University of Pittsburgh | School of Computing and Information

Dissertation Defense

Doctor of Philosophy in Computer Science

"**Modeling Clinical Multivariate Event Time-Series with Neural Temporal Models**" by
Jeong Min Lee

Date: November 29, 2021
Time: 3:00PM – 5:00PM
Place: 5317 Seminar Room, Sennott Square, 210 S Bouquet St, Pittsburgh, PA 15213
https://pitt.co1.qualtrics.com/jfe/form/SV_eWZDFMtHABgTnYa

Committee:

- Milos Hauskrecht, Professor, Computer Science, University of Pittsburgh
- Erin Walker, Associate Professor, Computer Science, University of Pittsburgh
- Adriana Kovashka, Assistant Professor, Computer Science, University of Pittsburgh
- Shyam Visweswaran, Associate Professor, Department of Biomedical Informatics, University of Pittsburgh

Abstract:
Electronic health records (EHRs) contain detailed clinical information about the patient's condition, patient management, and patient outcomes as multivariate event time-series. The ability to accurately predict future events has great importance as it can help to improve patient care. However, it is challenging due to inherent complexities of the clinical multivariate event time-series: (1) high dimensionality: the event time-series consist of several thousands of different clinical events including administrations of medications, orders and results of lab tests, orders of medical procedures, etc; (2) different temporal characteristics: each event in the multivariate time-series has different temporal behavior and different temporal ranges of dependencies for precursor events; (3) patient-specific variability: based on underlying clinical complications, each patient's sequence may consist of different sets of clinical events. However, population-based models learned from such sequences may not accurately predict patient-specific dynamics of event sequences.

In this thesis, we propose novel autoregressive models that can predict future occurrences of the clinical multivariate event time-series. First, we propose a new model that handles different temporal dependencies using multiple temporal mechanisms that cover various timescales. Second, we propose a new temporal mechanism that can model repeatedly occurring events through a specialized external memory module. Third, we develop a new adaptive event sequence prediction framework that learns to adjust its prediction for individual patients through an online model update. We evaluate our proposed models on the real-world clinical data derived from EHRs of critical care patients. We show that our new models equipped with the above temporal mechanisms and the personalized learning framework lead to improved prediction performance compared to multiple baselines.