290. Extracting Dynamical Structure Embedded in Motor Preparatory Activity

Byron M. Yu1, Afsheen Afshar1, Gopal Santhanam1, Stephen I. Ryu1, Krishna V. Shenoy1, Maneesh Sahani2
1 Stanford University,2 Gatsby Computational Neuroscience Unit, UCL

At present, the best view of the activity of a neural circuit is provided by multiple-electrode extracellular recording technologies. These technologies are able to simultaneously measure spike trains from up to a few hundred cells in one or more brain areas during each trial. However, characterizing the dynamics of the circuit from such data is made difficult, both by the point-like nature of the spikes and by the variability in the responses to identical repeated trials. By their nature, spikes give us only an occasional view of the process from which they are generated.

The classic approach to handling both problems is to average responses from different trials. There is little alternative to this approach when recordings are made one cell at a time, even when the dynamics of the system are the subject of study. Unfortunately, such averaging can obscure important internal features of the response. In many experiments stimulus events provide the trigger for activity, but the resulting time-course of the response is internally regulated and may not be identical on each trial. This is especially important during cognitive processing such as decision making or planning. In this case, the mean trajectory obtained by averaging across trials may not reflect the true trial-by-trial dynamics. For example, a sharp change in firing rate that occurs with varying latency might appear as a slow smooth transition in the mean. An alternative approach is to adopt latent variable methods from the statistics and machine learning literatures, and to identify a hidden dynamical system that can summarize and explain the recorded spike trains. This is made possible by simultaneously recording from multiple neurons during each trial. The central idea is that the responses of different neurons reflect different views of a common dynamical process in the network, whose effective dimensionality may be lower than the number of units in the network. While the underlying state trajectory may be slightly different on each trial, the commonalities among these trajectories can be captured by the network's parameters, which are shared across trials. These parameters define how the network evolves over time, as well as how the observed spike trains relate to the network's state at each time point.

The use of latent variable models with hidden dynamics for neural data has, thus far, been limited. In [1-2], small groups of cells in the frontal cortex were modeled using hidden Markov models in which the latent dynamical system is assumed to transition between a set of discrete states. In [3], a state space model with linear hidden dynamics and point-process outputs were applied to simulated data. However, these restricted latent models cannot capture the richness of dynamics that recurrent networks exhibit. In particular, systems that converge toward point or line attractors, exhibit limit cycle oscillations, or even transition into chaotic regimes have long been of interest in neural modeling. If such systems are relevant to real neural data, we must seek to identify hidden models capable of reflecting this range of behaviors.

In this work, we developed a latent variable model having (1) hidden underlying recurrent structure with continuous-valued states, and (2) Poisson-distributed output spike counts, conditioned on the state. We applied the model to neural data recorded using a 96-electrode array chronically implanted in the dorsal premotor cortex (PMd) of a rhesus monkey. The monkey performed 26 center-out delayed reaches to each of seven possible, pseudo-randomly chosen radial targets. In this task, a ``planning'' period, whose duration is also pseudo-randomly chosen, separates movement-target specification and the movement-initiation cue. PMd neurons are believed to be involved in movement preparation and are active during this delay period.

Recent evidence from our laboratory suggests that PMd neural activity settles to a movement-specific state during the delay period. We found that the firing rate variability, measured {\\it across} trials and normalized by mean firing rate, drops during the delay period. This finding is consistent with the hypothesis that the process underlying motor planning is initially variable across trials, but settles during the delay period toward a subspace appropriate for the upcoming movement. However, it is difficult to view the settling process on a trial-by-trial basis using this variability measure.

In an attempt to observe this settling process on a trial-by-trial basis, we fit the latent variable model with underlying recurrent structure and Poisson outputs to PMd spiking data during the delay period from 100 simultaneously-recorded single- and multi-units. We used the Expectation-Maximization algorithm, which iteratively (1) infers a unique underlying state trajectory for each trial and (2) learns the system parameters, which are shared across trials. Although approximations were made to make the nonlinearities in the model tractable, this approach has provided the first glimpse of the hypothesized settling process toward a restricted subspace on a single-trial basis.

This work was supported by NDSEG and NSF Graduate Research Fellowships (BMY and GS), the Gatsby Charitable Foundation (BMY and MS), the Medical Scientist Training Program (AA), the Christopher Reeve Paralysis Foundation (SIR and KVS), and the following awards to KVS: NSF Center for Neuromorphic Systems Engineering at Caltech, ONR, Whitaker Foundation, Center for Integrated Systems at Stanford, Sloan Foundation, and Burroughs Wellcome Fund Career Award in the Biomedical Sciences.

References

[1] M. Abeles, H. Bergman, I. Gat, I. Meilijson, E. Seidemann, N. Tishby, and E. Vaadia. Cortical activity flips among quasi-stationary states. Proc Natl Acad Sci USA, 92:8616-8620, 1995.

[2] I. Gat, N. Tishby, and M. Abeles. Hidden Markov modelling of simultaneously recorded cells in the associative cortex of behaving monkeys. Network, 8(3):297-322, 1997.

[3] A. C. Smith and E. N. Brown. Estimating a state-space model from point process observations. Neural Comput, 15(5):965-991, 2003.