Sparse Bayesian inference and experimental design for synaptic weights and locations

Ari Pakman, Carl Smith and Liam Paninski
Columbia University

A detailed understanding of the organization of local neural circuits requires determining not only the connectivity pattern in a neural population, but also the exact location and strength of synaptic interactions in the dendritic tree.

In a previous work [1], we showed how to approach this problem by combining the ability to stimulate individual presynaptic neurons with simultaneous imaging of postsynaptic neurons at subcellular resolution. This work extends our previous results in two directions. On the one hand we revisit the inference method used to extract the locations and strengths of the synaptic weights from the observed data. While in [1] the synaptic weights were the maximum a posteriori (MAP) solution of a state-space model with an $L1$ prior (the "Lasso" model), in this work we also obtain confidence intervals by adopting a fully Bayesian approach. In particular, we compare the results of several popular sparsity-inducing priors for the synaptic weights: the Bayesian Lasso [2], the Horseshoe [3] and the Spike-and-Slab [4]. Particular emphasis is placed on the constraint imposed by Dale's law, which states that the synaptic weights have a definite sign, thus leading to truncated probability distributions.

Equipped with the full posterior distribution of the synaptic weights, our second contribution explores optimal experimental design. We extend the type of voltage measurements from localized observations to linear combinations of voltages across several locations with random coefficients. This setting corresponds to a "compressed sensing" sampling scheme [5], which yields an impressive reduction in the number of measurements required to infer the synaptic weights. In particular, we show how to choose the correlation among the random coefficients to offset the correlation between successive measurements imposed by the neuron dynamics. We illustrate our results on simulated measurements in toy and real neurons.

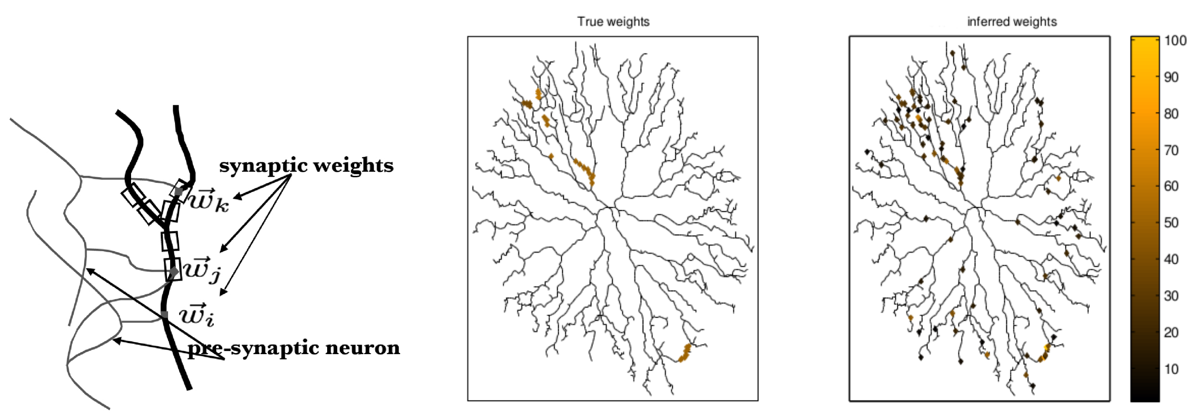

Figure 1: Schematic of proposed method. By observing a noisy, subsampled spatiotemporal voltage signal on the dendritic tree, we can infer the strength of a given presynaptic cell's inputs at each location on the postsynaptic cell's dendritic tree.

