

2018 Research Interest/Project Ideas

Elias Bareinboim

<http://causalai.net>

Explainability and Causal Machine Learning

Causality is central to scientific inquiry across many disciplines. Researchers are usually interested not only in knowing how two events are correlated, but also whether one causes the other and, if so, how. In general, scientific practice seeks not just a surface description of the observed data, but rather deeper explanations, ranging from predicting the effect of new interventions to estimating retrospective counterfactuals. Ultimately, scientists aim to understand the mechanisms underlying an observed phenomenon. With the unprecedented accumulation of data, researchers are becoming increasingly aware that traditional techniques must be enriched with two additional ingredients in order to construct such explanations: 1. the ability to integrate data from multiple, heterogeneous sources, and, 2. the ability to distinguish causal from associational relationships. In the last decades, machine learning has focused, almost invariably, on issues of correlations, neither causation nor explanation. In this project, we develop a mathematical framework for assisting the data scientist in leveraging multiple heterogeneous, non-causal datasets to construct causal explanations that are human-interpretable, robust, and consistent with the scientific method.