(Fast) Rates for Estimating Optimal Individualized Treatment Rules

Alex Luedtke, University of Washington

This talk provides guarantees on the performance of several strategies for estimating optimal treatment rules. Here, a treatment rule is defined as a rule that takes as input an individual-level covariate and outputs a treatment decision. An optimal rule maximizes some user-defined value function, such as the counterfactual mean outcome if that rule were implemented in the population. The regret of a rule is defined as the difference between the value of the rule and the value of an optimal rule. I give a margin-free result showing that the regret decay for using a certain class of supervised machine learning procedures is second-order, with regret decaying at a faster rate than the standard error of an efficient estimator of the value of an optimal rule. I also present a result giving guarantees on the regret decay of optimal rule estimators for the case that the data is generated by local perturbations of a fixed distribution -- this result is useful for describing the finite-sample performance of optimal treatment rule estimators. Finally, I give a result from the classification literature that shows that faster regret decay is possible via an alternative estimation strategy if a margin condition holds.