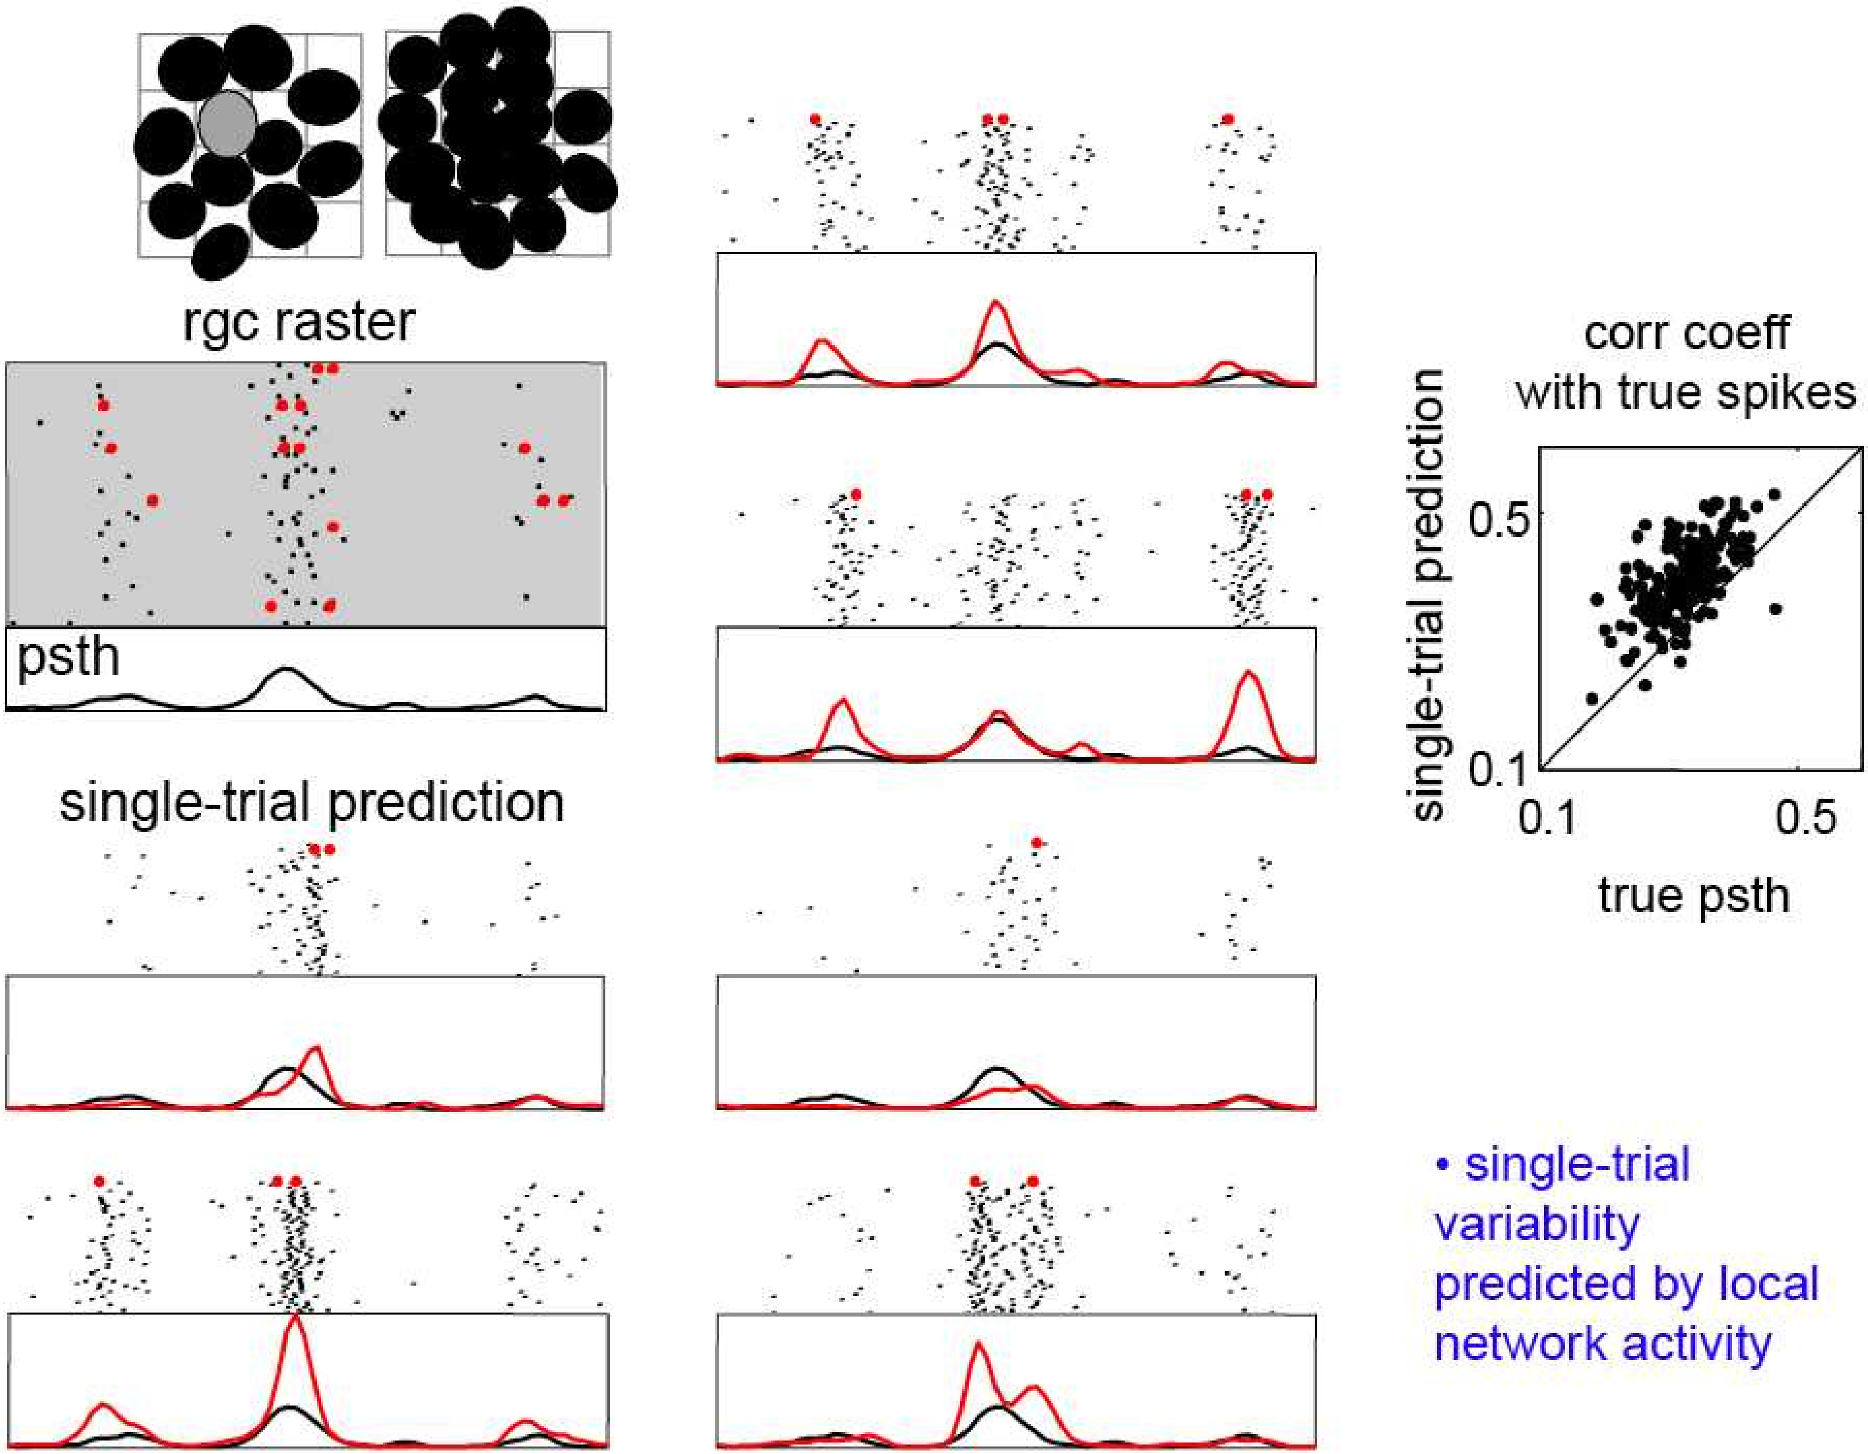
rgc raster
psth
single-trial prediction
corr coeff
with true spikes
single-trial prediction
0.5
0.1
0.1
0.5
true psth
• single-trial
variability
predicted by local
network activity

