Statistical challenges and opportunities for reliable CNS interfaces

Liam Paninski

Department of Statistics and Center for Theoretical Neuroscience
Columbia University
http://www.stat.columbia.edu/∼liam
liam@stat.columbia.edu
April 8, 2011

A different perspective: optimal decoding of retinal spike train data

Preparation: macaque retina *in vitro* (Litke et al '04)

— extracellularly-recorded responses of populations of ganglion cells

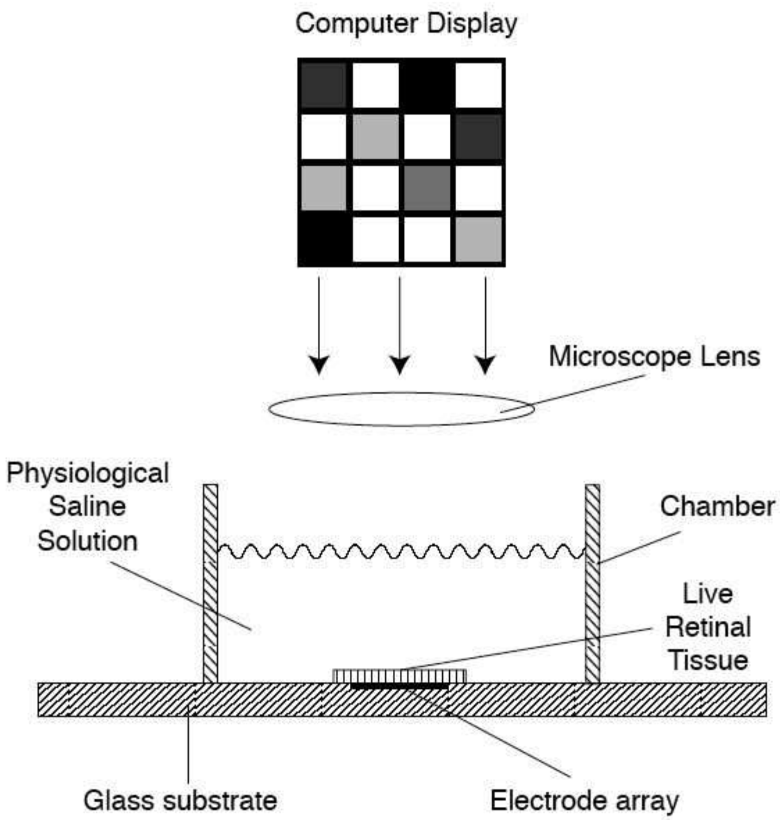

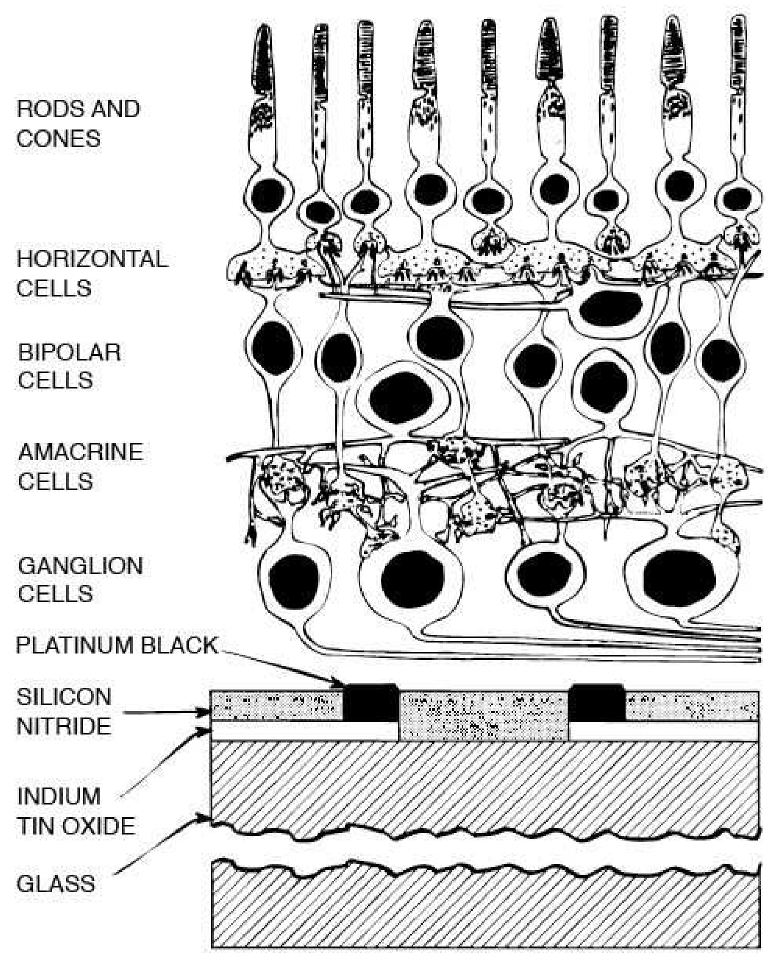

Decoding the spatiotemporally-filtered stimulus via fast Bayesian methods

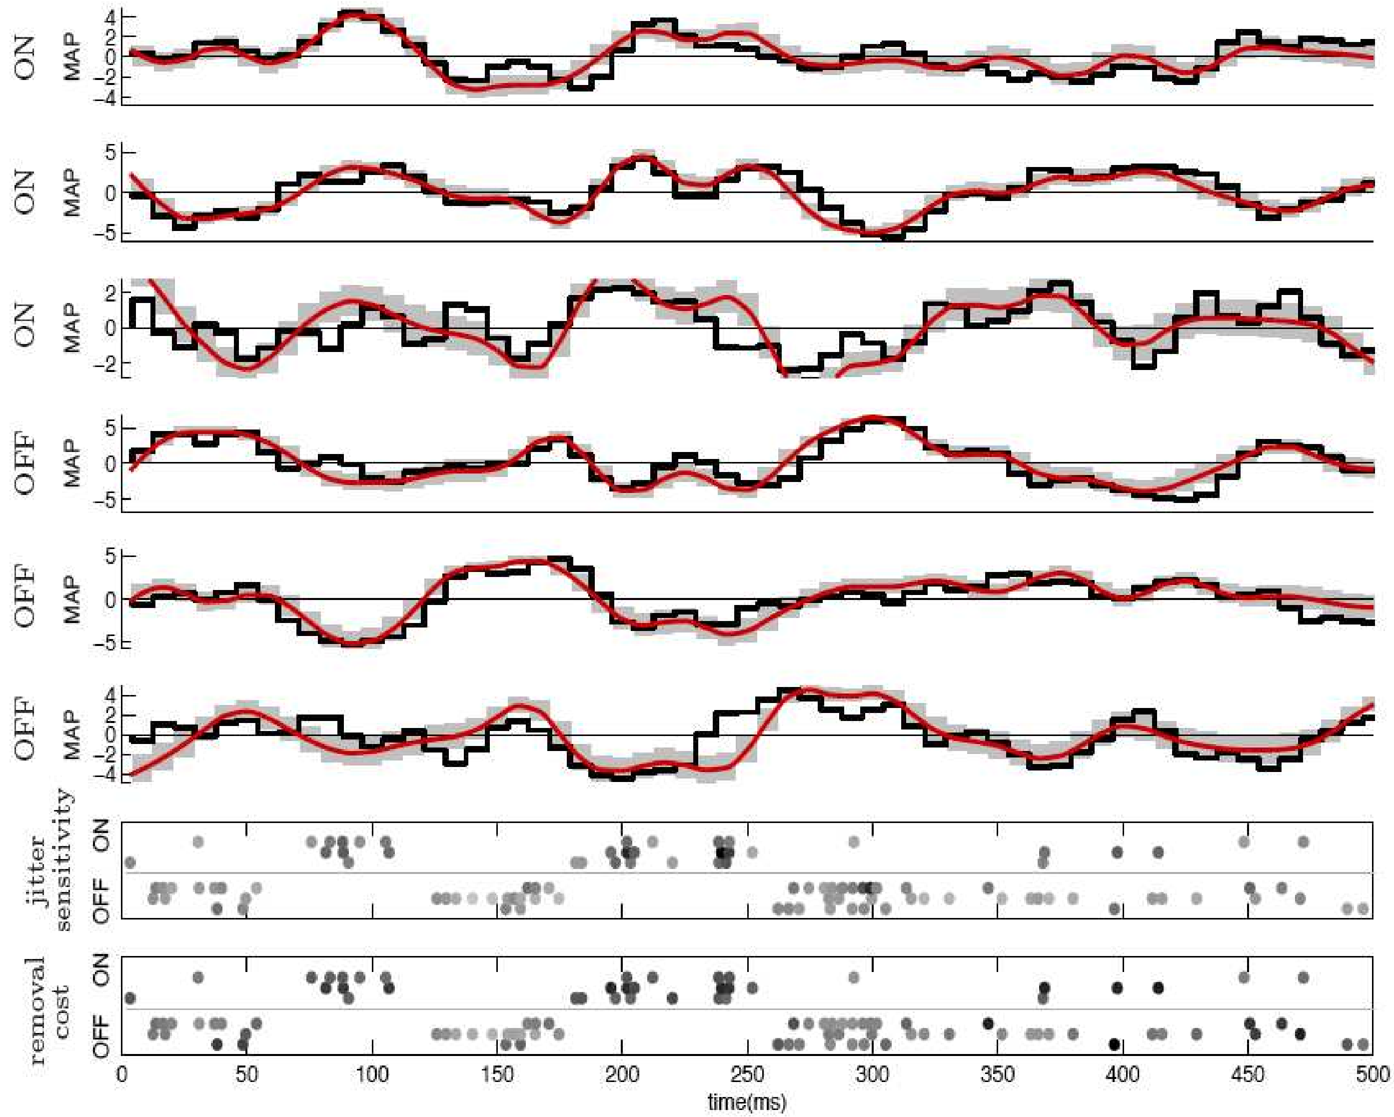

(Ahmadian et al '11); note: motor decoders provide much lower bandwidth.

Outline

- The state of the art: state-space models
- Fast estimation of nonlinear, nonstationary encoding models
