

Learning causal mechanisms

Bernhard Schölkopf¹

Abstract: In machine learning, we use data to automatically find dependences in the world, with the goal of predicting future observations. Most machine learning methods build on statistics, but one can also try to go beyond this, assaying causal structures underlying statistical dependences. Can such causal knowledge help prediction in machine learning tasks? We argue that this is indeed the case, due to the fact that causal models are more robust to changes that occur in real world datasets. We discuss implications of causal models for machine learning tasks, focusing on an assumption of ‘independent mechanisms’, and discuss an application in the field of exoplanet discovery.

¹ Max-Planck-Institut für Intelligente Systeme, Tübingen bs@tuebingen.mpg.de