Engineering Robust Decision-Support Systems for Online Resource Allocation under Distribution Shifts and Confounding

The objective of this proposal is to develop new analytical and algorithmic foundations for engineering reliable online resource allocation systems learned from offline observational data collected under confounding and deployed under distribution shift. Many critical service infrastructures — including housing, healthcare, and education — operate as complex decision-driven environments in which individuals and scarce resources arrive over time and must be matched in real time using standardized assessments, administrative records, and expert judgment. Allocation rules are learned from historical data and deployed in evolving institutional settings. In such systems, populations receiving different interventions differ systematically, and key drivers of both allocation decisions and downstream outcomes are often unobserved, inducing selection bias and confounding. Institutional changes further generate distribution shifts that can degrade performance, leading to increased congestion, under-utilization of capacity, and erosion of stakeholder trust. Ensuring operational stability under these conditions is a key challenge in modern data-driven service systems.

This proposal addresses the following scientific question: **Can we design optimization-based, causal decision-support systems for online resource allocation—learned from offline observational data—that maintain provable system-level performance guarantees and operational stability under distribution shifts and unobserved confounding?**

To answer this question, the project develops: *(i)* mathematical models of distribution shift and unobserved confounding calibrated using domain expertise; *(ii)* computationally efficient methods for bounding resource effectiveness under these uncertainties; and *(iii)* scalable algorithms for learning prioritization rules and static and dynamic allocation policies with explicit robustness guarantees. The resulting frameworks integrate (distributionally) robust optimization and causal machine learning to produce decision-support tools with quantifiable deployment guarantees. By addressing confounding and non-stationarity, the project advances foundational methods for engineering data-driven operational systems and contributes to robust causal inference, robust machine learning, robust queuing theory, and robust Markov Decision Processes (MDPs).

The central application guiding this research is the allocation of scarce housing to individuals experiencing homelessness. This setting provides a rich engineered service testbed characterized by high-dimensional, noisy administrative data, severe resource constraints, evolving institutional processes, and substantial human discretion — which induce confounding and distribution shift. Through longstanding collaborations with the Los Angeles Homeless Services Authority (LAHSA) and the Missouri Balance-of-State (BoS) Continuum of Care (CoC), the team is embedded in real-world service contexts. The project will convene advisory boards of policymakers, service providers, and people with lived experience to inform its assumptions and models and to ensure that resulting engineering frameworks remain grounded in operational realities. These partnerships establish a clear pathway for translating methodological advances into decision-support tools that can be responsibly integrated into operational workflows and inform allocation practices.

The proposed methods apply to broad classes of engineered service and infrastructure systems. Through her leadership of the Center for AI in Society (CAIS), the PI will pursue applications in substance use prevention and biodiversity conservation. The project will translate methodological advances into deployable computational tools emphasizing usability and extensibility, strengthening adoption of the methods and supporting reliable, adaptable decision infrastructures. Educationally, the project will train students at the intersection of optimization, machine learning, and engineered service systems. Research activities will be integrated into the NSF National Research Traineeship (NRT) Operations Research and Artificial Intelligence (ORAI) program, which the PI co-directs, fostering a workforce capable of designing reliable data-driven decision infrastructures.