Model captures spatiotemporal cross-corrs

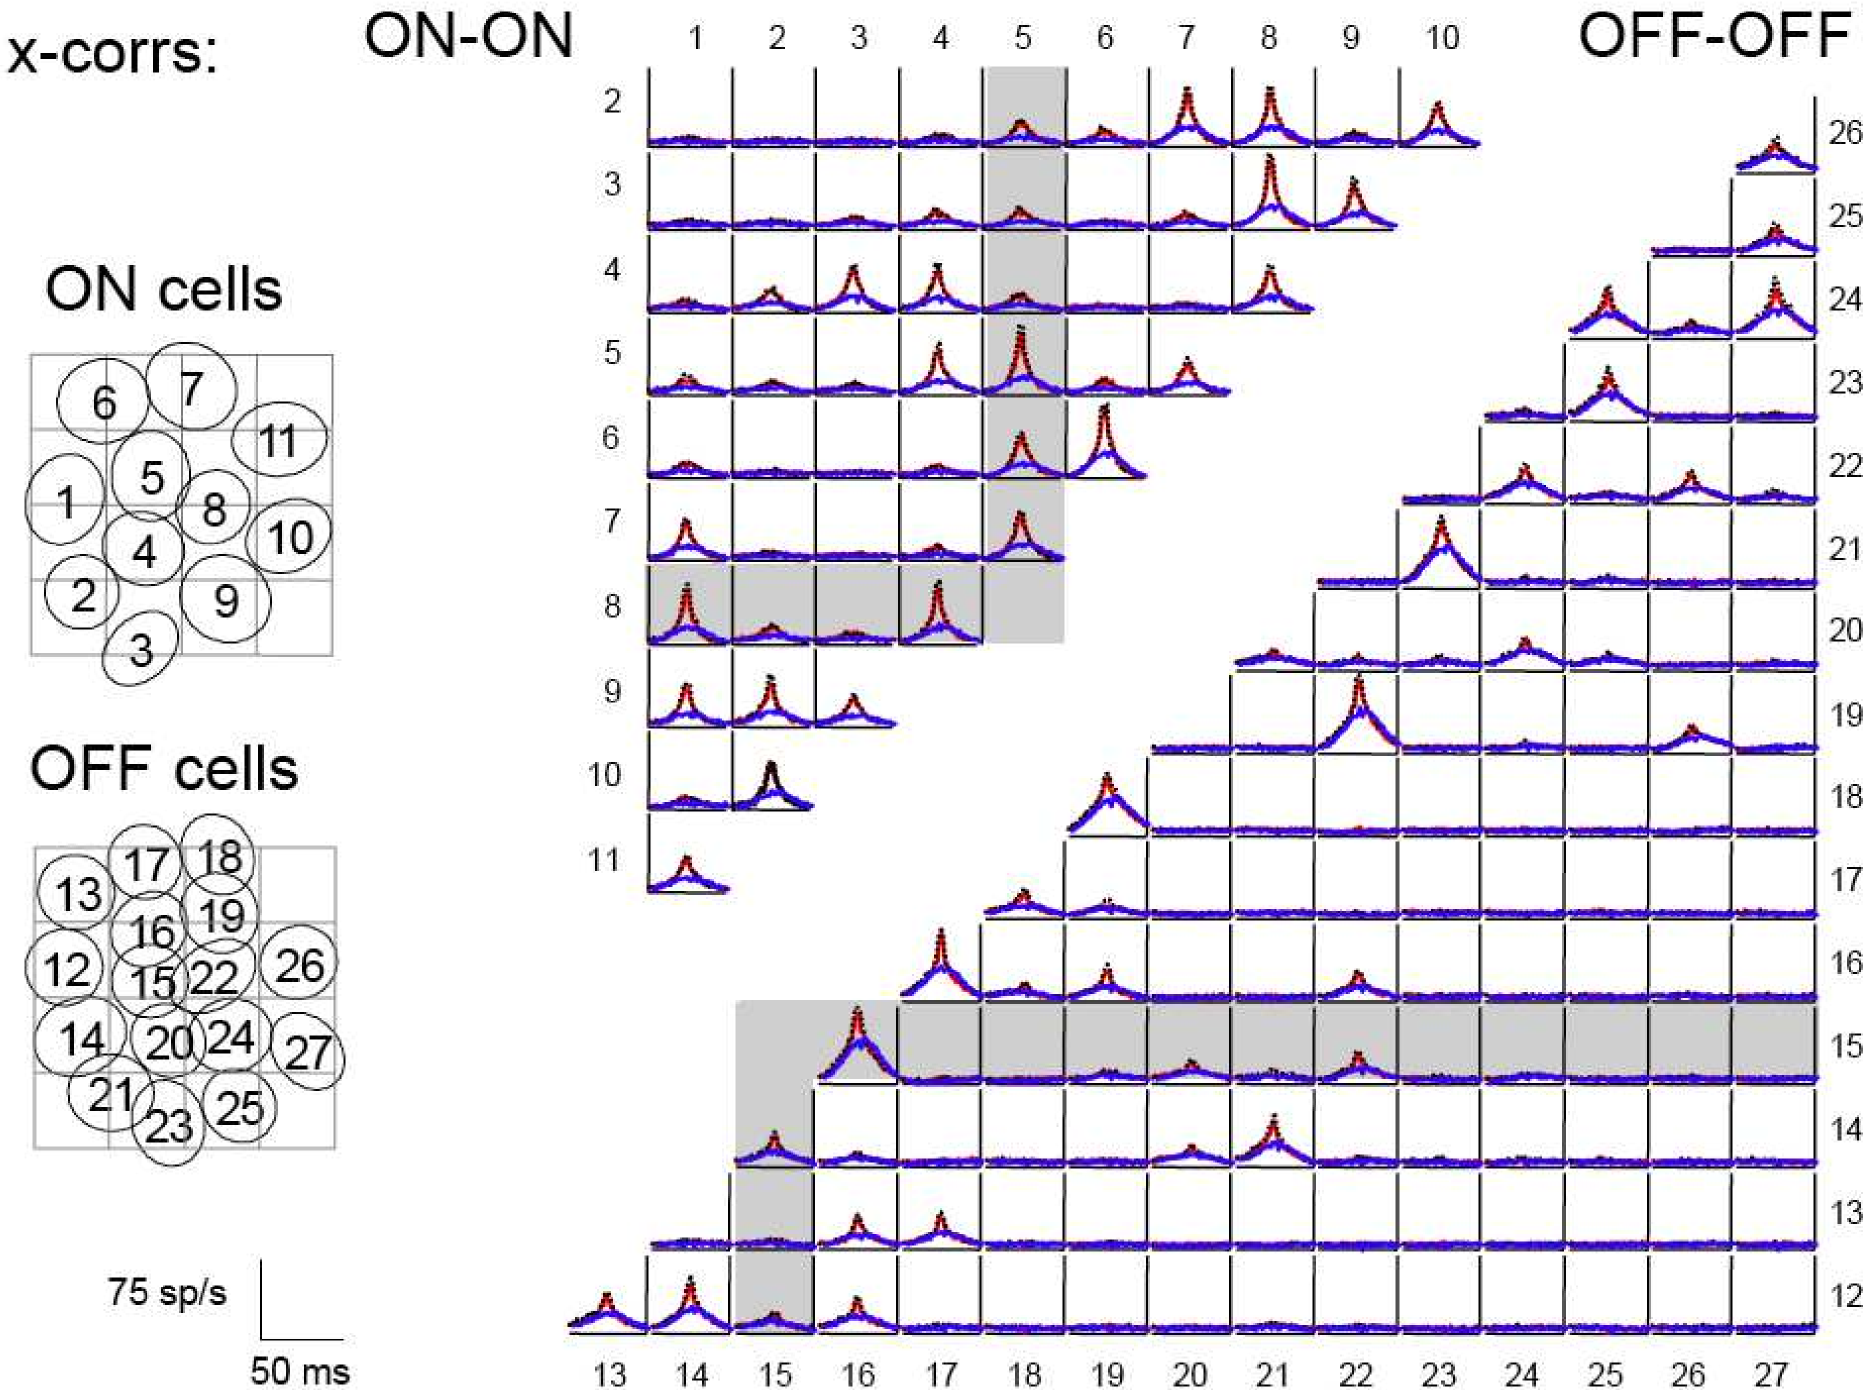

Maximum a posteriori decoding

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

— $\log P(spikes|\vec{x})$ is concave in $\vec{x}$: concave optimization again. (In fact, can be done in linear time.)

