

Hidden Markov models applied toward the inference of neural states and the improved estimation of linear receptive fields

Sean Escola and Liam Paninski

Columbia University

Several recent experimental results suggest that neurons are associated with multiple firing regimes, or states (e.g. tonic and burst modes of LGN cells and up-and-down states in the cortex). We develop a general framework for estimating neural receptive fields (RFs) from paired spike-train and stimulus data assuming that neurons transition between several discrete hidden states. Previous approaches applied to RF estimation such as the spike-triggered average, the linear-nonlinear-Poisson (LNP) point-process model, and information theoretic techniques are predicated on the assumption that all spikes are equally informative about the stimulus. If, instead, some spikes occur while the neuron is in a stimulus-ignoring state, then including those spikes in the RF estimation will necessarily worsen the estimate. Furthermore, if the neuron moves between several states, each of which responds to different features of the stimulus, then these techniques will discover some composite RF that may differ significantly from all of the individual, state-specific RFs. By discovering the hidden state of the neuron at every point in time, each of the individual RFs can be estimated for each of the hidden states.

We have modified the traditional hidden Markov model (HMM) theoretical framework (Rabiner, 1989) to allow for point-process observables (i.e. spike-trains) and to be parameterized by N^2 linear filters of the time-varying stimulus (where N is the number of states) rather than static, conditional probability tables. Specifically, each state is associated with its own N filters: $N - 1$ of these determine the transition rates from the current state to the other states, and the remaining filter determines the firing rate for the current state. This latter filter is the canonical RF (although now there are N of them, one for each state), while the former filters are “RFs” for the state dynamics, a new concept in sensory neurophysiology. The actual transition and firing rates are the result of nonlinear transformations of the dot-products of the filters and the stimulus (i.e. as in LNP models). Stimulus-dependent transition rates are required to ensure that the neuron can be prompted by the stimulus to enter a particular state needed to respond to a feature of the stimulus.

The filters are learned using expectation-maximization (EM)—specifically the Baum-Welch algorithm—to maximize the log-likelihood as with traditional HMMs. Assuming the nonlinearities used in the model conform to a certain class of functions (Paninski, 2004), the M-step of EM is concave in the parameter space with a unique solution easily found via gradient ascent. We show the results of training from a number of simulated spike-train and stimulus pairs. The linear filters recovered by our algorithm nicely match the filters used to generate the data.

References

- Paninski, L. (2004). Maximum likelihood estimation of cascade point-process neural encoding models. *Network: Computation in Neural Systems*, 15:243–262.
- Rabiner, L. (1989). A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77:257–286.