University of Pittsburgh | School of Computing and Information

Dissertation Defense
Doctor of Philosophy in Computer Science

"**Utility-aware Privacy and Model Integrity Analytics for IoT Systems**"
by **Ajesh Koyatan Chathoth**

Date: May 15, 2025
Time: 10:00 a.m. – 12:00 p.m.
Place: 6106 Sennott Square, 210 S Bouquet St., Pittsburgh, PA 15213

Committee:

- Dr. Stephen Lee, Committee Chair, Assistant Professor, Department of Computer Science, SCI
- Dr. Daniel Mosse, Professor, Department of Computer Science, SCI
- Dr. James B.D. Joshi, Professor, Department of Informatics and Networked Systems, SCI
- Dr. Amy Babay, Assistant Professor, Department of Computer Science, SCI

Abstract:

Internet of Things (IoT) devices have become ubiquitous in our daily lives, integrating intelligence through seamless data exchanges across distributed environments. From smart homes and industrial systems to wearables, IoT systems collect and process vast amounts of data that are pivotal for enabling monitoring and detecting anomalies to improve overall efficiency. However, while enabling utility, this data-centric paradigm also introduces critical challenges at the intersection of privacy preservation and model integrity. On the one hand, ensuring utility-aware privacy is essential to protect sensitive user data without compromising the functionality and effectiveness of data-driven services. On the other hand, the increasing reliance on machine learning models within IoT ecosystems exposes them to integrity threats, such as backdoor attacks, where adversaries embed malicious behavior during training, and model drift, where model performance degrades over time due to changing data distributions.

This dissertation first addresses the issue of utility degradation that occurs when privacy enhancement techniques are applied in heterogeneous data-driven IoT systems with varying privacy requirements. Specifically, I design and implement methods to improve the overall performance of IoT systems while satisfying user privacy needs within a differentially private federated learning framework. Additionally, I create a utility-aware mechanism that allows users to dynamically select their privacy preferences when sharing data with cloud-based services, such as those used for sensor-based activity recognition. Furthermore, I investigate the integrity-related challenges faced by data-driven systems to understand their vulnerability to backdoor adversarial attacks. To achieve this, I design and implement an attack model and conduct model integrity analytics. These analytics provide insights into how intrusion detection systems and sensor-based activity recognition models can be susceptible to backdoor attacks, revealing weaknesses in their integrity and highlighting the importance of developing robust data-driven systems.