Distribution Generalization with Instrumental Variables

Niklas Pfister

University of Copenhagen

Abstract

Causal models can provide robust predictions even under distributional shifts. This observation has lead to the development of various methods that use some level of causal learning to solve prediction tasks. In this talk, we will consider a class of such methods based on instrumental variable (IV) models. IV models allow us to identify a causal function between covariates X and a response Y, even in the presence of unobserved confounding. However, in many practical prediction tasks the causal function is not fully identifiable. We will investigate under which assumptions these models have show that even when the causal model is not identifiable it is still possible to learn a predictive function that similar to the causal function is able to guard against a large class of distribution shifts.

We prove that the proposed estimator is invariant to distributional shifts on the instruments and worst-case optimal whenever these shifts are sufficiently strong. These results hold even in the under-identified case where the instruments are not sufficiently rich to identify the causal function.