Detecting changes in neural dynamics within single trials

Biljana Petreska[1], Byron M. Yu[1,2,4], John P. Cunningham[3,4], Gopal Santhanam[4], Stephen I. Ryu[4], Krishna V. Shenoy[4] & Maneesh Sahani[1]
[1]Gatsby Computational Neuroscience Unit, UCL, [2]Carnegie Mellon University, [3]University of Cambridge & [4]Stanford University

Motivation

Extract from recorded spike-trains the low-dimensional nonlinear dynamical rules that drive single-trial network activity.

Framework

- Network activity is organised around the (relatively) low-dimensional evolution of a set of latent order parameters or computational state.
- The trajectory taken through this computational state space may differ on each trial.
- Both the dynamics in the computational state, and its relationship to neural firing, may be nonlinear.

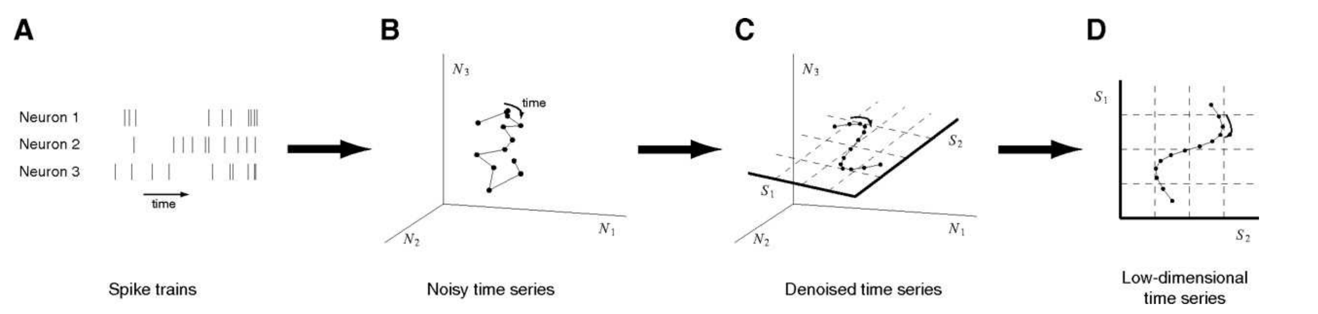

We attempt to recover noisy low-dimensional dynamics reflected in the recorded data.

Task

Delayed reach in 7 directions: 0–315°; 2 distances: 60 and 100 mm; various delays: 200–900 ms; 52 trials per direction.

Recordings made with a 96-channel silicon electrode array.

Smoothed latent trajectories

Gaussian Process Factor Analysis (GPFA) [Yu et al., J Neurophysiol, 2009] visualises low-dimensional trajectories that capture network evolution. Assumes:

- linear evolution of state (and linear embedding)
- non-Markov dynamics (so GPFA state $\neq$ phase state)
- independent latent dimensions

Dynamical model

A **Hidden Switching Linear Dynamical Systems (HSLDS) model** provides nonlinear, Markov, coupled dynamics in the latent space.

Transition dynamics: p-dimensional state $x_t \in \mathbb{R}^{p\times 1}$ evolves according to one of S linear dynamical laws with Gaussian innovations:

$$x_t = A_{s_t} x_{t-1} + \eta^x \qquad \eta^x \sim \mathcal{N}(0, Q_{s_t})$$

with both drift A_{s_t} and innovations covariance matrix Q_{s_t} indexed by $s_t \in \{1, ..., S\}$.

Switching dynamics: s_t is Markov with transition matrix T that is learned from the data:

$$s_t \sim \text{Discrete}(T_{\cdot, s_{t-1}}) \, .$$

This drives switching between the different dynamical laws.

The trial is thus divided into segments, with latent dynamics in each segment following a different linear dynamical law. Switching between the available linear systems provides a piecewise approximation to nonlinear network dynamics, as well as capturing changes in the dynamics.

Observation process. As in GPFA, the neural state x_t is related to the recorded spikes $y_t \in \mathbb{R}^{q\times 1}$ $(q > p)$ through a linear-Gaussian relationship:

$$y_t = C x_t + \eta^y \qquad \eta^y \sim \mathcal{N}(d, R)$$

with observation matrix C and a bias term d independent of s_t. R is diagonal, compelling shared variability between neurons to be modelled by the latent process.

