III-6. Precise estimates of single-trial neural population state in motor cortex via deep learning methods

Chethan Pandarinath[1,2] CPANDAR@EMORY.EDU
Jasmine Collins[3] JLCOLLINS@GOOGLE.COM
Rafal Jozefowicz[3] RAFAL@OPENAI.COM
Sergey Stavisky[4] SERGEY.STAVISKY@STANFORD.EDU
Jonathan Kao[4] JCYKAO@STANFORD.EDU
Mark Churchland[5] MC3502@CUMC.COLUMBIA.EDU
Matthew Kaufman[6] MKAUFMAN@CSHL.EDU
Stephen Ryu SEOULMAN@STANFORD.EDU
Jaimie Henderson[4] HENDERJ@STANFORD.EDU
Krishna Shenoy[4] SHENOY@STANFORD.EDU
Laurence Abbott[5] LFA2103@CUMC.COLUMBIA.EDU
David Sussillo[3,4] SUSSILLO@GOOGLE.COM

[1]Emory University
[2]Georgia Institute of Technology
[3]Google Brain
[4]Stanford University
[5]Columbia University
[6]Cold Spring Harbor Laboratory

Neuroscience is experiencing a data revolution in which many hundreds or thousands of neurons are recorded simultaneously. This often reveals structure in the population activity that is not apparent from single neuron responses. However, understanding this structure on a single-trial basis is often challenging due to limited observations of the neural population, trial-to-trial variability, and the inherent noise of action potential arrival times. Here we introduce Latent Factor Analysis via Dynamical Systems (LFADS), a deep-learning method to infer latent dynamics from simultaneously recorded, single-trial, high-dimensional neural spiking data. LFADS is a sequential model based on a variational auto-encoder (Kingma & Welling, 2013). By making a dynamical systems hypothesis regarding the generation of the observed data, LFADS reduces observed spiking to a set of low-dimensional temporal factors, per-trial initial conditions, and inferred inputs. Here we apply LFADS to a variety of datasets from monkey motor cortex. We show that LFADS's estimates of neural population state are more informative about behavioral variables than population activity itself. In addition, LFADS uncovers multiple known dynamic features of single-trial motor cortical firing rates, including slow oscillations (1-3Hz) that accompany the transition from pre- to peri-movement activity (Churchland et al., Nature 2012), and high-frequency oscillations (15-45 Hz) that occur during the pre-movement period (Donoghue et al., J Neurophys 1998). In cases where the neural data's dynamics cannot be modeled by an initial state alone (e.g., unexpected perturbations), LFADS infers time-varying external inputs that correlate with behavioral outcomes. Finally, we apply LFADS to an unstructured dataset (no precise timing, free-paced reaching movements, no repeated conditions) and show that it uncovers precise state estimates and inputs from unstructured activity. These results showcase the ability of LFADS to infer precise estimates of single-trial dynamics on multiple timescales and uncover inputs that correlate with behavioral choices.

III-7. Sequence replay in model neural networks without Hebbian plasticity

Richard Pang RPANG@UW.EDU
Adrienne Fairhall FAIRHALL@U.WASHINGTON.EDU

University of Washington

The reactivation of neuronal activity patterns outside the context in which they originally occurred is thought to play a role in mediating memory, but the mechanisms underlying reactivation are not well understood. Especially