

Distributionally Robust Policy Learning



Abstract: Policy learning using historical observational data is an important problem that has found widespread applications. Examples include selecting offers, prices, advertisements to send to customers, as well as selecting which medication to prescribe to a patient. However, existing literature rests on the crucial assumption that the future environment where the learned policy will be deployed is the same as the past environment that has generated the data -- an assumption that is often false or too coarse an approximation. In this paper, we lift this assumption and aim to learn a distributionally robust policy with incomplete observational data. We first present a policy evaluation procedure that allows us to assess how well the policy does under the worst-case environment shift. We then establish a central limit theorem type guarantee for this proposed policy evaluation scheme. Leveraging this evaluation scheme, we further propose a novel learning algorithm that is able to learn a policy that is robust to adversarial perturbations and unknown covariate shifts with a performance guarantee based on the theory of uniform convergence. Finally, we empirically test the effectiveness of our proposed algorithm in synthetic datasets and demonstrate that it provides the robustness that is missing using standard policy learning algorithms. We conclude the paper by providing a comprehensive application of our methods in the context of a real-world voting dataset.

Bio: Nian Si recently joined the University of Chicago, Booth School of Business as a postdoctoral principal researcher. He finished PhD in Operations Research in the Department of Management Science and Engineering (MS&E) at Stanford University. He received a B.A. degree in Economics and a B.S. degree in Mathematics from Peking University in 2017. His research lies at the interface of applied probability, simulation, and machine learning and he develops methods for trustworthy decision making. He is also interested in real-world problems arising from online platforms and some of his proposed solutions have been deployed in the industry.