Inference and learning. Inference in the HSLDS is intractable and can be approximated with a forward and a backward pass using Assumed Density Filtering and Expectation Correction [Barber, JMLR, 2006]. This leads to an approximate EM algorithm for learning.

Modelling details. Spikes were binned at 10 ms. The total number of recorded neurons was $q = 105$. HSLDS models were fit to subsets of trials; only neurons that fired at least once during the trials in question were modelled. The linear dynamics most likely to be active was computed using the Viterbi algorithm.

Results

Detection of Behavioral Epochs. The HSLDS models nonlinear dynamics by identifying multiple different regimes within each trial. Some of the switches between linear regimes correlate well with the timing of behavioural events on the corresponding trial.

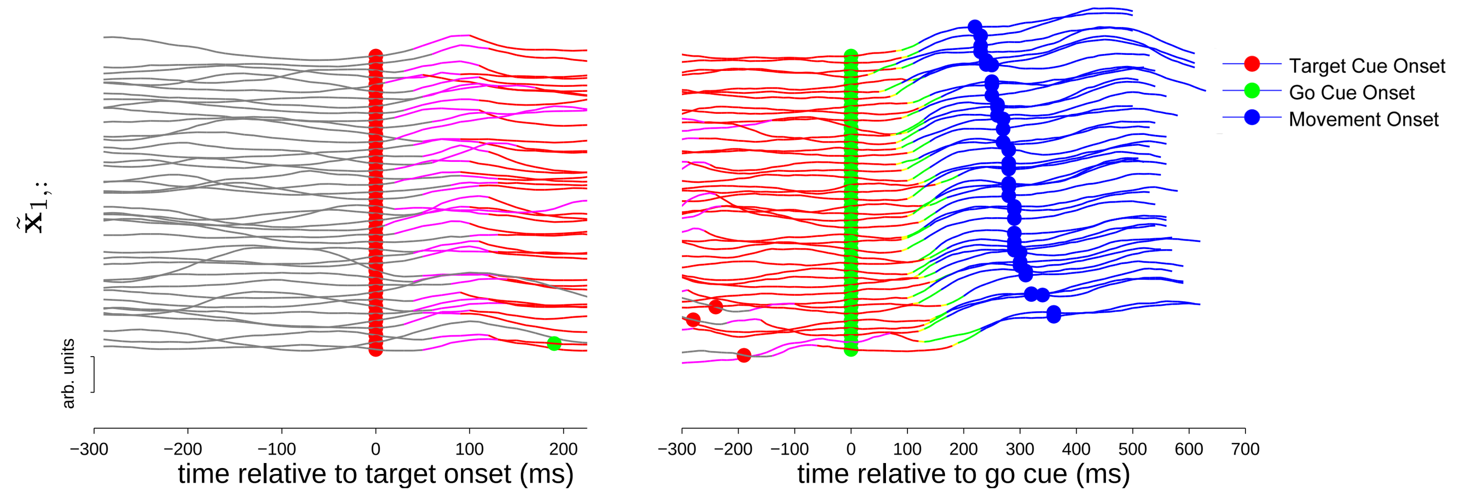

First latent dimension $\tilde{x}_{1,:}$ of the orthonormalised neural trajectories found by HSLDS (one movement direction, p = 7, S = 7), trials sorted by reaction time and aligned to behavioural events. Switches follow the target onset (left) and precede the movement onset (right) reliably. Each colour represents the use of a different linear dynamical system.

Modelling multiple directions. Sufficiently different movement directions are associated with different dynamics.

