Jonas Peters
University of Copenhagen

Abstract

Title: Causality in Data Science

Causality enters data science in different ways. Often, we are interested in knowing how a system reacts under a specific intervention, e.g., when considering gene knock-outs or a change of policy. The goal of causal discovery is to learn causal relationships from data. Other practical problems in data science focus on prediction. But as soon as we want to predict in a scenario that differs from the one which generated the available data (we may think about a different country or experiment), it might still be beneficial to apply causality related ideas. We present assumptions, under which causal structure becomes identifiable from data and methods that are robust under distributional shifts. No knowledge of causality is required.