allow us to identify three stereotypical linear filter types other than the STA which modulate the firing of most cells in stereotyped ways, shedding some light into visual features other than the classical receptive field which influence ganglion cell responses.

doi:

I-33. Detecting changes in neural dynamics within single trials

Biljana Petreska[1] BILJANA.PETRESKA@A3.EPFL.CH
Byron M. Yu[2] BYRONYU@CMU.EDU
John P. Cunningham[3] JCUNNIN@STANFORD.EDU
Gopal Santhanam[4] GOPAL@NERUR.COM
Stephen I. Ryu[4] SEOULMAN@STANFORD.EDU
Krishna V. Shenoy[4] SHENOY@STANFORD.EDU
Maneesh Sahani[1] MANEESH@GATSBY.UCL.AC.UK

[1]University College London
[2]Carnegie Mellon University
[3]University of Cambridge
[4]Stanford University

We hypothesize that the high-dimensional spiking activity recorded in the brain is driven by a smooth low-dimensional process (or "neural state") which reflects the dynamical evolution of the underlying network. In the context of movement for example, spikes recorded in motor-related areas reflect low-dimensional control signals that drive the dynamics and kinematics of the hand. Gaussian Process Factor Analysis (GPFA) is a method that effectively captures this neural state through simultaneous denoising and dimensionality reduction, facilitating the visualization and analysis of multineuron recordings (Yu et al., J Neurophysiol, 2009). Here, we extend the latent model beyond that of GPFA, introducing explicit and non-Gaussian dynamics in the form of a switching linear dynamical system (SLDS). The SLDS state evolves according to one of a set of different linear-Gaussian dynamical laws, switching between these laws both to reflect true changes in the underlying network dynamics and to approximate any dynamical non-linearities. The model is identified by an approximate Expectation-Maximisation algorithm using Assumed Density Filtering in the E-step (Barber, JMLR, 2006). We investigated how well the SDLS model captured the dynamics of firing in a population of 104 neurons recorded in the monkey premotor and motor cortices during a delayed-reach task. Using the cross-prediction approach of Yu et al., we found that the SLDS model described the spiking data better than did GPFA for all latent dimensionalities studied. Some switches in the latent dynamics reflected non-linear approximations. However some switches were reliably correlated with trial-by-trial behavioural events, with temporal lags appropriate for causality, even when the model was learned without supervision from the spiking data alone. Thus, SLDS models appear to successfully capture neural dynamics, and behaviourally-related changes in those dynamics, in population recordings.

doi:

I-34. Modelling low-dimensional dynamics in recorded spiking populations

Jakob Macke[1] JAKOB.MACKE@GMAIL.COM
Lars Büsing[1] LARS@GATSBY.UCL.AC.UK
John P. Cunningham[2] JCUNNIN@STANFORD.EDU
Byron M. Yu[3] BYRONYU@CMU.EDU
Krishna V. Shenoy[4] SHENOY@STANFORD.EDU
Maneesh Sahani[1] MANEESH@GATSBY.UCL.AC.UK

[1]University College London
[2]University of Cambridge

[3]Carnegie Mellon University
[4]Stanford University

Neural population activity reflects not only variations in stimulus drive (captured by many neural encoding models) but also the rich computational dynamics of recurrent neural circuitry. Identifying this dynamical structure, and relating it to external stimuli and behavioural events, is a crucial step towards understanding neural computation. One data-driven approach is to fit hidden low-dimensional dynamical systems models to the high-dimensional spiking observations collected by microelectrode arrays (Yu et al, 2006, 2009). This approach yields low-dimensional representations of population-activity, allowing analysis and visualization of population dynamics with single trial resolution. Here, we compare two models using latent linear dynamics, with the dependence of spiking observations on the dynamical state being either linear with Gaussian observations (GaussLDS), or generalised linear with Poisson observations and an exponential nonlinearity (PoissonLDS) (Kulkarni & Paninski, 2007). Both models were fit by Expectation-Maximisation to multi-electrode recordings from pre-motor cortex in behaving monkeys during the delay-period of a delayed reach task. We evaluated the accuracy of different approximations for the E-step necessary for PoissonLDS using elliptical slice sampling. We quantified model-performance using a cross-prediction approach (Yu et al). Although only the Poisson noise model takes the discrete nature of spiking into account, we found no consistent improvement of the Poisson-model over GaussLDS: PoissonLDS was generally more accurate for low dimensions, but slightly under-performed GaussLDS in higher dimensions (cf. Lawhern et al. 2010). We also examined the ability of such models to capture conventional population metrics such as pairwise correlations and the distribution of synchronous spikes counts. We found that both models were able to reproduce these quantities with very low dynamical dimension, although the non-positivity of the Gaussian model introduced a bias. Thus, despite its verisimilitude, the Poisson observation model does not always yield more accurate predictions in real data.

doi:

I-35. Reservoir dynamics: Feedback and chaos in the network solution of a complex cognitive task

Loic Matthey LOIC.MATTHEY@GATSBY.UCL.AC.UK
Peter Dayan DAYAN@GATSBY.UCL.AC.UK

University College London

Amorphous, non-linear, recurrent microcircuits have been shown to retain and combine information about their input trajectories, and thus provide a rich temporal basis for representing complex functions (Jaeger 2004; Maass 2004; Abbott 2009). However, the effect of key parameters, including their degree of chaos and feedback strength, is incompletely understood. Further, visualization of their state trajectories is as primitive as for experimental observations of the activities of multiple real neurons. We trained a collection of reservoir networks in a simplified version of the popular 12AX task involving a precisely calibrated requirement for working memory, for which we could establish an unambiguous notion of a high-level network state. The networks had to remember the past appearance of a specific symbol and ignore intervening random symbols. All networks comprised 500 neuron-like rate units, and had sparse recurrent connections. They differed according to the degree of intrinsic chaos (controlled by the scaling of the recurrent weights) and the feedback strength. Even the weakest network could retain information for around 10 sec, despite its 10ms membrane time constant. Non-chaotic networks with strong feedback retained information nearly indefinitely. With weaker feedback, memory decayed roughly linearly with time. Chaotic networks, even with weak feedback, transiently held perfect information, before decaying again linearly. We used dimensionality reduction methods to examine and visualize the overall state of the networks. PCA analysis of network activity revealed clear clusters in the dynamics, associated with the structure in the underlying task. The relative positions of the clusters indicated different dependencies on inputs and feedback, and provided hints as to the differing computational strategies of the networks. More complex tasks, such as 12AX, which impose hierarchical demands on memory and computation will provide a flexible and stringent testbed for determining the power and modus operandi of such networks.