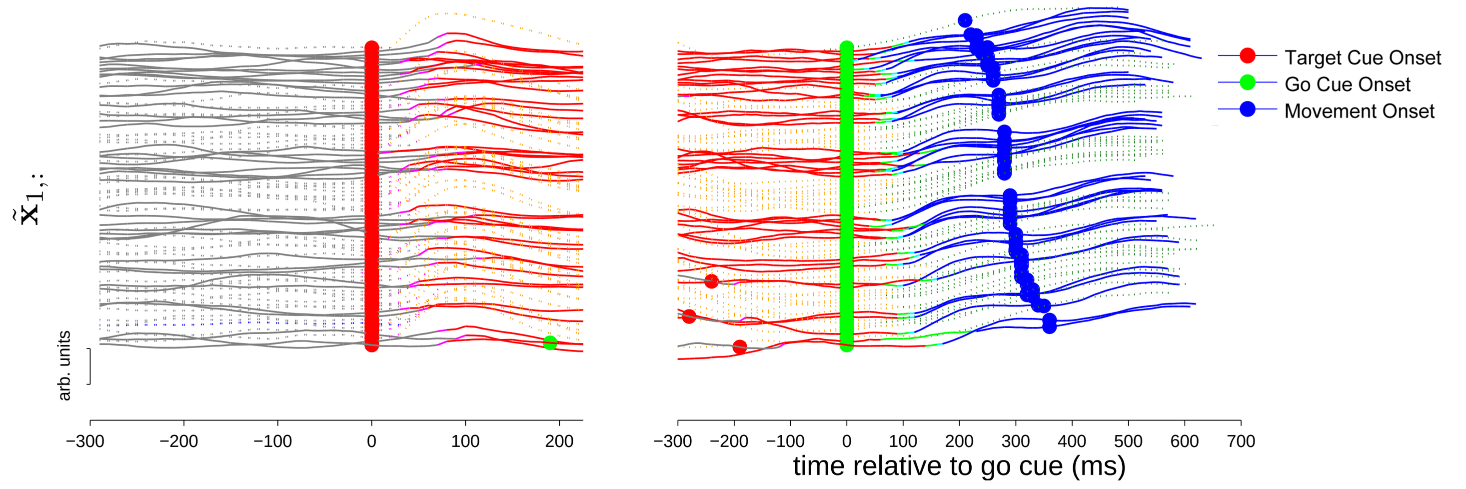

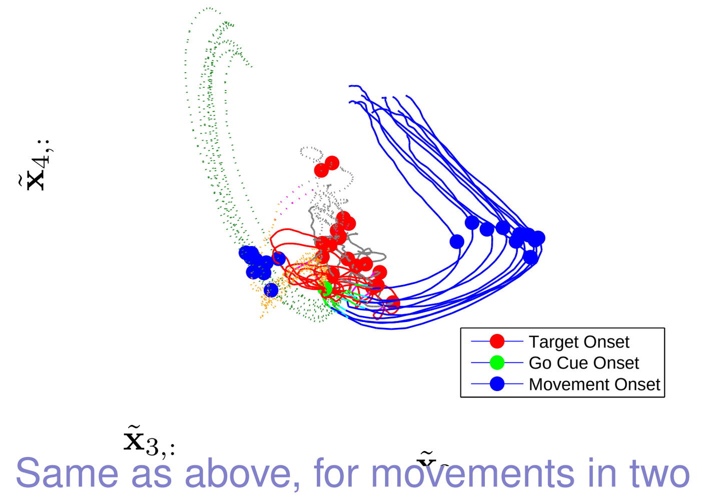

Same as above, for movements in two different directions (full vs dotted line). Note that the first (pre-delay) state is the same for both movement types.

Validation of the Model. Latent models capture shared variance. Thus, the firing rate of one neuron can be predicted using measurements of the rest. HSLDS models capture shared structure better than GPFA or a single hidden LDS.

