(BCI), which allowed us to directly specify which co-modulation patterns we would like the subject to show. We implanted a monkey with a 96-channel electrode array in the primary motor cortex. The subject modulated his neural activity to control the velocity of a cursor on a center-out BCI task. Each session, the subject first used an intuitive BCI mapping (i.e., the relation between neural activity and cursor velocity) to control the cursor. We then perturbed the mapping guided by the intrinsic manifold (IM), which is a subspace of the multi-neuronal firing rate space defined by the frequently generated co-modulation patterns. Some perturbations required the monkey to combine the basis co-modulation patterns of the IM in new ways. Other perturbations required the monkey to generate new co-modulation patterns outside the IM. We hypothesized that it would be easier for the subject to learn to control within-manifold perturbations than outside-manifold perturbations. Across all sessions, the data supported this hypothesis. This suggests that the IM reflects constraints imposed by the underlying circuitry, which prevented the monkey from readily generating co-modulation patterns outside the IM. This work provides a novel network-level explanation for why learning some tasks might be easier than others.

T-30. State-space models for cortical-muscle transformations

Jeffrey Seely[1]   JSSEELY@GMAIL.COM
Matthew Kaufman[2]   ANTIMATT@GMAIL.COM
Christopher J Cueva[1]   CCUEVA@GMAIL.COM
Liam Paninski[1]   LIAM@STAT.COLUMBIA.EDU
Krishna Shenoy[3]   SHENOY@STANFORD.EDU
Mark Churchland[1]   MC3502@COLUMBIA.EDU

[1]Columbia University
[2]Cold Spring Harbor Laboratory
[3]Stanford University

A common desire in systems neuroscience is to infer models that link recorded inputs and outputs. Input-output models can be static—mapping inputs at each time to outputs at that time—or dynamic, where outputs may depend, via internal dynamics, on the history of inputs. A paradigmatic example of input/output transformations is from motor cortex to muscle activity. Over the last 45 years, the nature of this transformation has been hotly debated. Does motor cortex code kinematic parameters, which are then transformed into muscle activity by the spinal cord and other subcortical structures? Or does motor cortex produce muscle activity directly? Both possibilities might involve static or dynamic transformations. Prior studies have sometimes fit muscle activity with a temporally filtered version of neural activity, consistent with the possibility that spinal dynamics may provide filtering. However, such studies have not applied standard system identification techniques to determine (1) whether dynamics provide additional explanatory power or (2) the nature of any underlying dynamics. We therefore compared the performance of a linear state-space model with that of a static model. Analysis was based on data collected from one monkey performing a delayed reach task (180 single unit recordings from motor cortex, 8 EMG recordings from muscles, 27 reach conditions). The state-space model was identified using subspace identification followed by an optimization routine. The dynamical model fit 85% of the output variance using 11 input dimensions, translating to 130 system parameters. The static model required 42 input dimensions to achieve 85% fit, translating to 336 model parameters. Thus, the state-space model can fit outputs using many fewer model parameters. Generalization performance was also superior for the state space models. Finally, examination of the underlying dynamics revealed that the improved performance was due not just to exponential filtering but to more complex dynamics as well.