- Robust subspace identification
- New methods for optimal Bayesian decoding: fast approximate Kalman-based methods and sequential Markov chain Monte Carlo
- Non-myopic optimal experimental design
- Using all the available information: fast approximate methods for hierarchical regularized models
- Modeling and exploiting co-adaptation.

State-space models

Some notation:

x_t = kinematic state

r_t = neural response

encoding model: $p(r_t|x_t)$

kinematics model: $p(x_t|x_{t-1})$

optimal Bayesian decoder: $p(x_t|r_{1:t})$

This approach is flexible, computationally efficient (because the decoder is computed recursively at each time step t), and currently provides the best available performance.

Challenge: estimating the encoder $p(r_t|x_t)$

— Simple, near-linear models have sufficed so far in 2d planar hand tracking studies; no longer true for $>$ 20-DOF movements (see, e.g., Vargas-Irwin et al '10)

— Nonstationarity is a key, unavoidable issue: the subject adapts to the controller as the controller adapts to the subject. Re-estimating parameters every once in a while is a suboptimal solution.

— Speed is essential, since we have to estimate an encoding model for each observed neural channel.

— Reliability requirements translate into a need to avoid local minima in parameter searches: convex optimization is key tool

New methods for tractably estimating nonlinear models

Simplest example: higher-order terms (e.g., quadratic; Li '09):

$$E(r_t|x_t) = b + k^T x_t + x_t^T A x_t.$$

Problem: number of parameters explodes as we add more terms.

