University of Pittsburgh | School of Computing and Information

## Dissertation Defense

***Doctor of Philosophy in Intelligent Systems***

"**Understanding and Quantifying Dataset Shift in Clinical Predictive Modeling**" by
**Mohammadamin Tajgardoon**

**Date:** December 4, 2025
**Time:** 3:30 – 5 p.m.
**Place:** Room 407A, DBMI Building, 5607 Baum Blvd, Pittsburgh, PA 15206
*** There is new security system when entering the building. From the parking garage, the new door code is 5225# (small box beside the old, bigger box at door). You will be able to take the elevator to the 4th floor and go to the classroom**.

**Committee:**

- Shyam Visweswaran, Professor, Intelligent Systems Program, School of Computing and Information
- Harry Hochheiser, Associate Professor, Intelligent Systems Program, School of Computing and Information
- Ahmad Tafti, Assistant Professor, Intelligent Systems Program, School of Computing and Information
- Vanathi Gopalakrishnan, Associate Professor, Biomedical Informatics, School of Medicine

**Abstract:**
Predictive models used in clinical decision support can become less reliable over time due to dataset shift, which refers to changes in data distributions between training and deployment. These shifts result from changes in patient populations, clinical practices, or data collection methods. Over time, they can erode model accuracy and fairness, often going unnoticed until trust or patient safety is compromised. This dissertation addresses that challenge by systematically examining how distributional changes impact model performance and by evaluating unsupervised divergence metrics as early-warning signals of degradation. A principled evaluation framework is introduced, combining controlled synthetic experiments with real-world validation on the MIMIC-IV critical care database. In the synthetic setting, we simulate healthcare data under defined covariate, label, and concept drifts to assess how different divergence metrics correspond to model degradation across multiple classifiers. These findings are extended through retrospective MIMIC-IV analyses, where temporal, demographic, and operational shifts are used to evaluate the robustness and interpretability of divergence signals in a realistic clinical context. Preliminary analyses using the HiDenIC ICU dataset at UPMC further demonstrate that divergence metrics can reflect meaningful changes in model reliability, supporting their feasibility as early-warning indicators.

The results provide empirical insights, methodological innovations, and practical guidance for managing dataset shift in clinical machine learning. Case-mix changes, which involve shifts in both features and outcomes, cause major drops in discrimination and calibration. These are reliably detected by feature-based and score-based divergence metrics. Conceptual shifts related to workflow or practice evolution also reduce performance, with representation-level diagnostics proving most sensitive. Demographic subgroup shifts, such as those defined by race or gender, frequently induce fairness and calibration failures, underscoring the need for equity-aware monitoring. By integrating synthetic and real-world evidence, the dissertation identifies the most reliable divergence metrics and derives actionable thresholds for triggering model reviews. An ensemble alarm strategy is also proposed, combining multiple diagnostics to deliver accurate and interpretable early warnings. Together, these contributions define a framework for trustworthy model monitoring, helping practitioners anticipate degradation, allocate resources for retraining, and maintain reliability in dynamic healthcare environments.