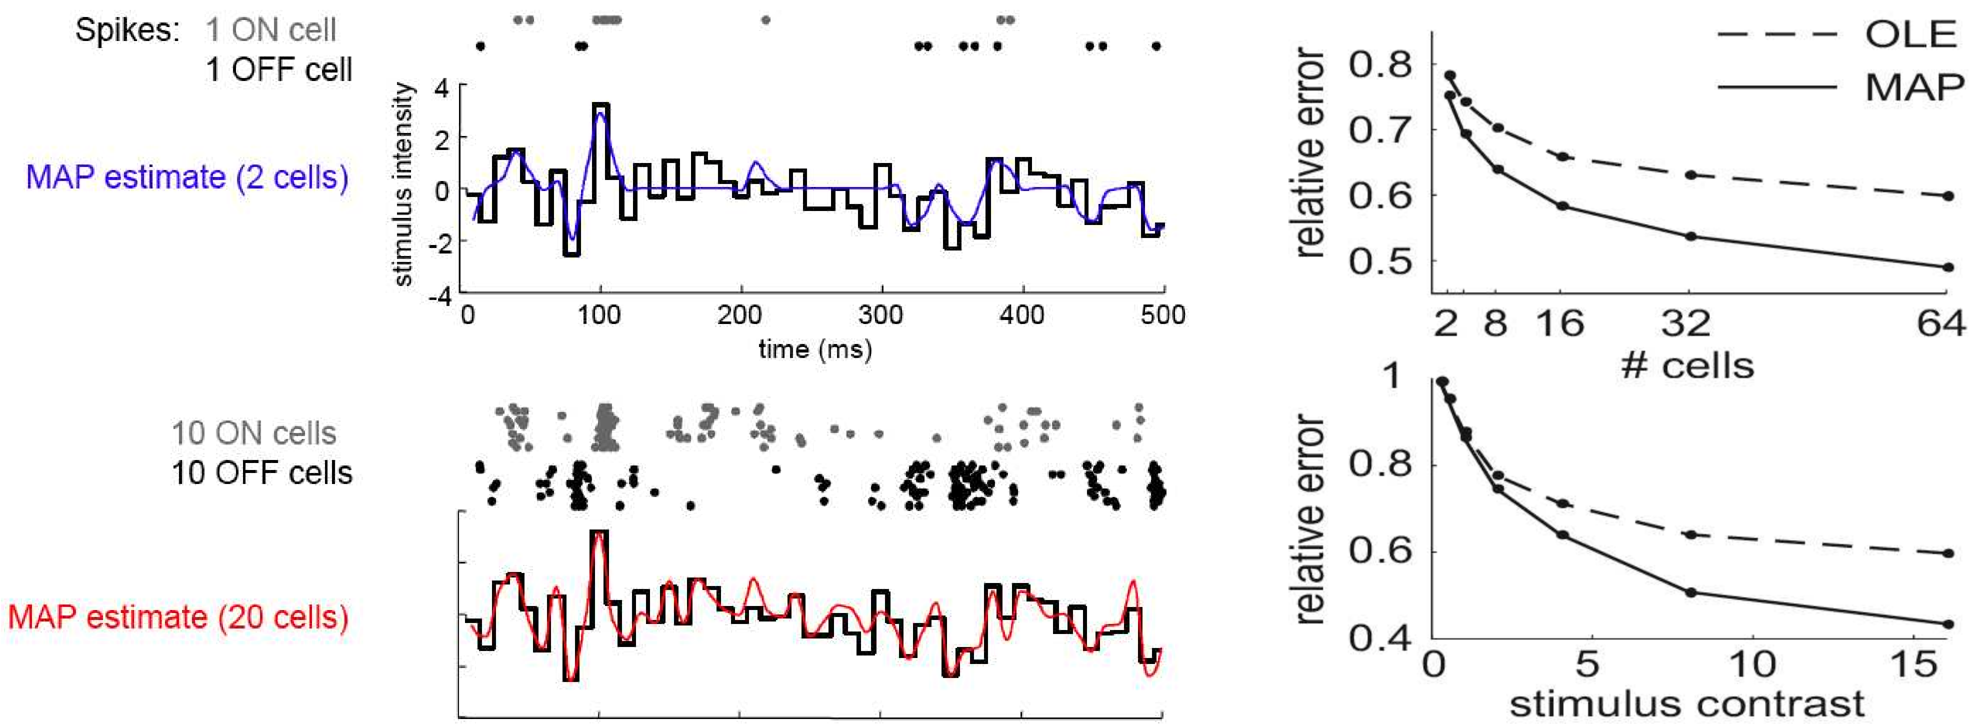

Does including correlations improve decoding?

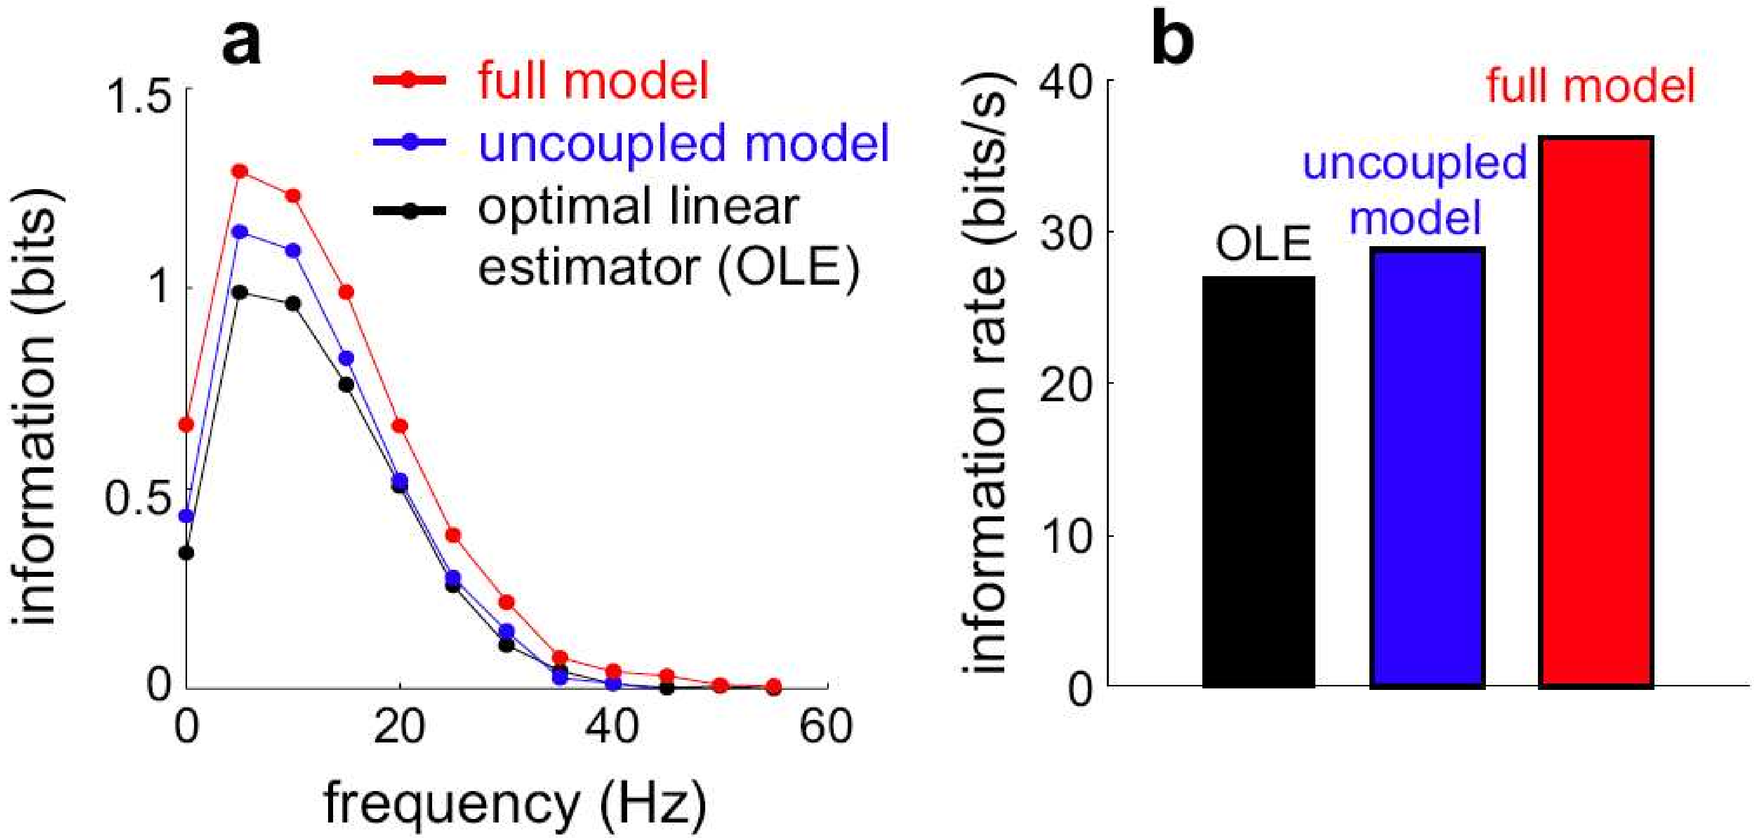

— Including correlations improves decoding accuracy.

How important is timing?