However, we can exploit new convex optimization ideas from the machine learning literature (low-rank tensor completion) to avoid overfitting and obtain quite robust and accurate estimates:

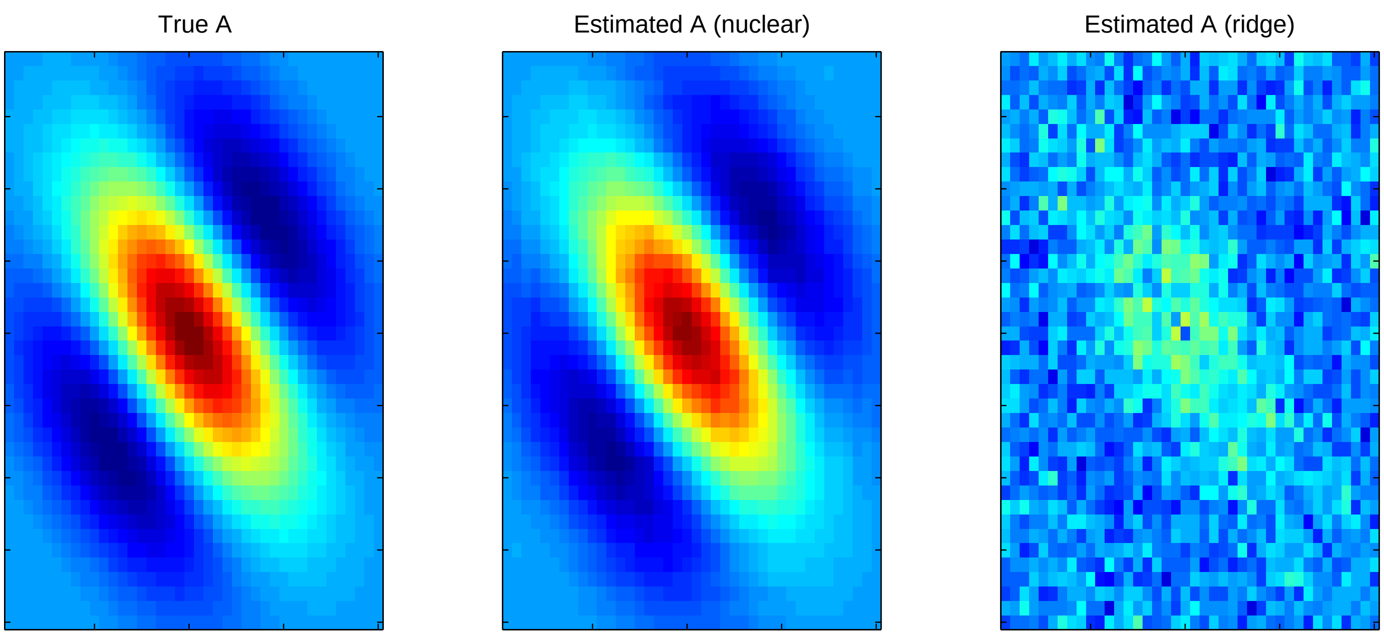

New methods for fast optimal Bayesian nonstationarity tracking

A flexible framework: generalized additive models, $E(r_t) = F\left[\sum_{i=1}^{d} a_i(t) g_i(x_t)\right]$.

As $a_i(t)$ changes, so do the response properties.

Goal: track the a_i's given limited, noisy data.

Fast, robust methods for tracking nonstationarities (optimal Bayesian inference requires just $O(dT)$ time; Paninski et al, '11)

