

# Causal Representation Learning across Multiple Environments

[Phillip Lippe](#)

Identifying causal variables and their relations from high-dimensional observations is of great interest in applications like robotics and embodied AI. As it has been shown that causal variables are not necessarily identifiable in the most general setting, recent research focused on using observations from multiple, slightly perturbed environments (e.g. by interventions) to enable identifiability. In this talk, we present the current state and open challenges of causal representation learning in multi-environment settings. We will first review existing methods, including our work on CITRIS and BISCUIT, which leverage intervention-based data to learn causal representations. However, real-world scenarios often involve dynamic environments with varying causal structures and observation functions. To address this, we will discuss our ongoing research on learning causal representations that generalize to unseen environments. Specifically, we will explore under which settings causal representations can identify multiple environments from samples of a joint observation distribution. Furthermore, we explore the usage of object-centric encodings to enable zero-shot generalization to novel, compositional environments.