Advances in ILP Theory and Implementations

Stephen Muggleton
Department of Computer Science,
University of York,
Heslington, York, YO1 5DD,
United Kingdom.

No Institute Given

Abstract. A strong linkage exists between advances in applications, implementations and theory within Inductive Logic Programming (ILP). Early ILP systems, such as FOIL, Golem and LINUS learned single predicate definitions from positive and negative examples and extensional background knowledge. They also employed strong learning biases such as ij-determinacy. Although these systems found a number of applications, they had problems in areas such as molecular biology and natural language learning.

General mechanisms for inverting entailment have now been developed which support the use-of non-ground background knowledge, and the revision of multiple inter-related predicates. ILP theory results concerning complete refinement graph operators now allow efficient admissible searches. The absolute requirement for negative examples (rare within natural language domains) has been eased by Bayesian analysis of learning from positive-only examples. Bayesian approaches have also supported sample complexity analysis of predicate invention within the framework of repeat learning. In this framework it is assumed that the learner's prior is not equivalent to the distribution from which the teacher is sampling targets. By providing a series of sessions the learner is able to update the initial prior by adding and deleting background predicates. Within the Bayesian framework stochastic logic program representations have been used to estimate the distribution of examples over the instance space. Stochastic logic programs are a generalisation of hidden Markov models and stochastic grammars.

Apart from a few special cases PAC-learning results have been largely negative for ILP. This is in large part due to the fact that testing satisfiability is intractrable for most interesting subsets of first-order Horn logic. The development of Bayesian approaches to ILP supported the development of U-learnability, which allows classes of distributions over the hypotheses. Here it was shown that for any exponential-decay distribution the class of time-bounded logic-programs is polynomially U-learnable. The use of such bounds on proof depth is common within ILP systems. Although logically impure, this approach allows general-purpose flexible representations, while maintaining termination guarantees.