
Learning the Time Complexity of Logic Programs

Stephen Muggleton

SHM@DOC.IC.AC.UK

Computational Bioinformatics Laboratory
Department of Computing
Imperial College, London, UK

One of the key difficulties in machine learning recursive logic programs is associated with the testing of examples. Inefficient hypotheses, though usually of least interest to the learner, take more time to test. Almost by definition, efficient learners require a bias towards low-complexity hypotheses. However, it is unclear how such a bias can be implemented. To these ends, in this presentation we address the problem of recognising inefficient logic programs. To the author's knowledge, this problem has not been considered previously in the literature. A successful approach to recognition of inefficient logic programs should bring the prospects of effective machine learning of recursive logic programs closer. In general the problem of recognising the time complexity of an arbitrary program is incomputable since halting is undecidable. However, partial solutions cannot be ruled out. In this presentation we provide an initial investigation of the problem based on developing a framework for machine learning higher-order patterns associated with various complexity orders.