
Preface



IT has been a great pleasure to be asked to write the preface for the book based on Kristian Kersting's thesis. There is no doubt in my mind that this is a remarkable and outstanding piece of work.

In his thesis Kristian has made an assault on one of the hardest integration problems at the heart of Artificial Intelligence research. This involves taking three disparate major areas of research and attempting a fusion among them. The three areas are: Logic Programming, Uncertainty Reasoning and Machine Learning. Every one of these is a major sub-area of research with its own associated international research conferences. Having taken on such a Herculean task, Kristian has produced a series of widely published results which are now at the core of a newly emerging area: Probabilistic Inductive Logic Programming. The new area is closely tied to, though strictly subsumes, a new field known as "Statistical Relational Learning" which has in the last few years gained major prominence in the American Artificial Intelligence research community.

Within his thesis Kristian makes several major contributions, many of which have already been published in refereed conference and journal papers. Firstly, Kristian introduces a series of definitions which circumscribe the new area formed by extending Inductive Logic Programming to the case in which clauses are annotated with probability values. This represents a new and powerful framework which supersedes a number of influential papers and research areas in Artificial Intelligence. Secondly, Kristian introduces Bayesian Logic Programs (BLPs). These represent an elegantly defined lifting of Judea Pearl's Bayesian networks to the logic programming level. Since Kristian's introduction of BLPs, a number of results indicate that BLPs generalise many previously defined representations, not the least of which are Bayesian networks, Logic Programs, Probabilistic Relation Models and Stochastic Logic Programs. Next Kristian investigates the approach of Learning from proofs. This is an interesting new learning framework which is the first to go beyond the two standard semantic frameworks of Inductive Logic Programming.

Kristian then looks at the problem of upgrading HMMs to logical HMMs. Hidden Markov Models (HMMs) are one of the most widely used machine learning technologies in Statistical Linguistics and Bioinformatics, and allow the representation of probabilistic finite automata. Kristian has upgraded standard HMMs to allow relational descriptions to be included within the description of the automata. The three standard HMM estimation algorithms are also upgraded. He has demonstrated the power of such representations using biological predictive modelling problems, and shown performance increases over alternative approaches.

Kristian next considers the issue of upgrading Fisher Kernels to Relational Fisher kernels. Fisher kernels have been widely used within statistics and more recently

in support vector machines. Building on his previous approaches involving lifting propositional representations Kristian shows how relations can be usefully included within this context. The approach was empirically tested on protein fold prediction and shown to have high predictive accuracy relative to logical HMMs.

Lastly, Kristian introduces Markov decision programs. As a final demonstration of his general approach Kristian shows how temporal descriptions involving action can be introduced by lifting Markov decision processes to logical Markov decision programs. Kristian demonstrates how these can be learned using relational reinforcement algorithms which he tests empirically in a Blocks World setting.

In summary, this thesis represents an extremely powerful and wide-ranging study which has made strong contributions right across the intellectual landscape. Both Kristian and his thesis supervisor Luc De Raedt, should be highly commended for this important contribution.

London, July 2006

Stephen H. Muggleton