Jonas Peters, MPI Tübingen, Germany

Invariances and Causality

Abstract: In causal inference, we are interested in the system's behavior under an intervention: what happens to gene B, for example, if we delete gene A? If gene A is causal for gene B, then predicting B from A works equally well in the intervened and in the non-intervened scenario; this does not hold if gene B is causal for gene A, for example.
In our talk, we discuss two implications of this well-known observation:

[1] Causal models are more stable against a change of environments and may be used for problems in domain adaptation.

[2] If we are given data from different environments, we can look for such invariant models and draw conclusions about the underlying causal structure.