Tuesday 23 July 2024

09:00-10:30 Invited Session 4 (Main Room)

Optimal individualized treatment rules (Chairs: Giorgos Bakoyannis, Rodolphe Thiebaux)

Policy learning with distributional welfare

Sukjin Han (University of Bristol, UK)

In this paper, we explore optimal treatment allocation policies that target distributional welfare. Most literature on treatment choice has considered utilitarian welfare based on the conditional average treatment effect (ATE). While average welfare is intuitive, it may yield undesirable allocations especially when individuals are heterogeneous (e.g., with outliers)---the very reason individualized treatments were introduced in the first place. This observation motivates us to propose an optimal policy that allocates the treatment based on the conditional *quantile of individual treatment effects* (QoTE). Depending on the choice of the quantile probability, this criterion can accommodate a policymaker who is either prudent or negligent. The challenge of identifying the QoTE lies in its requirement for knowledge of the joint distribution of the counterfactual outcomes, which is generally hard to recover even with experimental data. Therefore, we introduce minimax policies that are robust to model uncertainty. A range of identifying assumptions can be used to yield more informative policies. For both stochastic and deterministic policies, we establish the asymptotic bound on the regret of implementing the proposed policies. In simulations and two empirical applications, we compare optimal decisions based on the QoTE with decisions based on other criteria. The framework can be generalized to any setting where welfare is defined as a functional of the joint distribution of the potential outcomes.