

Exposing the Causal Structure of Processes by Learning CP-Logic Programs

Hendrik Blockeel

K.U. Leuven, Belgium and Leiden University, the Netherlands

Abstract. Since the late nineties there has been an increased interest in probabilistic logic learning, an area within AI that combines machine learning with logic-based knowledge representation and uncertainty reasoning. Several different formalisms for combining first-order logic with probability reasoning have been proposed, and it has been studied how models in these formalisms can be automatically learned from data.

This talk starts with a brief introduction to probabilistic logic learning, after which we will focus on a relatively new formalism known as CP-logic. CP-logic stands for “causal probabilistic logic”. It is a knowledge representation formalism that allows us to write down rules that indicate that a certain combination of conditions may cause certain effects with a particular probability (e.g., tossing a coin may cause a result of heads or tails, each with 50% probability). Besides the fact that this formalism is interesting for knowledge representation in itself, it also offers interesting opportunities from the machine learning point of view. Indeed, given the semantics of these CP-logic programs, learning them from data amounts to extracting probabilistic causal influences from data. We will discuss recent research on learning CP-logic programs, including: algorithms for learning them; how they relate to graphical models; and applications of learning CP-logic programs.