An Efficient Algorithm for Sequential Optimal Design of Neurophysiology Experiments

Jeremy Lewi[1], Robert Butera[1], and Liam Paninski[2]

[1]Georgia Tech, [2]Columbia University
http://www.prism.gatech.edu/∼gtg120z/rtmodel

We present an efficient algorithm for selecting optimal stimuli for estimating the conditional response function of a neuron. Stimuli are selected by optimizing an objective function which quantifies the expected reduction in uncertainty about the unknown response function. Our objective function is based on mutual information and leads to an optimality criterion known as D-optimality [1]. Our implementation overcomes the computational hurdles of sequential optimal experimental design in this setting. Simulations show that using optimally chosen stimuli can reduce the number of trials needed to estimate the conditional response function by more than an order of magnitude.

Our algorithm has three main components. 1) We model the conditional response parametrically using General Linear Models (GLMs). This a very flexible nonlinear model which can capture many firing rate statistics of a neuron. Furthermore, it is possible to incorporate effects such as adaptation, refractory periods, and burstiness. We consider a restricted class of nonlinearities for the GLM which ensures the log-likelihood is concave [2]. Concavity of the likelihood improves the tractability of many of the necessary computations. 2) We approximate the posterior distribution on the parameters of the conditional response function as a Gaussian distribution. Asymptotically this approximation is accurate [3]. The normal approximation makes it easier to update the estimated parameters and to compute the mutual information. 3) We show that choosing the stimulus to maximize the mutual information requires at worst one 2-d optimization per trial.

We present a number of simulations to demonstrate the potential utility and applicability of our algorithm to neurophysiology experiments. A theoretical and empirical analysis shows that the running time of our algorithm grows on average as the square of the dimensionality. We compared the estimated parameters using stimuli drawn according to our algorithm to the estimated parameters using I.I.D stimuli drawn from a uniform distribution. These simulations looked at the performance when: 1) the parameters are very high dimensional 2) spike history effects are included and 3) the parameters are non-stationary. The validity of our Gaussian approximation is tested using Monte-Carlo methods to measure the Kullback-Leibler distance to the true posterior. Finally we show that asymptotically our uncertainty about the unknown parameters decreases at a rate near that predicted by a theoretical analysis of the information maximizing approach.

Acknowledgments JL is supported by the DOE and by the NSF. LP is supported by the NEI and the Gatsby Foundation. We thank P. Latham for helpful conversations.

References

[1] V. V. Fedorov, *Theory of Optimal Experiments* (Academic Press, 1972).

[2] L. Paninski, *Network: Computation in Neural Systems* **15**, 243 (2004).

[3] L. Paninski, *Neural Computation* **17**, 1480 (2005).