University of Pittsburgh | School of Computing and Information

Dissertation Defense

Doctor of Philosophy in Information Science

"**An AI-driven Framework of Flood Preparedness in Urban Navigation**" by **Yue Qin**

Date: July 29, 2026
Time: 10 a.m. – 12:30 p.m.
Place: Room 6106, Sennott Square, 210 S. Bouquet Street, Pittsburgh, PA 15260

Committee:

- Rosta Farzan, Professor, Department of Informatics and Networked Systems, School of Computing and Information, University of Pittsburgh
- Xiaowei Jia, Associate Professor, Department of Computer Science, School of Computing and Information, University of Pittsburgh
- Na Du, Assistant Professor, Department of Informatics and Networked Systems, School of Computing and Information, University of Pittsburgh
- Peter Brusilovsky, Professor, Department of Informatics and Networked Systems, School of Computing and Information, University of Pittsburgh
- Tian Zheng, Professor of Statistics, Department of Statistics, Columbia University

Abstract:
Climate change and rapid urbanization are exposing more population under the threat of urban floods around the globe, with greater losses and damages expected in the near future. One way of mitigating risk and adapting to urban floods is through effective preparedness measures, which involve taking actions and making plans when floods are anticipated. Official sources such as FEMA provide guidelines for individual preparedness; however, these resources often do not provide support for urban mobility. Urban residents often do not have the right information or appropriate tools to prepare for their daily urban navigation when flood disruptions happen. In this dissertation, I approach the core question: "how can we support flood preparedness in urban navigation harnessing the capabilities of AI" by constructing an AI-driven framework consisting of four major components: *fine-grained data*, *data-driven and AI-assisted models*, *bias evaluation of AI models* (i.e., large language models, LLMs), and *application deployment* (i.e., multi-agent LLMs).

Data on urban floods lays the foundation of urban flood modeling and human decision-making. However, existing urban flood data either lacks fine-grained spatiotemporal resolution, or is not publicly available. To address the gap, the first goal of my dissertation has been focused on building a technical pipeline to extract and fuse street-level flood events from heterogeneous open data sources. Using Pittsburgh and New York City (NYC) as the case studies, I employed the proposed pipeline and constructed two datasets with street-level resolution of urban flood occurrences and made this data publicly available. *Models* on urban floods mostly adopt observations from physical or social sensors, suffering from limited spatial transferability and biased reporting in crowdsourcing. To overcome the limitations, I proposed and developed a latent-state urban flood model capable of both predicting true, unobserved flood occurrences at the street-segment scale and uncovering heterogeneous reporting behaviors. Results from systematic evaluation show that the proposed model achieves reasonable prediction performance while offering interpretable insights into disparities in crowdsourced reporting.

University of Pittsburgh | School of Computing and Information

While LLMs are becoming more powerful and increasingly adopted to assist decision-making during hazards, it is important to understand their limitations and weaknesses, especially possible biases embedded in the large data models used by these AI tools. Focusing on socioeconomic biases in hazard adaptation generation, I systematically evaluate six LLMs in generating flood adaptation recommendation and simulation under different prompts and trip scenarios. The findings illustrate the substantial amount of socioeconomic bias exists in different LLMs under different evaluation scenarios.

The final goal of the dissertation is to develop an application that can bring together the data and models developed earlier and support urban residents in their decision making when their daily navigation is interrupted by urban floods. By building a multi-agent LLM system, I demonstrate the possibility of harnessing the capabilities of agentic LLMs in hazard decision-support.

Taken together, the outcomes of this dissertation contribute both to the academic research and real-world decision-making among multiple stakeholders. Researchers and practitioners can apply the data processing pipeline and flood prediction model to promote open data sharing, train prediction models to understand flood patterns and make adaptation plans. Insights from LLM evaluation offer guidelines of employing LLMs in hazard mitigation and response, while the application provides novel ways of integrating agentic AI in assisting urban navigation during mobility disruptions.