

Scalable, Sustainable, Generalizable, and Responsible AI for Public Sector

Zheyuan Ryan Shi

Department of Computer Science, University of Pittsburgh, Pittsburgh, PA
ryanshi@pitt.edu

Talk Outline

AI for public sector research is about using AI to tackle the numerous challenges faced by public sector organizations when they are out there making our world a better place. AI for public sector research is use-inspired research. It differs from traditional AI research first and foremost in its key objective being measurable societal impact. AI for public sector research contributes to the computing community by proposing new problem models, raising complexities that challenge abstractions which often leads to new methodologies, and introducing new contexts for evaluation. However, fulfilling this promise is easier said than done. This talk consists of three parts about our preliminary work in our long-term quest to make AI for public sector **operationally scalable, financially sustainable, technically generalizable, and socially responsible**. We will cover (1) a concrete AI for public sector project from problem scoping to field trials and deployment, (2) a generalizable algorithm applicable to various public sector domains, and (3) an overview of our work in a wide variety of applications.

Part 1: Volunteer Engagement for Food Security Food waste and food insecurity arise simultaneously in many parts of the world. Food rescue organizations (FR) handle real-time food donations and match them to organizations that serve underprivileged communities. FRs rely on volunteers to transport the food, yet they also bring uncertainty. With 412 Food Rescue (412FR), we formulated a data-driven optimization problem to suggest when FRs should send push notifications to advertise the rescue (Shi et al. 2020). We then designed an online rescue-specific recommender algorithm to advertise each rescue to volunteers through push notifications (Shi, Lizarondo, and Fang 2021). We also designed large language model-based algorithm to predict and present rescue difficulty to volunteers (Shi et al. 2024). A randomized controlled trial with 412FR showed that our algorithm significantly improved the notification hit rate by 44% and rescue claim rate by 9%. The algorithm has now been fully integrated into 412FR’s management platform, engaging over 45,000 volunteers, helping the food insecure population served by over 8,000 community organizations, facilitating over 6 million meals delivered per year.

Part 2: Bandit Data-Driven Optimization As we worked on these applied AI for public sector projects, we realized that many of them have shared technical challenges. As one such example, we proposed Bandit Data-Driven Optimization as the first iterative prediction-prescription framework to address four major pain points in AI for public sector applications (Shi et al. 2022): small data, uniform embedded interventions, incomplete policy objectives, and unintended consequences. We proposed a no-regret algorithm called PRedict-then-Optimize with Optimism in Face of uncertainty, and showed that it consistently outperforms bandit baseline by a margin of 15%-300%.

Part 3: The Big Picture and Future Directions We will conclude the talk with an overview of our work and future directions. Alongside completing the first survey of AI for good (Shi, Wang, and Fang 2020), we have worked on other applications such as environmental conservation with successful real-world deployment, and made methodological contributions in game theory and sequential decision making. Since starting the faculty position in 2024, we have established collaborations with five different nonprofit organizations working in agriculture, healthcare, education, and food security. With these efforts, we hope to prove in the near future that it is possible to achieve scalable, sustainable, generalizable, and responsible AI for public sector research.

References

- Shi, Z. R.; Lizarondo, L.; and Fang, F. 2021. A recommender system for crowdsourcing food rescue platforms. In *Proceedings of the Web Conference 2021*, 857–865.
- Shi, Z. R.; Wang, C.; and Fang, F. 2020. Artificial intelligence for social good: A survey. *arXiv preprint arXiv:2001.01818*.
- Shi, Z. R.; Wu, Z. S.; Ghani, R.; and Fang, F. 2022. Bandit Data-Driven Optimization for Crowdsourcing Food Rescue Platforms. In *AAAI*, volume 36, 12154–12162.
- Shi, Z. R.; Yuan, Y.; Lo, K.; Lizarondo, L.; and Fang, F. 2020. Improving Efficiency of Volunteer-Based Food Rescue Operations. *IAAI*.
- Shi, Z. R.; Zhi, J.; Zeng, S.; Zhang, Z.; Kapoor, A.; Hudson, S.; Shen, H.; and Fang, F. 2024. Predicting and Presenting Task Difficulty for Crowdsourcing Food Rescue Platforms. In *Proceedings of the Web Conference 2024*.