

Learning to Forecast Dynamical Systems from Streaming Data

R. Ward¹

¹*UT Austin*

Kernel analog forecasting (KAF) is a powerful methodology for data-driven, non-parametric forecasting of dynamically generated time series data. This approach has a rigorous foundation in Koopman operator theory and it produces good forecasts in practice, but it suffers from the heavy computational costs common to kernel methods. In this talk, we will discuss a streaming algorithm for KAF that only requires a single pass over the training data. This algorithm dramatically reduces the costs of training and prediction without sacrificing forecasting quality. Computational experiments demonstrate that the streaming KAF method can successfully forecast several classes of dynamical systems. The overall methodology may have wider interest as a new template for streaming kernel regression.

This is joint work with Dimitris Giannakis, Amelia Henriksen, and Joel Tropp.