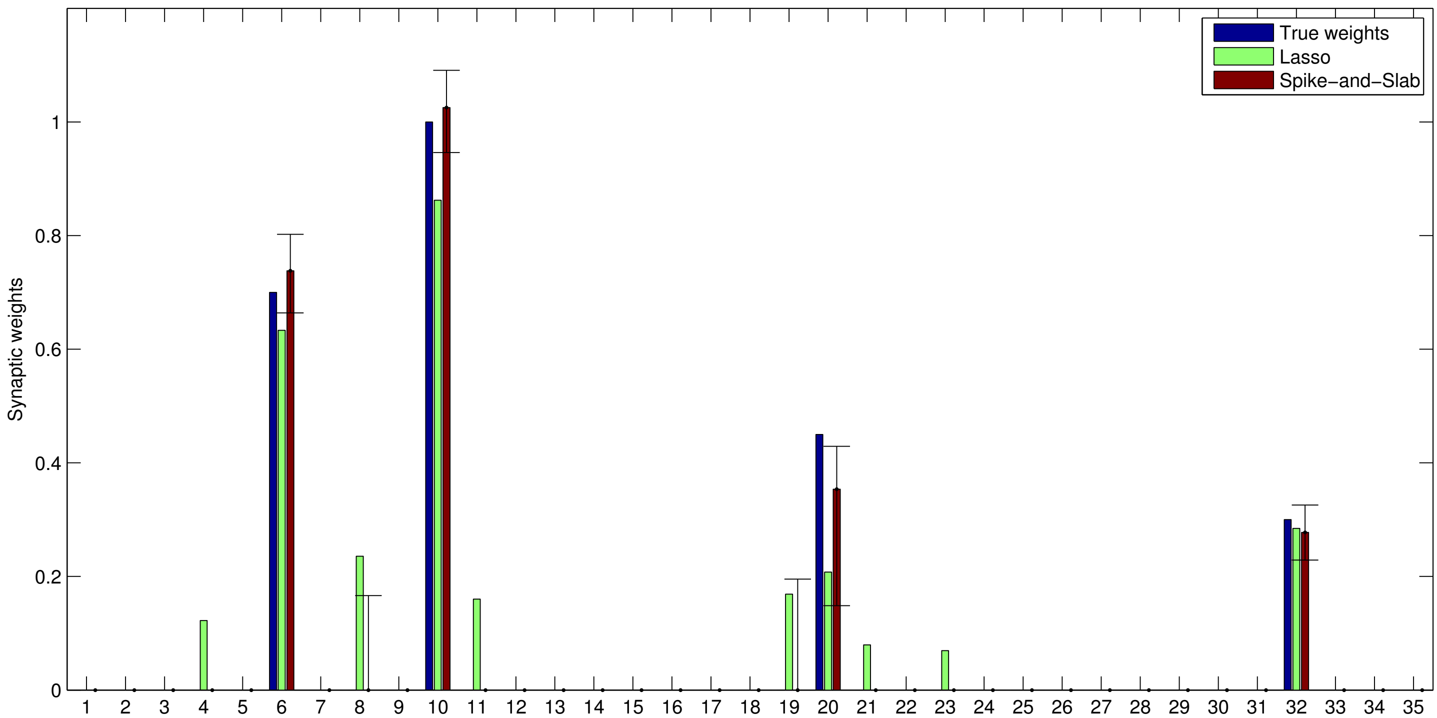

Figure 2: Inference of synaptic weights using sparsity-inducing priors. Synaptic weights inferred in a T-shaped toy neuron with 35 compartments similar to that studied in [1]. The neuron had non-zero synaptic weights at compartments 6, 10, 20 and 32. The green bars show the Lasso estimates using the techniques developed in [1]. The heights of the red bars are the medians, at each compartment, of the posterior distribution using a Spike-and-Slab prior. The .25 and .75 quantiles are also indicated. The Spike-and-Slab prior was superior to the Bayesian Lasso and the Horseshoe priors (not shown).

[1] A. Pakman, J. Huggins, and L. Paninski. Fast penalized state-space methods for inferring dendritic synaptic connectivity. *Under revision. Available at http://stat.columbia.edu/∼liam/research/pubs/pakman-huggins-synaptic.pdf.*

[2] T. Park and G. Casella. The Bayesian lasso. *Journal of the American Statistical Association*, 103(482):681–686, 2008.

[3] C.M. Carvalho, N.G. Polson, and J.G. Scott. The horseshoe estimator for sparse signals. *Biometrika*, 97(2):465–480, 2010.

[4] T.J. Mitchell and J.J. Beauchamp. Bayesian variable selection in linear regression. *Journal of the American Statistical Association*, 83(404):1023–1032, 1988.

[5] V. Studer, J. Bobin, M. Chahid, H. Moussavi, E. Candes, and M. Dahan. Compressive fluorescence microscopy for biological and hyperspectral imaging. *Preprint arXiv:1201.0650*, 2012.