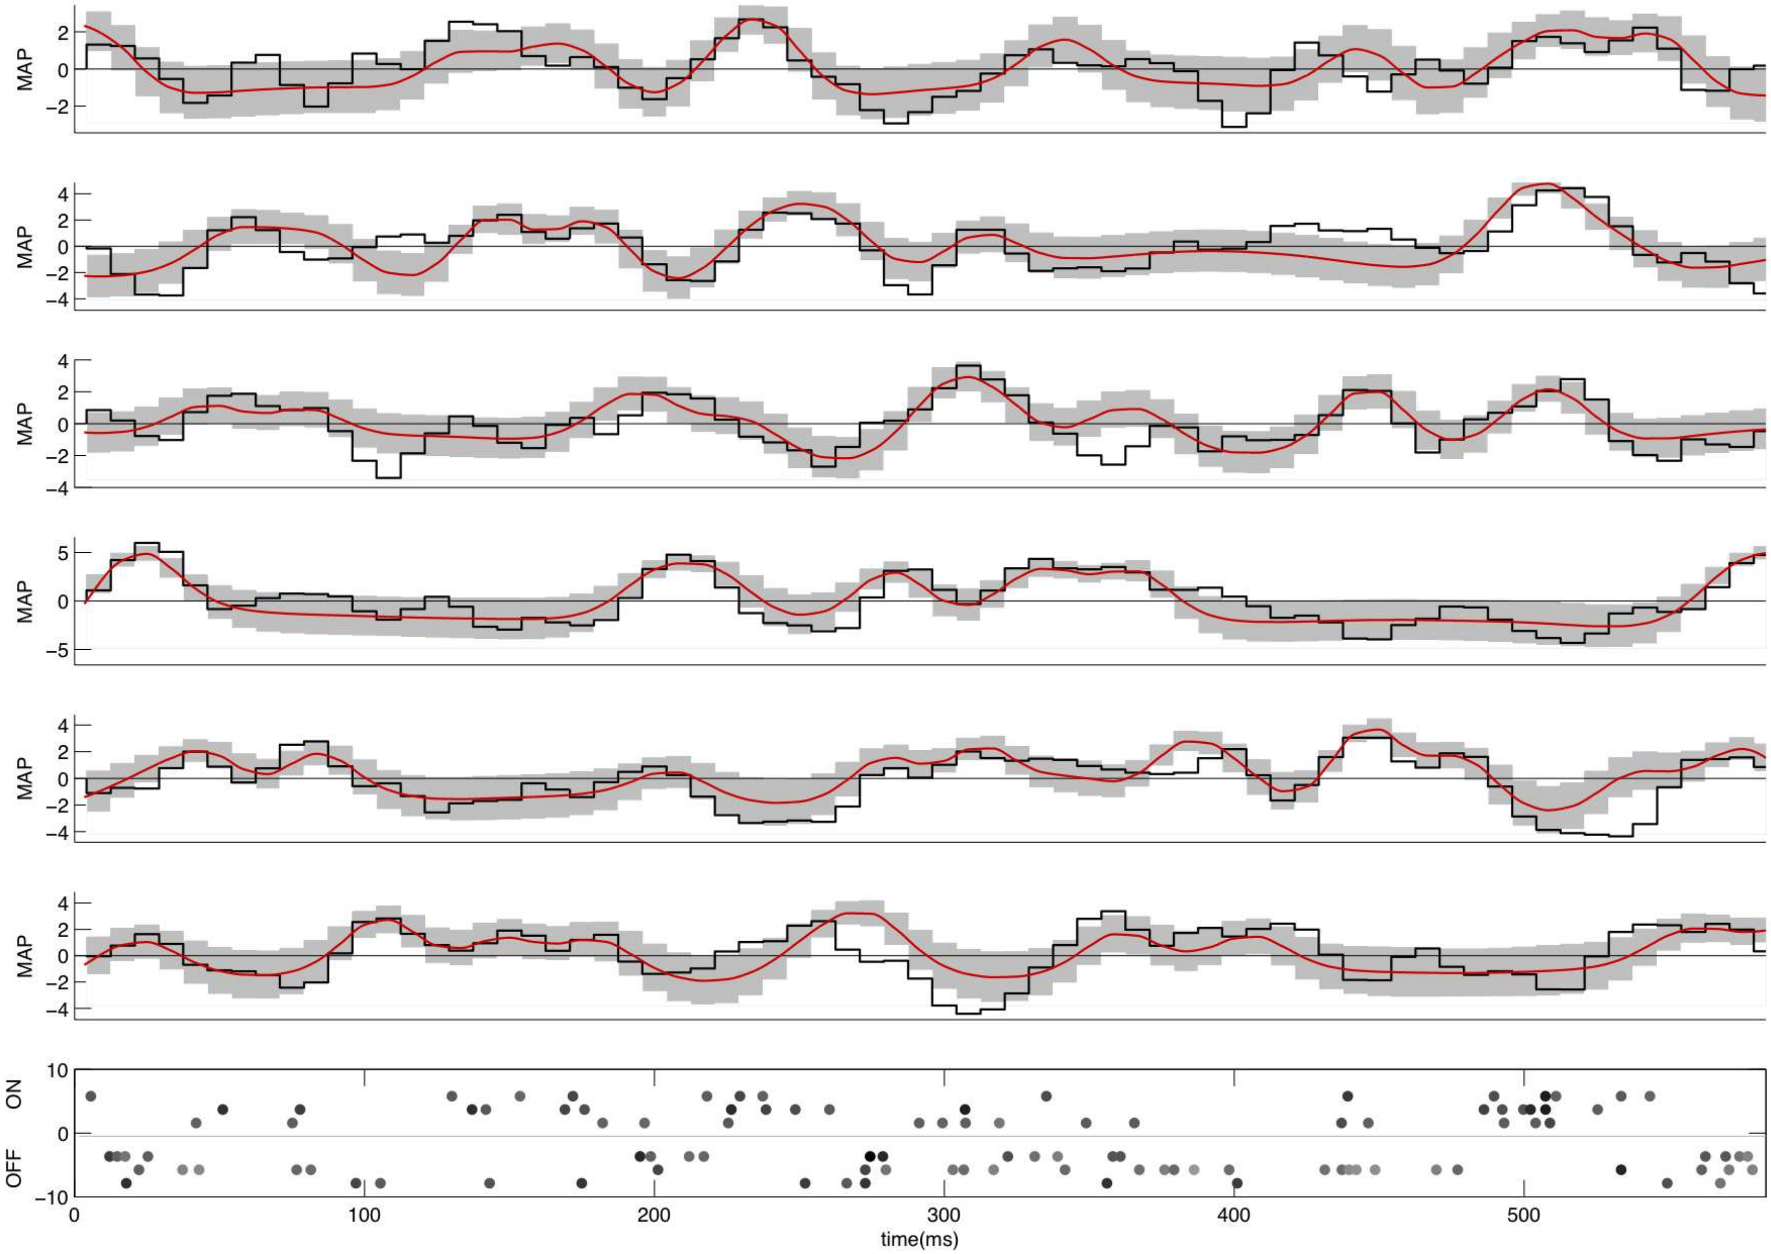

(Ahmadian et al., 2008)