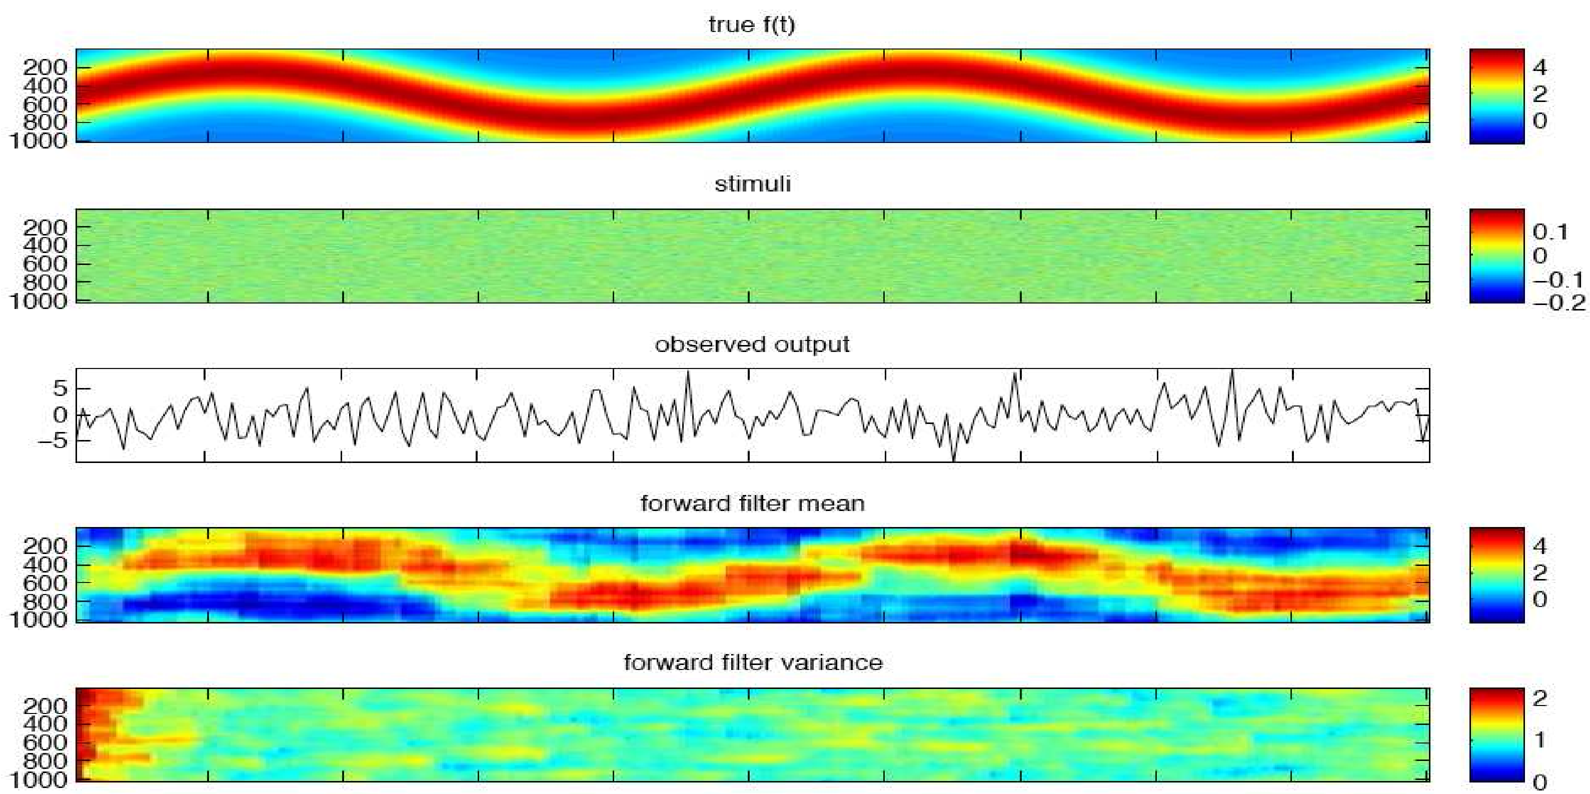

Another challenge: latent variable models

$$E(r_t) = F(x_t, z_t)$$

z_t = latent state variable whose dynamics are restricted to a low-dimensional subspace. Captures the fact that firing rates are modulated by many variables that we don't observe directly (i.e., not just x_t).

Similar models proposed by Yu, Sahani, Shenoy et al and Wu, Paninski et al: significant improvements over models with no z_t term

Problem: estimation of these models is more challenging, since z_t is never observed.

Subspace identification

Previous work applied iterative methods (expectation-maximization) for estimating the latent model parameters.

These methods are slow and non-robust: many iterations, prone to local optima.

Idea: borrow methods from control literature (e.g., Liu and Vandenberghe '09). Use matrix completion again to identify the subspace using limited, noisy data.

Convex problem: no local optima.

Robust decoding

Once model is identified, how to decode?

If $p(x_t|r_{1:t})$ is unimodal, then Kalman-based methods suffice. — we can exploit our new fast $O(d)$ Kalman methods here as well.

If $p(x_t|x_{t-1})$ or $p(r_t|x_t)$ is strongly multimodal, Monte Carlo methods are necessary.

