

GPU-Accelerated Bayesian Learning and Forecasting in Simultaneous Graphical Dynamic Linear Models

Lutz Gruber

Mike West

Abstract

We discuss Bayesian analysis of dynamic models customized to learning and prediction with increasingly high-dimensional time series. A new framework of simultaneous graphical dynamic models allows the decoupling of analyses into those of a parallel set of univariate time series dynamic models, while flexibly modeling time-varying, cross-series dependencies and volatilities. The strategy allows for exact analysis of univariate time series models that are then coherently linked to represent the full multivariate model. Computation uses importance sampling and variational Bayes ideas, and is ideally suited to GPU-based parallelization. The analysis and its GPU-accelerated implementation is scalable with time series dimension, as we demonstrate in an analysis of a 400-dimensional financial time series.

The full article is available online at <http://projecteuclid.org/euclid.ba/1425304898>