Constructing a metric between spike trains

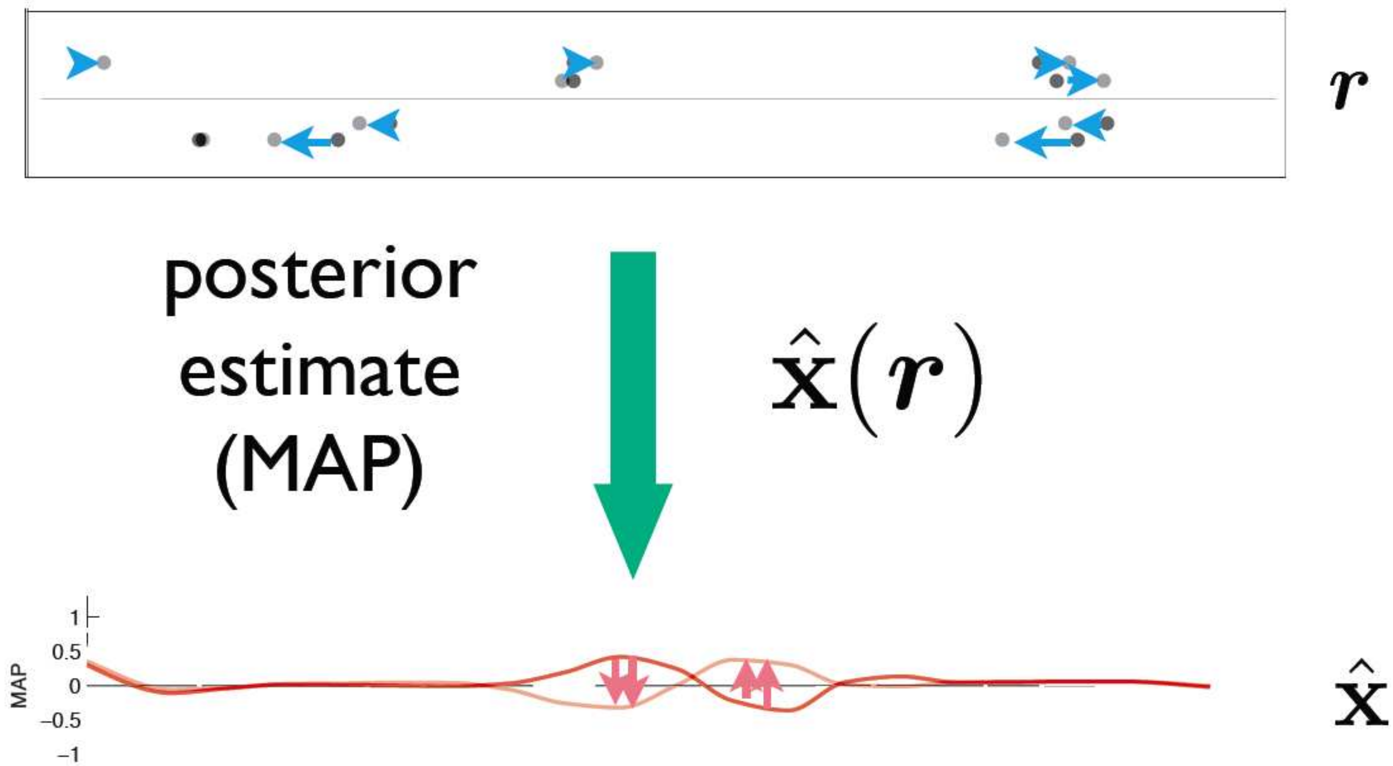

$$d(r_1, r_2) \equiv d_x(x_1, x_2)$$

Locally, $d(r, r + \delta r) = \delta r^T G_r \delta r$: interesting information in G_r.

Effects of jitter on spike trains

Look at degradations as we add Gaussian noise with covariance:

1. $C \propto G^{-1}$ (optimal)
2. $C \propto diag(G)^{-1}$ (perturb less important spikes more)
3. $C \propto I$ (simplest)

Non-correlated perturbations (2,3) are about 2.5$\times$ more costly.

Can also add/remove spikes:
cost of spike addition/deletion $\approx$ cost of jittering by 10 ms.

Optimal velocity decoding

How to decode behaviorally-relevant signals, e.g., image velocity?

If image I is known, use Bayesian estimate (Weiss et al., 2002):

$$p(v|D, I) \propto p(v)p(D|v, I)$$

If image is unknown, we have to integrate out:

$$p(v|D) \propto p(v)p(D|v) = p(v) \int p(I)p(D|v, I)dI;$$

$p(I)$ denotes *a priori* image distribution.

— connections to standard energy models
(Frechette et al., 2005; Lalor et al., 2008)

