




DISSERTATION DEFENSE



Shengpu Tang

Towards Clinically Applicable
Reinforcement Learning

Monday, June 3, 2024

9:00am – 11:00am

3941 Beyster

Hybrid – [Zoom](#) Passcode: UM-CSE

ABSTRACT: In healthcare, clinicians constantly make decisions about when and how to treat each patient. These decisions are based on medical training and clinical experience, but they may not always be optimal. Reinforcement learning (RL) offers an appealing framework to create decision support tools that could assist clinicians in selecting appropriate treatments. Despite recent successes of RL in other domains such as games and chat-bots, existing approaches are often incompatible with clinical decision-making problems in healthcare. In this thesis, we develop several methods to improve the applicability of RL, in particular, offline RL, for clinical settings.

First, to enable clinician-in-the-loop decision-making, we propose to learn *set-valued policies* to capture multiple actions that achieve similar outcomes (e.g., survival). By providing users with near-equivalent choices instead of a single best action, we allow clinicians to incorporate additional contextual knowledge (e.g., patient preferences, cost and availability) when making final treatment decisions. Second, when multiple types of treatments need to be selected simultaneously, a standard approach for policy learning is inefficient as it considers a combinatorially large number of actions. To better leverage the structure of *factored action spaces*, we propose a linear decomposition of the Q-function. We show both theoretically and empirically that our approach makes more efficient use of limited data, leading to better policies. Third, we present a *practical pipeline of model selection for offline RL* - an issue often skirted by past work. Based on an in-depth analysis of various methods for offline policy evaluation (OPE), we propose a *two-stage* procedure to balance the trade-offs between accuracy and computational cost of different OPE methods. Finally, to bridge the gap between offline and online evaluation, we propose a novel *semi-offline* evaluation framework that augments the observed patient trajectories with *counterfactual annotations* of unobserved treatments. We demonstrate the advantages of this approach through theoretical analyses and empirical simulations. Our framework could be implemented safely with minimal disruption to improve confidence in newly learned policies. By recognizing and addressing the unique challenges at the intersection of RL and healthcare, the proposed approaches in this thesis can help enable intelligent decision support systems that are applicable to real-world clinical settings.

CHAIR: Prof. Jenna Wiens