COMPUTER SCIENCE & ENGINEERING
UNIVERSITY OF MICHIGAN

DISSERTATION DEFENSE

Christopher Clarke

Towards Enhanced Human-AI Interaction: A Holistic Approach to Personalization in Natural Language Processing

Tuesday, December 3, 2024
6:00pm – 8:00pm
3725 Beyster
Hybrid – Zoom

ABSTRACT: Traditional NLP approaches lean towards developing universal models designed to cater to a wide spectrum of tasks and user demographics. These models prioritize broad applicability, effectively homogenizing user interactions into a one-size-fits-all framework. While practical for many common applications, this one-size-fits-all approach often fails to address the rich tapestry of human diversity and individual needs needed to build truly interactive systems.

This dissertation argues for a paradigm shift towards personalized NLP enabling systems that can adapt to individual users' preferences, needs, and contexts. Personalization is a critical aspect of human-AI interaction, as it enables AI systems to better understand and cater to individual users' unique requirements. In this dissertation, I demonstrate how personalization can be integrated into modern NLP systems to enhance user experiences from a holistic perspective. I showcase a series of works for personalized NLP that encompass four key aspects: 1) Approaches for incorporating user perspective, 2) Adaptive Learning & Feedback for Personalization, 3) Interactive Interfaces for Personalization, and 4) Datasets & Benchmarks for Personalization. First, I explore techniques for incorporating user perspectives into large language models (LLMs), enabling models to better understand user preferences and needs. Secondly, I investigate adaptive learning and feedback mechanisms that allow LLMs to adapt to user feedback and improve over time. Thirdly, I explore interactive interfaces that facilitate user-AI collaboration, enabling users to provide feedback and guidance. Lastly, I discuss the importance of datasets and benchmarks for evaluating personalized LLMs, highlighting the need for diverse and representative datasets to ensure the robustness and generalizability of personalized models.

CHAIR: Prof. Jason Mars and Prof. Lingjia Tang