Fast low-SNR high-dimensional optimal filtering, applied to inference of dynamic receptive fields

Liam Paninski Kamiar Rahnama Rad Jonathan Huggins
Columbia University

The Kalman filter implements the optimal Bayesian filter in the linear-Gaussian setting and is therefore a workhorse of statistical time series analysis. In addition, straightforward non-Gaussian generalizations are available; the resulting optimal filters provide state-of-the-art methods for tracking nonstationary tuning curves and for decoding time-varying behavior in a variety of neural settings.

However, the state variable in many problems is very high-dimensional. Standard implementations of the Kalman filter require $O(N^3)$ time and $O(N^2)$ space per time step, where N is the dimensionality of the state variable, and are therefore impractical when N is large. For example, we might need $N > 100$ basis functions to adequately parameterize a tuning curve in a two-dimensional space (such as a place field or a grid field), and N might need to be much larger ($N > 1000$) to adequately parameterize tuning curves for higher-dimensional variables.

In this paper we note that if a relatively small number of low-SNR observations are available per time step (as is frequently the case in neural applications), then the Kalman equations may be very well approximated in terms of a low-rank perturbation of the steady-state zero-SNR solution. This approximation may in turn be computed and updated very efficiently (often in just $O(N)$ or $O(N \log N)$ time and space per time step), using fast methods from numerical linear algebra. This opens up the possibility of real-time adaptive experimental design and optimal control in systems of much larger dimensionality than was previously feasible. We have previously described an application of these methods to the special case of optimal smoothing of spatiotemporal voltage observations on large dendritic trees, but the basic idea turns out to be much more general. We discuss a number of examples, focusing on efficient estimation of dynamic receptive field and place field properties, with applications to both simulated and real data.