"Particle filtering" is the most common approach; in principle, these methods are very general. However, in practice they are often slow and very non-robust.

Speed issues are solvable: method is embarrassingly parallel at each timestep t, and could be implemented on a GPU.

Robustness issues are more fundamental: standard methods put particles in the wrong place in many cases

Sequential MCMC methods are much more robust than the particle filter

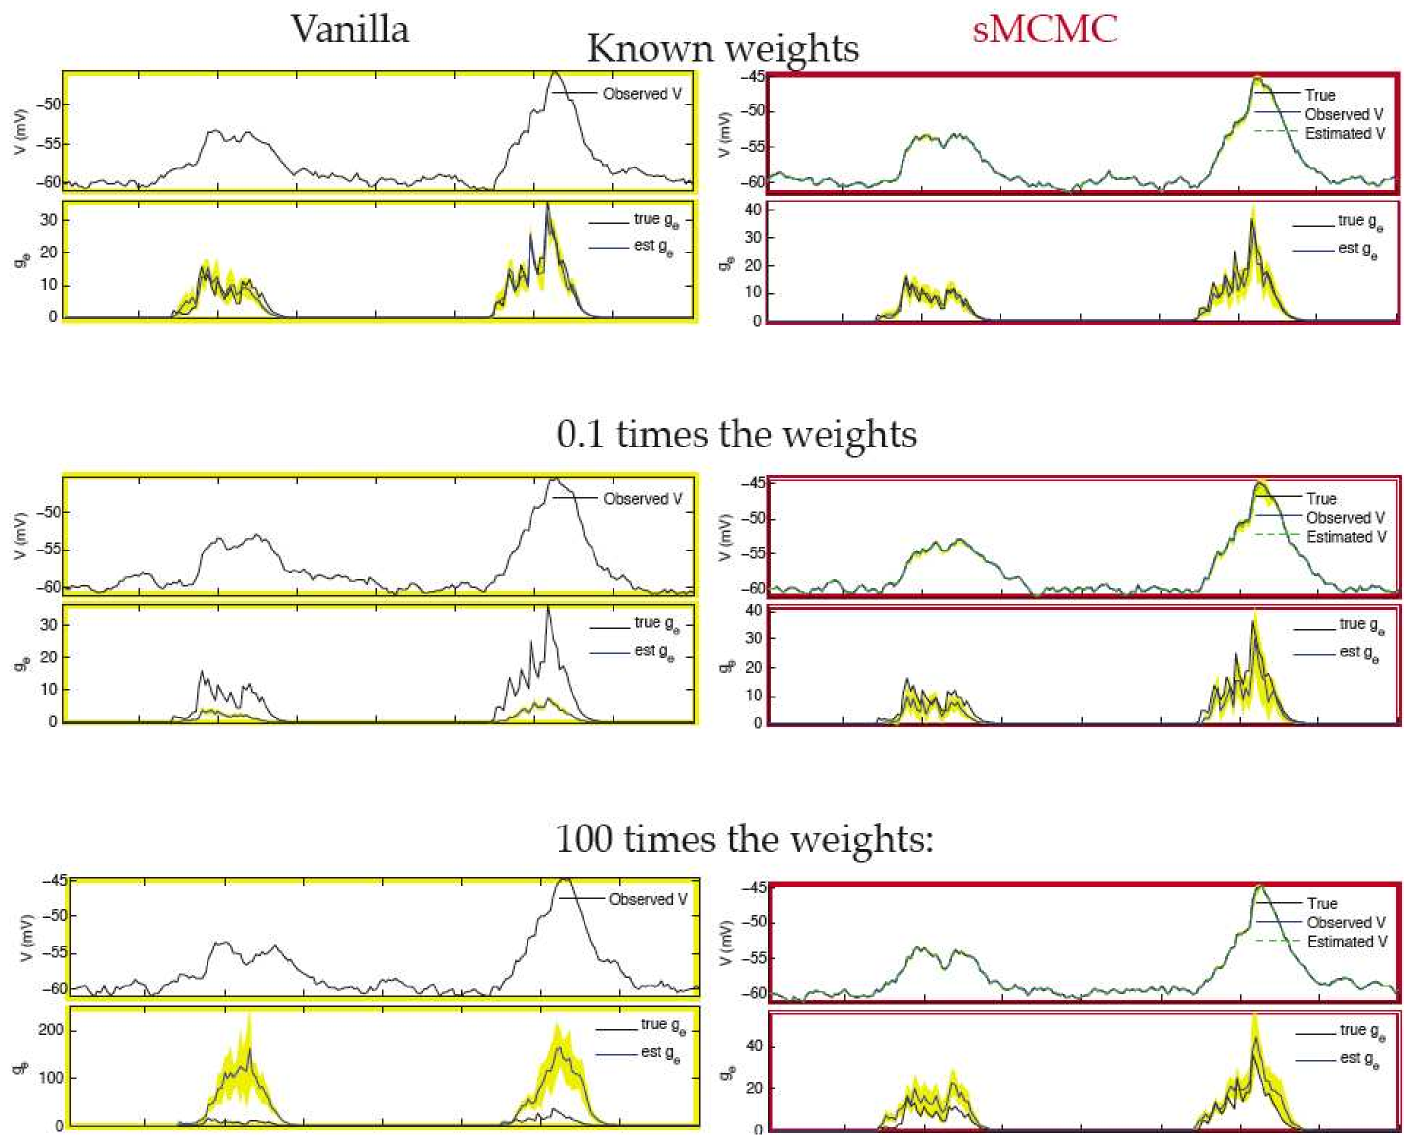

(Vidne and Paninski '11)

Optimal experimental design