Optimal velocity decoding

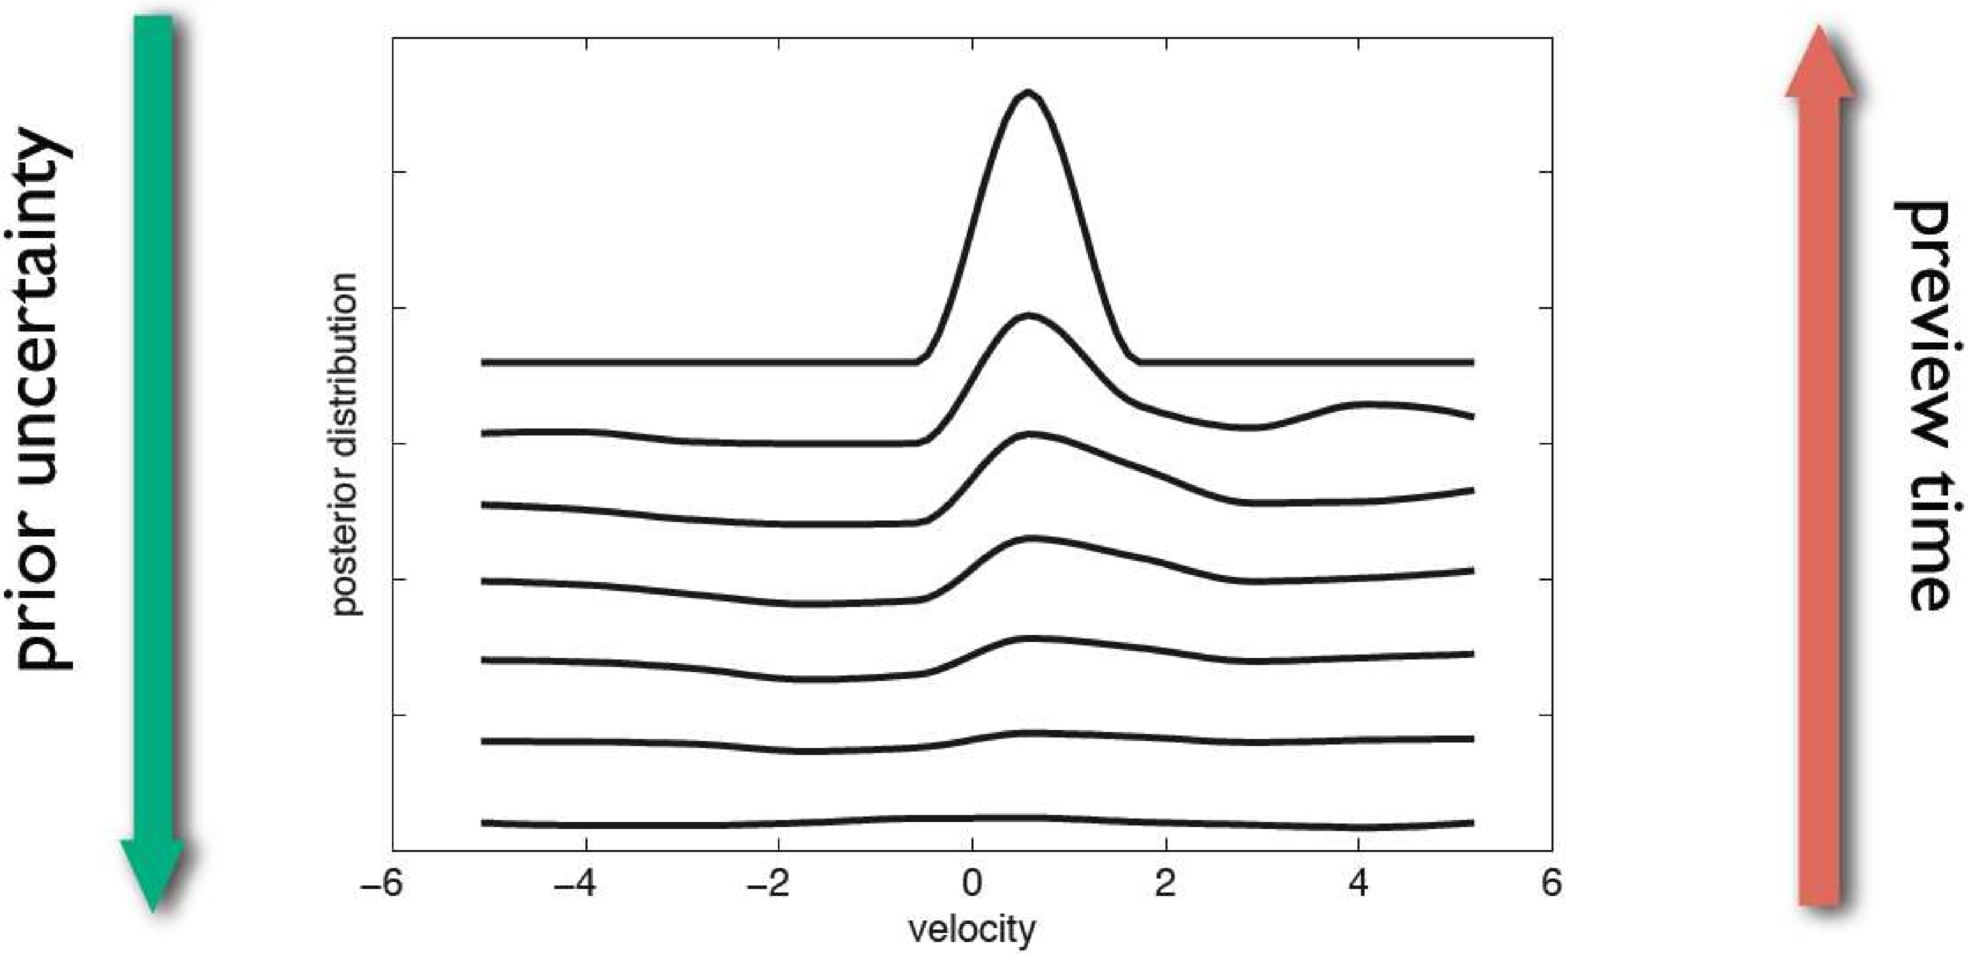

— estimation improves with knowledge of image

Image stabilization is a significant problem

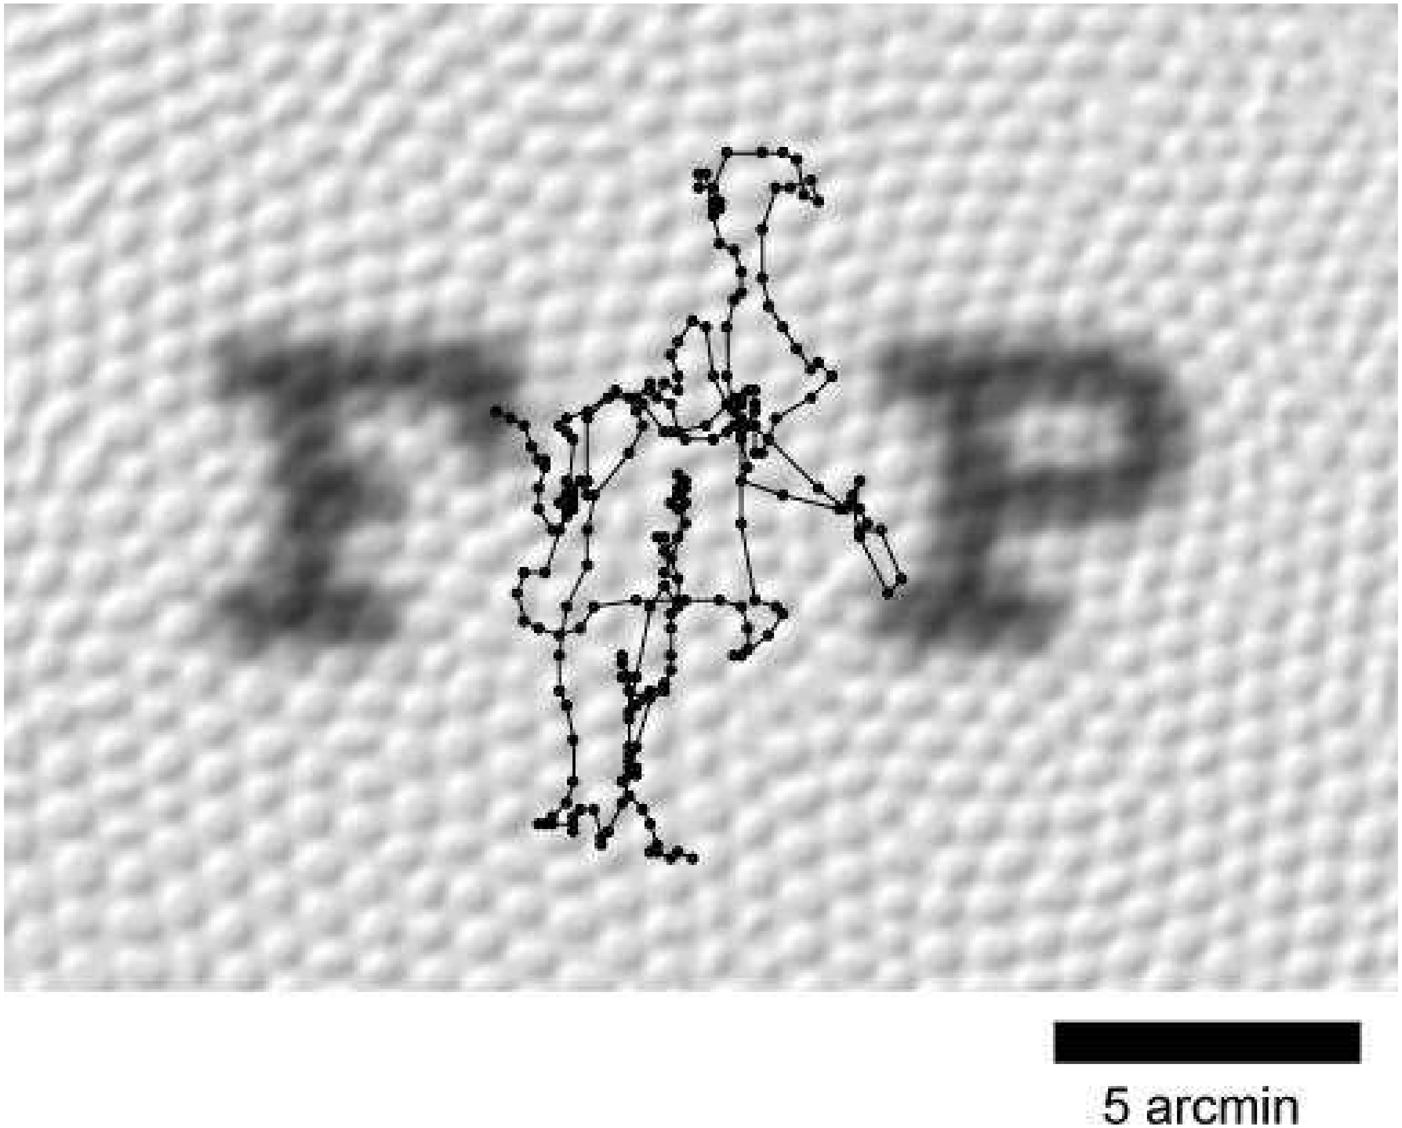

From (Pitkow et al., 2007): neighboring letters on the 20/20 line of the Snellen eye chart. Trace shows 500 ms of eye movement.

Bayesian methods for image stabilization

Similar marginalization idea as in velocity estimation:

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

e denotes eye jitter path.

