Deciphering correlations: Bayesian decoding of multi-neuronal spike trains in primate retina

Jonathan W. Pillow[1], Jonathon Shlens[2], Liam Paninski[3], Alexander Sher[4], Alan M. Litke[4], E.J. Chichilnisky[2], Eero P. Simoncelli[5]

[1]Gatsby Computational Neuroscience Unit, UCL; [2]The Salk Institute; [3]Columbia University; [4]University of California, Santa Cruz [5]HHMI, Center for Neural Science, and Courant Inst. Math. Sciences, NYU

Groups of neurons are known to exhibit correlated spiking activity, but the relevance of correlations for information processing is a topic of much current debate. We examine this issue by developing an optimal decoder for reconstructing a visual stimulus from the observed spike trains of a population of retinal ganglion cells. The decoder is based on a generalized linear model (GLM) that accurately represents the probabilistic encoding of stimulus information in the spiking responses of the population. The model for each cell consists of a bank of linear filters that capture stimulus dependence, spike history dependence, and dependence on the spikes of other cells, followed by an exponential nonlinearity and instantaneous (Poisson) spike generation. The filters operating on spike trains of other cells serve as "coupling terms" that capture interactions between cells, allowing for time-varying correlations (and anti-correlations) beyond those induced by the stimulus.

We fit this model to a population of several dozen simultaneously-recorded macaque retinal ganglion cells (ON and OFF parasol cells), whose receptive fields completely covered several square degrees of visual space. We used a regularized maximum likelihood fitting procedure to determine functional connectivity, and found strong positive coupling between pairs of ON cells and pairs of OFF cells, and negative coupling between ON-OFF pairs, with the strongest coupling occurring between nearby pairs of cells. We also fit a version of the model in which coupling filters were eliminated, making the response of each model cell conditionally independent of its neighbors (i.e., dependent only on the stimulus and its own spiking history).

We then used the model to perform optimal Bayesian decoding (reconstruction) of a novel stimulus from a set of observed spike times. By comparing the performance of the uncoupled model and with that of the full (coupled) model, we assessed the importance of correlated spiking for encoding and decoding. We found that: (1) spatial receptive fields are significantly smaller under the coupled model, suggesting that some portion of the classical receptive field can be more accurately explained in terms of the spiking activity of nearby cells; (2) the coupled model provides significantly more accurate predictions of multi-neuronal spike trains, faithfully accounting for second and higher-order correlations; (3) model-based decoding was significantly more accurate than optimal linear decoding; and (4) stimuli can be decoded $\sim$10% more accurately using reconstruction based on the coupled model compared to the uncoupled model. These results indicate that correlations in retinal spiking activity cannot be ignored without losing information.

Acknowledgments
Funding was provided a Royal Society USA/Canada Research Fellowship (JP); NSF IGERT DGE-03345 (JS); NEI grant EY018003 (EJC, LP, & EPS); Gatsby Foundation Pilot Grant (LP); US National Science Foundation grant PHY-0417175 (AML); McKnight Foundation (AML & EJC); and HHMI (EPS).