Idea: choose test movements to best constrain the model parameters (simple example: Cunningham et al '08).

Previous approaches: choose one test movement at a time, in a greedy (myopic) way (Mackay, '92).

New approach: exploit connection to classical problem in information theory to compute the globally optimal sequence: no greedy local optimization required.

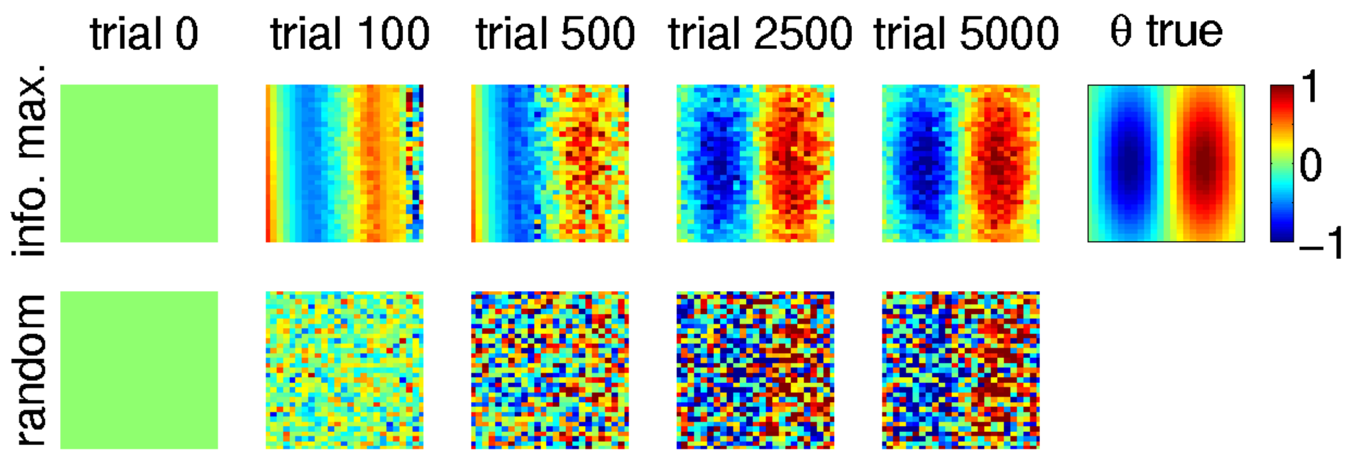

(Lewi et al '09)

Hierarchical models

We record from many units simultaneously, but typically estimate encoding models one unit at a time: this is suboptimal.

