

CRUNCH Seminars at Brown, Division of Applied Mathematics

Friday – June 8, 2018

Deep and confident prediction for time series at Uber

(Zhu & Laptev, arXiv:1709.01907)

Yang Liu

Reliable uncertainty estimation for time series prediction is critical in many fields, including physics, biology, and manufacturing. At Uber, probabilistic time series forecasting is used for robust prediction of number of trips during special events, driver incentive allocation, as well as real-time anomaly detection across millions of metrics. Classical time series models are often used in conjunction with a probabilistic formulation for uncertainty estimation. However, such models are hard to tune, scale, and add exogenous variables to. Motivated by the recent resurgence of Long Short Term Memory networks, we propose a novel end-to-end Bayesian deep model that provides time series prediction along with uncertainty estimation. We provide detailed experiments of the proposed solution on completed trips data, and successfully apply it to large-scale time series anomaly detection at Uber.