Transforming Healthcare Decision Making Using Artificial Intelligence

Shengpu Tang
Emory University
Department of Computer Science
shengpu.tang@emory.edu

Decision making is at the core of healthcare: clinicians constantly make complex decisions that span diagnosis, treatment, care coordination, and resource allocation. Yet, human decisions are never perfect, leading to suboptimal patient care. My research uses AI to augment and improve decision-making in healthcare, following a synergistic approach that combines novel AI methods with practical, real-world implementation. Here, I will explore two key themes: *Application-Inspired AI Innovations*, focused on novel AI methods grounded in practical healthcare problems; and *Path to Deployment and Impact*, which addresses AI integration into clinical workflows for real-world improvements.

Part 1. Application-Inspired AI Innovations. Among AI methods, reinforcement learning (RL) offers a promising framework to create decision support tools that could assist clinicians in selecting appropriate treatments. Despite RL's success in areas such as gaming and chat-bots, many existing approaches not well-suited to the unique challenges of clinical decision-making problems in healthcare. In my recent work, I have developed several methods to improve the applicability of RL, particularly offline RL, for clinical settings (Tang et al. 2020a; Tang and Wiens 2021; Tang et al. 2022; Tang and Wiens 2023; Rebello et al. 2023). For example, by leveraging compositionality of real-world problems, such as *factored action spaces* (Tang et al. 2022; Rebello et al. 2023), we show learning and evaluation can be made more data-efficient. To support clinician-in-the-loop decision making, I introduced new algorithms to learn *set-valued policies*, which capture near-equivalent actions that achieve similar outcomes (e.g., survival) (Tang et al. 2020a). Additionally, to improve offline RL policy evaluation before online deployment, I proposed model selection pipelines (Tang and Wiens 2021) and techniques to incorporate human annotations of counterfactual outcomes (Tang and Wiens 2023).

Part 2. Path to Deployment and Impact. While numerous clinical prediction models have been built, such as in my own work (Tang et al. 2020b; Tang, Hur et al. 2021), most have yet to impact decision making in hospitals. In this part, I will share my experience deploying two such models at Michigan Medicine (MM). The first model, M-CURES (Tang, Kamran et al. 2022), was developed during the COVID-19 pandemic to predict early deterioration in emergency department patients. After extensive validation using data from 13 U.S. hospitals, we deployed M-CURES within MM's electronic health record system, and the rapid response team now uses the model in their daily rounding decisions. In another notable project tackling *C. difficile* infections, I designed an AI-guided infection prevention protocol to reduce both pathogen exposure and host susceptibility. A three-year quasi-experimental study at MM demonstrated a tangible impact of AI in reducing infections and improving antimicrobial stewardship (Tang et al. 2024).

Future Directions. I will discuss promising directions with open questions in areas such as *compositionality*, *human-centered designs*, and *ML deployment pipelines* in hospitals – all in the context of *clinically relevant problems*.

References

Rebello, A. P.; Tang, S.; Wiens, J.; and Parbhoo, S. 2023. Leveraging Factored Actions for Off-Policy Evaluation. *ICML workshop: Counterfactuals in Minds & Machines*.

Tang, S.; Hur, J.; et al. 2021. Predicting postoperative opioid use with machine learning and insurance claims in opioid-naïve patients. *The American Journal of Surgery*.

Tang, S.; Kamran, F.; et al. 2022. Early identification of patients admitted to hospital for covid-19 at risk of clinical deterioration. *BMJ*.

Tang, S.; Makar, M.; Sjoding, M.; Doshi-Velez, F.; and Wiens, J. 2022. Leveraging Factored Actions for Efficient Offline Reinforcement Learning in Healthcare. *NeurIPS*.

Tang, S.; Modi, A.; Sjoding, M.; and Wiens, J. 2020a. Clinician-in-the-Loop Decision Making: Reinforcement Learning with Near-Optimal Set-Valued Policies. *ICML*.

Tang, S.; and Wiens, J. 2021. Model Selection for Offline RL: Practical Considerations for Healthcare. *MLHC*.

Tang, S.; and Wiens, J. 2023. Counterfactual-Augmented Importance Sampling for Semi-Offline Evaluation. *NeurIPS*.

Tang, S.; et al. 2020b. Predicting aGVHD Using ML and Longitudinal Vital Signs. *JCO Clinical Cancer Informatics*.

Tang, S.; et al. 2024. Targeting *C. difficile* Infection Prevention Efforts with AI. *IDWeek 2024*.