
Causal Inference using Invariant Prediction^{*}

Jonas Peters

Abstract

Why are we interested in the causal structure of a data-generating process? In a classical regression problem, for example, we include a variable into the model if it improves the prediction; it seems that no causal knowledge is required. In many situations, however, we are interested in the system's behavior under a change of environment, i.e., under a different distribution. Here, causal structures become important because they are usually considered invariant under those changes and can therefore be used to answer questions about this new distribution. For example, a causal prediction (which uses only direct causes of the target variable as predictors) remains valid even if we intervene on predictor variables or change the whole experimental setting. We propose a method that exploits the invariance principle when data from different environments are available. This talk concentrates on ideas and concepts and does not require any prior knowledge about causality.

^{*}Machine Learning External Seminar, Gatsby Unit, January 28, 2015.