When the problem becomes richer than supervised learning. A real-world use of causality in machine learning.

Onno Zoeter (Booking.com)

The classic supervised learning problem that is taught in machine learning courses and is the subject of many machine learning competitions is often too narrow to reflect the problems that we face in practice. Historical datasets typically reflect a combination of a source of randomness (for example customers making browsing and buying decisions) and a controlling mechanism such as a ranker or highlighting heuristics (badges, promotions, etc.). Or there might be a selection mechanism (such as the decision to not accept transactions with high fraud risk) that influences the training data. A straightforward regression approach would not be able to disentangle the causal influence of the controller and the phenomenon under study. As a result it risks making incorrect predictions as the controller is changed. In practice however, such problems are typically treated as a classic regression problem in a first iteration and attempts to identify and correct for these complications come as afterthoughts or are not undertaken at all. Ideally there is a rigorous and flexible formalism that captures the correct framing of the problem from the very start, accompanied by a set of practical algorithms that work well in practice for each of the identified cases. In our initial set of successes, structural causal models have proven to be an effective language to express the understanding of the phenomena and to make accurate causal predictions for changes just in the part that is under control, e.g. the ranker, the acceptance policy, etc. This overall research objective is the main goal of the Mercury Machine Learning Lab, one of the labs within Booking AI Research. The Mercury lab is a collaboration between the University of Amsterdam, the Technical University of Delft and Booking.com. It brings together the fields of information retrieval, causality and reinforcement learning where the topic is studied under different names e.g. off-line evaluation, transportability and s-recoverability and off-policy learning. This presentation will sketch the problem, highlight some of the theoretical results so far and describe a significant real-world application.