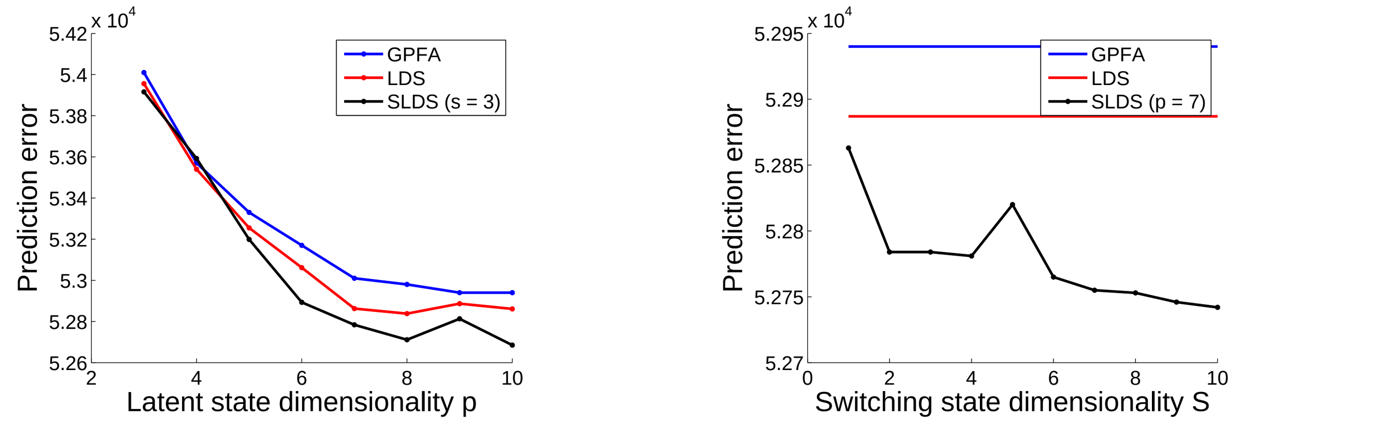

The HSLDS model outperforms GPFA and a single LDS at all latent dimensions. More linear laws appear to be better.

Relationship between HSLDS switches and behavioural events. The time of first switch following the go cue depends on reaction time, and interpolates between the sensory go-signal and the onset of movement.

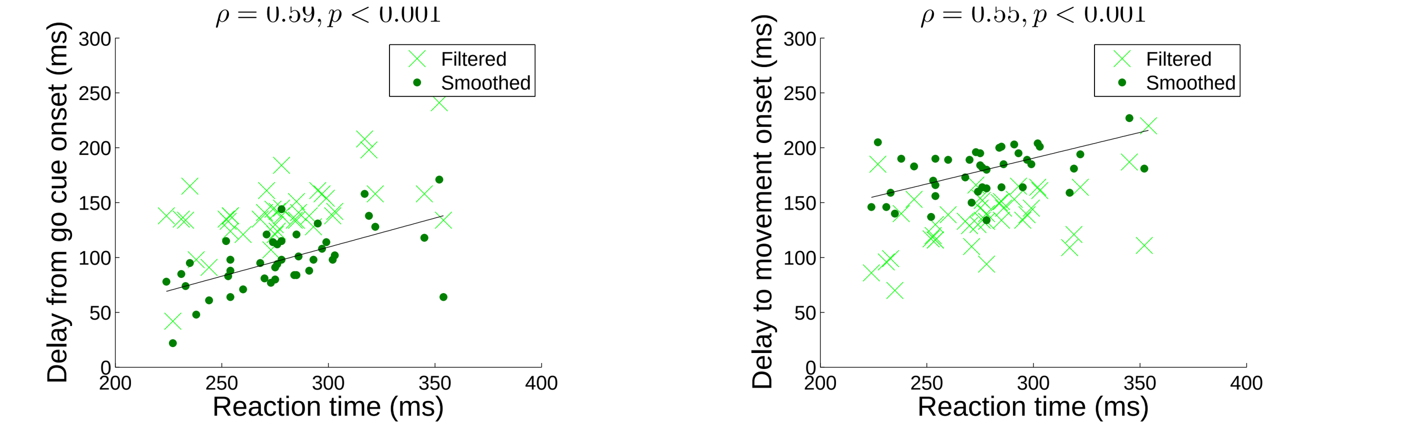

HSLDS switches correlate both with time of preceding go-signal, and time of subsequent movement.

Discussion

HSLDS models successfully capture the neural dynamics in populations of recorded neurons and do so within single trials. We show that HSLDS models: (1) detect behavioural-related changes and nonlinearities in the neural dynamics, (2) attribute different dynamics to movements in different directions and (3) perform better than other dimensionality reduction and denoising methods.

Support: Gatsby Charitable Foundation; DARPA REPAIR (N66001-10-C-2010); NIH CRCNS R01-NS054283; NIH Pioneer 1DP1OD006409.