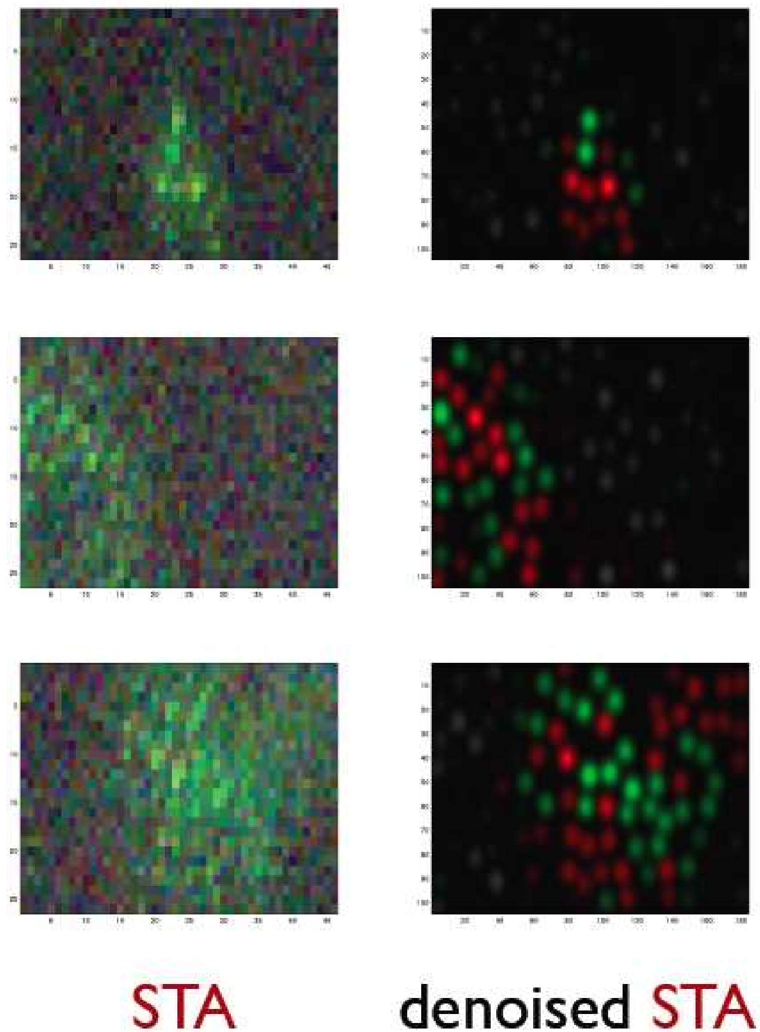

(Field et al, Nature '10; Sadeghi et al, in preparation)

Modeling and exploiting co-adaptation

A number of intriguing results, but no good quantitative models (to my knowledge).

State-space ideas provide a possible starting point.

Simplest case: linear filtering: $\hat{x}_t = \theta_t^T r_t$

Subspace model: $E(r_t) = K u_t$: subject can only influence r_t within a subspace K of dim = DOF.

Co-adaptation: experimentalist optimizes θ_t to optimize accuracy; subject tries to infer θ_t based on recent history, in order to best choose the control signal u_t.

Modeling and exploiting co-adaptation

Natural state-space model: subject tracks mapping $\hat{\theta}_t$ (with uncertainty) via Bayesian updates. Similar models in motor psychophysics literature on adaptation of sensorimotor maps.

Goal: exploit co-adaptation instead of fighting it:

— natural hierarchical model for tracking encoding models

— improved decoder: incorporate subject's uncertainty about θ

— connection to optimal design: the subject will try control signals u_t that balance exploration (to optimize information about θ) and minimization of error. Qualitatively consistent with nonstationarities during tracking (e.g., results from Carmena, Schwartz labs).

Thanks!

Joint work with E. Pnevmatikakis, K. Sadeghi, M. Vidne, J. Shi, Y. Ahmadian, K. Rahnama Rad, J. Huggins, J. Lewi, E. Doi

Support: Sloan Fellowship, NSF/NIH collaborative research in computational neuroscience, NSF CAREER, McKnight Scholar award.