University of Pittsburgh | School of Computing and Information

Proposal Defense

Doctor of Philosophy in Computer Science

"**Neural Models for Multivariate Temporal Prediction**" by **Ankitkumar Joshi**

Date: November 25, 2025
Time: 10:00 am – 12:00 pm
Place: 6106 E. Lilly, Sennott Square, Pittsburgh PA 15260

Committee:

- Milos Hauskrecht, Professor, Computer Science Department, School of Computing and Information, University of Pittsburgh
- Adriana Kovashka, Associate Professor, Computer Science Department, School of Computing and Information, University of Pittsburgh
- Xiaowei Jia, Assistant Professor, Computer Science Department, School of Computing and Information, University of Pittsburgh
- Geoffrey J Gordon, Professor, Machine Learning Department, Carnegie Mellon University

Abstract:
Understanding natural and human-built systems requires analyzing how they change over time. Time series data, measurements collected over time across different aspects of a system, are fundamental for capturing system dynamics. They reveal patterns that explain behavior, enable prediction of future events, and support decisions based on these dynamics. For example, in healthcare, monitoring a patient's vital signs can provide early warnings of clinical deterioration. In meteorology, tracking temperature, pressure, and humidity over time allows forecasters to predict future conditions. Careful analysis of such temporal data is essential for building reliable and responsive systems across domains.

Predicting outcomes from real-world temporal data, however, remains challenging. Observations are often irregular, noisy, or incomplete due to missing sensors. Dependencies may be multi-scale, non-linear, and non-stationary, spanning both time and interacting variables. These characteristics violate common modeling assumptions, such as regular sampling or stationary dynamics, and limit the effectiveness of traditional prediction methods.

This thesis addresses the above challenges using modern neural models, focusing on three core research problems and their solutions:

1. **Irregular Time Series Modeling:** Many recent models for irregular temporal data are complex, yet it is unclear whether they outperform simpler, principled recurrent architectures. To examine this, we propose GRUwE, a Markov state approximation model that delivers strong performance in regression and event prediction while maintaining a compact state for efficient, real-time inference.

2. **Next-Event Distribution Modeling:** Existing models for predicting the next event often rely on assumptions that reduce accuracy, either oversimplifying dependencies or overfitting limited data. We explore methods to balance model complexity by sharing parameters across event types and regularizing arrival time representations.
3. **Modeling Nonstationary Dynamics:** Real-world temporal sequences frequently exhibit nonstationary behavior. Building on advances in frequency-domain forecasting, we develop models based on the Discrete Wavelet Transform, which capture both local transients and long-range patterns through multi-scale representations.