

Relational Machine Learning

Filip Železný

Czech Technical University in Prague, Czech Republic

Abstract. I will explain the main concepts of relational machine learning, or more precisely, those parts of it employing logic as the knowledge-representation formalism. The talk will not cover other relational approaches such as graph-mining. I will follow what I consider the three main stages of the field's historical development. First, I will visit the roots of relational learning lying in the area of inductive logic programming. Here, one learns logical theories from examples, formalizing the problem as search in a clause subsumption lattice. A newer stream of research called statistical relational learning extended the logical underpinnings with probabilistic inference. I will illustrate this with an example of a logical graphical probabilistic model. Most recently, relational learning has received a new impetus from the current revival of (deep) neural networks. I will exemplify some promising crossovers of the two fields, including the paradigm of Lifted Relational Neural Networks conceived in my lab.