Markov Chain Monte Carlo Methods for Decoding Neural Spike Trains

Yashar Ahmadian[1], Jonathan Pillow[2], and Liam Paninski[1]

[1]Columbia University, [2]University College London

Stimulus reconstruction or *decoding* methods provide an important tool for understanding how sensory and motor information is represented in neural activity. We address Bayesian decoding methods based on an *encoding* generalized linear model (GLM) [1, 2] that accurately describes how stimuli are transformed into the spike trains of a group of neurons. The log-concave GLM likelihood is combined with a prior distribution to yield the posterior distribution over the stimuli that possibly generated an observed set of spike responses. This posterior is log-concave so long as the prior is, meaning that the maximum *a posteriori* (MAP) stimulus estimte can be obtained using highly efficient optimization algorithms [3]. Unfortunately, however, the MAP estimate can have a relatively large average error when the posterior is highly non-Gaussian.

Here we introduce several Markov chain Monte Carlo (MCMC) algorithms that allow for the calculation of general Bayesian estimators involving posterior expectations (integrals). An efficient version of the "hit-and-run" algorithm [4], exploiting the log-concavity property, is shown to be superior to other MCMC methods when the prior distribution has sharp edges and corners. The Metropolis-adjusted Langevin algorithm (MALA) [5] was significantly superior to other MCMC methods for Gaussian priors. Using these algorithms we show that for the former class of priors the posterior mean estimate can have a considerably lower average error than MAP, whereas for Gaussian priors the two estimators have equal efficiency.

We also consider the effect of uncertainty in the GLM parameters on the posterior estimators. When this uncertainty increases, the posterior mean shifts towards the prior mean as the Bayesian decoder relies more heavily on the prior information and less on the observed spikes. Finally, by using MALA to calculate the mutual information between the stimulus and response exactly, we verify the validity of a computationally efficient Laplace approximation to this quantity for Gaussian priors in a wide range of model parameters. This makes direct model-based computation of the mutual information tractable even in the case of large observed neural populations, where methods based on binning the spike train fail.

Acknowledgments
We Thank M. Kennel, E. Simoncelli, E.J. Chichilnisky, Y. Ding, T. Teravainen, and K. Rahnama Rad for helpful conversations. LP and YA are supported by NEI R01 EY018003; LP is additionally supported by an Alfred P. Sloan Research Fellowship. JP is supported by a Royal Society USA/Canada Research Fellowship.

References

[1] McCullagh, P. and Nelder, J. (1989). *Generalized linear models*. Chapman and Hall, London.
[2] Paninski, L. (2004). Maximum likelihood estimation of cascade point-process neural encoding models. *Network: Computation in Neural Systems*, 15:243–262.
[3] Pillow, J. and Paninski, L. (2007). Model-based decoding, information estimation, and change-point detection in multi-neuron spike trains. *Under review, Neural Computation*.
[4] Lovasz, L. and Vempala, S. (2004). Hit-and-run from a corner. In *Proc. of the 36th ACM Symposium on the Theory of Computing (STOC '04)*, Chicago.
[5] Roberts, G. O. and Tweedie, R. L. (1996). Exponential convergence of langevin diffusions and their discrete approximations. *Biometrika*, 2:341–363.