Ph.D. Thesis

Title: Correlated Estimation Problems and the Ensemble Kalman Filter

Author: Jan Čurn

Abstract:

The Kalman filter is a recursive algorithm that estimates the state of a linear dynamic system from a sequence of noisy sensor measurements. Due to its relative simplicity, numerical efficiency and optimality, the Kalman filter and its variants have been applied to a wide range of problems in technology, notably in the areas of guidance, navigation, and control. The traditional definition of the Kalman filter is based on the assumption that at any given time, the errors associated with the predicted state estimate and the observation are statistically independent. However, in many practical problems, this assumption is not satisfied, and as such the Kalman filter may provide overconfident state estimates and diverge. This can have serious consequences in the context of safety-critical systems.

Although there are modifications of the Kalman filter that accommodate various types of correlation in the process and observation noises, these are not suitable in the situation where the correlation between the errors associated with the predicted state estimate and the observation is caused by the presence of common past information between the state estimate and the observation, which is characteristic of distributed sensor networks. On the contrary, existing methods that deal with the common past information problem either provide overly conservative estimates, or have too strict assumptions on the structure of the problem, such as the communication topology of the sensor network.

This thesis presents two new filters to address various correlated estimation problems that are based on the Ensemble Kalman filter, a Monte Carlo variant of the Kalman filter, which represents the state estimates and observations using sets of random samples instead of the conventional mean vectors and covariance matrices. Specifically, both of these filters provide a new generalised update rule that computes consistent state estimates even in the presence of correlation between the errors associated with the state estimate and the observation. This is only possible due to the fact that in the context of the Ensemble Kalman filter, the magnitude of such a correlation can be estimated from the random samples.

The new filters retain all of the important features of the Ensemble Kalman filter, such as scaling linearly with the number of state-space dimensions, and supporting non-linear process and observation models. An analysis of the numerical properties of the filters is provided, including a comparison with state-of-the-art methods in several benchmark scenarios. Furthermore, in order to demonstrate their practical utility, the new filters have been applied to three different real-world problems in the larger field of robot localisation: cooperative vehicle localisation, simultaneous localisation and mapping, and global satellite-based positioning.