University of California, Irvine
Statistics Seminar

Bayesian Structural Time Series

Steve Scott
Google

Thursday, Oct. 6, 2016
4 p.m., 6011 Bren Hall
(Bldg. #314 on campus map)

Bayesian structural time series models are a flexible and powerful tool for modeling time series data. The models have an additive structure which permits an analyst to add latent state components for handling trend, seasonal, regression, and other structural elements of a time series as needed. The additive structure makes it easy to place informative priors on individual components, such as a spike and slab prior on the regression component when working with large numbers of contemporaneous predictors. Efficient computation is possible for the time series component (using the Kalman filter) and the regression component (integrating out the coefficient vector) when the model is Gaussian. Data augmentation can be used to extend these efficient computational methods to many other model families including probit, binary and multinomial logit, Poisson, student T, Laplace, and support vector machines.

This talk describes the bsts R package, which was developed to allow Google employees who are non-experts in Bayesian modeling, or time series modeling, to work with these models. The package has been used for a variety of purposes, including nowcasting economic time series, anomaly detection, longer term forecasting, and causal inference.