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