Bayesian Receptive Fields and Neural Couplings with Sparsity Prior and Error Bars

Sebastian Gerwinn[1], Matthias Seeger[1], Günther Zeck[2], and Matthias Bethge[1]

[1]Max Planck Institute for Biological Cybernetics, Tübingen,
[2]Max Planck Institute of Neurobiology, Martinsried

Here we apply Bayesian system identification methods to infer stimulus-neuron and neuron-neuron dependencies. Rather than reporting only the most likely parameters, the posterior distribution obtained in the Bayesian approach informs us about the range of parameter values that are consistent with the observed data and the assumptions made. In other words, Bayesian receptive fields always come with error bars. In fact, we obtain the full posterior covariance, indicating conditional (in-)dependence between the weights of both, receptive fields and neural couplings. Since the amount of data from neural recordings is limited, such uncertainty information is as important as the usual point estimate of the receptive field itself.

We employ expectation propagation, a recently developed approximation of Bayesian inference, to a multi-cell response model consisting of a set of coupled units, each of which is a Linear-Nonlinear-Poisson (LNP) cascade neuron model. The instantaneous firing rate of each unit depends on both the spike train history of the units and the stimulus. Parameter fitting in this model has been shown to be a convex optimization problem [1], which can be solved efficiently. By doing inference in this model we can determine excitatory and inhibitory interactions between the neurons and the dependence of the stimulus on the firing rate. In addition to the uncertainty information (error bars) obtained within the Bayesian framework one can impose a sparsity-inducing prior on the parameter values. This forces weights actively to zero, if they are not relevant for explaining the data, leading to a more robust estimate of receptive fields and neural couplings, where only significant parameters are nonzero.

The approximative Bayesian inference technique is applied to both artificially generated data and to recordings from retinal ganglion cells (RGC) responding to white noise (m-sequence) stimulation. We compare the different results obtained with a Laplacian (sparsity) prior and a Gaussian (no sparsity) prior via Bayes factors and test set validation. For completeness, the receptive fields based on classical linear correlation analysis and maximum likelihood estimation are included into the comparison.

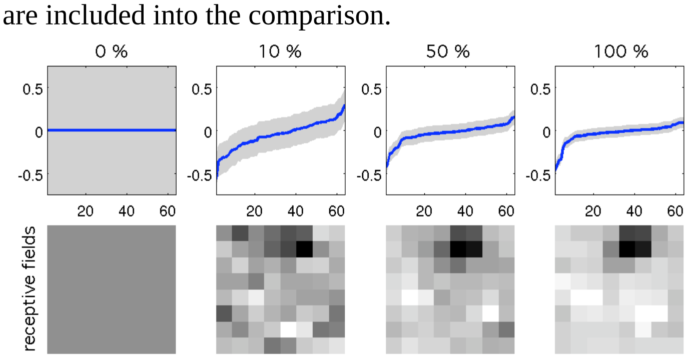

Illustration of Bayesian receptive field. The figure shows the inferred receptive fields of one neuron (lower) as well as the confidence range of the sorted pixel values (upper) when using a different fraction of the data (0, 10, 50, and 100% , n=8452 spikes).

References

[1] Paninski L (2004) Maximum likelihood estimation of cascade point process neural encoding models. *Network: Computation in Neural Systems* 15:243-262, 2004