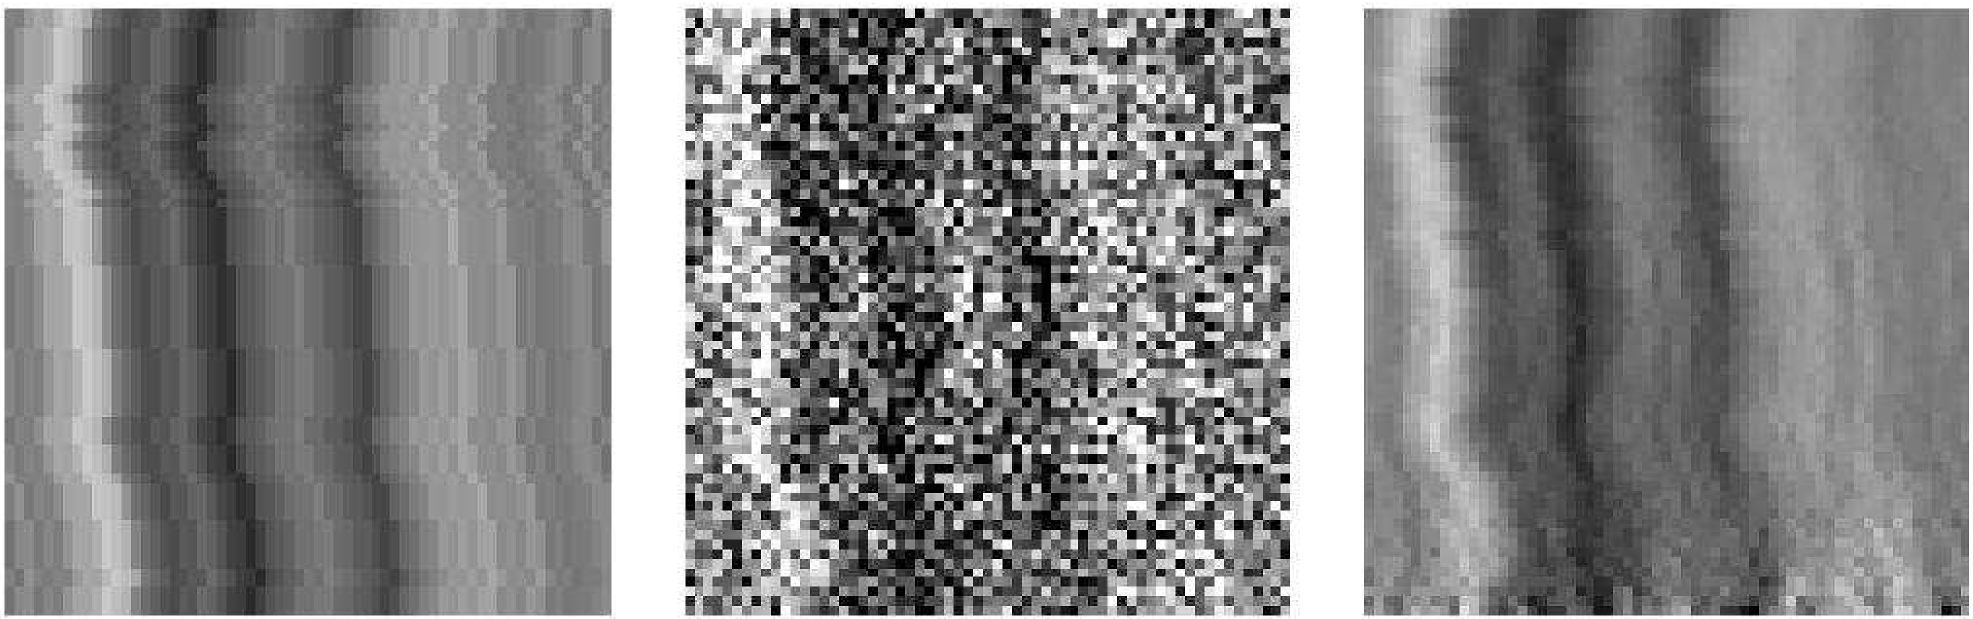

true image w/ translations; observed noisy retinal responses; estimated image.

Collaborators

Theory and numerical methods

- Y. Ahmadian, S. Escola, G. Fudenberg, Q. Huys, J. Kulkarni, M. Nikitchenko, X. Pitkow, K. Rahnama, G. Szirtes, T. Toyoizumi, Columbia
- E. Doi, E. Simoncelli, NYU
- E. Lalor, NKI
- A. Haith, C. Williams, Edinburgh
- M. Ahrens, J. Pillow, M. Sahani, Gatsby
- J. Lewi, Georgia Tech
- J. Vogelstein, Johns Hopkins

Retinal physiology

- E.J. Chichilnisky, J. Shlens, V. Uzzell, Salk

References

Ahmadian, Y., Pillow, J., and Paninski, L. (2008). 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.

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.

Pillow, J., Shlens, J., Paninski, L., Simoncelli, E., and Chichilnisky, E. (2008). Visual information coding in multi-neuronal spike trains. *Nature*, In press.

Pitkow, X., Sompolinsky, H., and Meister, M. (2007). A neural computation for visual acuity in the presence of eye movements. *PLOS Biology*, 5:e331.

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