

Data Efficient Causal Effect Estimation

Adith Swaminathan, Microsoft Research

While several methods for estimating Individual Treatment Effects have been recently suggested, these methods assume no constraints on the availability of data at the time of deployment. This assumption is unrealistic in settings where data about a test case has to be collected through an expensive process in order to predict the treatment effect. We'll highlight one recent approach (presented at AAAI'19) that achieves data efficiency at test time by distinguishing between causal discovery during training and effect prediction during testing. This is joint work with Maggie Makar and Emre Kiciman.