

Guest editors' introduction: Special issue on Inductive Logic Programming (ILP-2008)

Filip Železný · Nada Lavrač

Received: 29 April 2009 / Revised: 29 April 2009 / Accepted: 30 April 2009 /

Published online: 12 June 2009

Springer Science+Business Media, LLC 2009

The 18th International Conference on Inductive Logic Programming was held in Prague, September 10–12, 2008. Apart from four invited talks and a tutorial, the presented peer-reviewed papers consisted of 20 submissions accepted as long papers for the main technical track, and 21 papers accepted for the work-in-progress track. The long papers can be found in Volume 5194 of Springer's Lecture Notes in Artificial Intelligence series. Based on the reviews of the papers as well as their presentations at the conference, the ILP-2008 PC chairs invited the authors of 7 selected papers to submit a revised and significantly extended version for this special issue. As a result of an additional peer-review adopting the journal criteria, 5 of these papers were finally accepted.

While the last MLJ special issue on ILP emphasized approaches combining relational representations with probabilistic modeling and inference, the current issue reflects recent advances in the logical and graphical foundations of relational learning as well as novel applications of ILP in theory revision and natural language processing.

A number of first-order logic frameworks have been proposed by the ILP community and used for concept learning. The most popular ones are known as learning from interpretations, learning from entailment, and learning from satisfiability. In the article “Brave Induction: a Logical Framework for Learning from Incomplete Information”, Chiaki Sakama and Katsumi Inoue show that none of these frameworks is directly suited to elegantly deal with a certain type of incompleteness in training data. Inspired by *brave inference* known in nonmonotonic logics, the authors propose a novel logical setting termed *brave induction* for learning first-order clauses. They rely on minimal-model semantics instead of using the classical entailment relation to define when an example is covered by a hypothesized clause and background knowledge. The proposed framework is also applied to induction from non-monotonic logic programs.

F. Železný (✉)

Department of Cybernetics, Faculty of Electrical Engineering, Czech Technical University in Prague,
Prague, Czech Republic

e-mail: zelezny@fel.cvut.cz

N. Lavrač

Department of Knowledge Technologies, Jožef Stefan Institute, Ljubljana, Slovenia
e-mail: nada.lavrac@ijs.si

In the area of relational machine learning, ILP was always distinguished as a framework capable of precisely structuring its non-trivial hypothesis spaces by generality. This is thanks to the achieved exhaustive theoretical elaboration of clause subsumption lattices and of the corresponding refinement operators that allow searching it. However, a significant gap has emerged in the last years between theory and practice with the growing popularity of ILP systems which, for the sake of efficiency, limit their search space by a *bottom clause*. These systems search in strict substructures of the conventional subsumption lattice. The properties of such bounded spaces have not been hitherto systematically described. A theoretical insight into such spaces is the main contribution of this issue's second paper, authored by Alireza Tamaddoni-Nezhad and Stephen Muggleton, and titled “The lattice structure and refinement operators for the hypothesis space bounded by a bottom clause.”

Exploitation of the bottom clause has been so far constrained to standard concept learning settings, specifically in the framework of learning from entailment. The paper “Using the Bottom Clause and Mode Declarations on FOL Theory Revision from Examples” authored by Ana Luísa Duboc, Aline Paes and Gerson Zaverucha, demonstrates the benefits of using a bottom clause in the *theory revision* setting, which is likely to gain significant interest due to the currently growing knowledge bases in fields like bioinformatics, paralleled by the accumulation of empirical data, e.g., from high-throughput technologies. Specifically, the authors endow the classical system FORTE with revision operators constrained by a bottom clause. They show considerable improvements in the efficiency of the revision process without compromising on classification accuracy in real-life biochemical benchmark problems.

In ILP's early years, biochemistry was a domain of its killer applications. Through such applications leading to results unmatched at the time by competing methods, a research field gains perhaps even more significance than through theoretical rigor. In the paper “An Investigation into Feature Construction to Assist Word Sense Disambiguation” included in this issue, Lucia Specia et al. test the hypothesis that natural language processing could be a next killer application domain for ILP. More precisely, the authors take the challenge of word sense disambiguation, a long standing problem where a significant improvement in accuracy can be—by common agreement—achieved only by principled incorporation of background knowledge. For this sake, the authors pragmatically combine state-of-the-art ILP approaches (feature construction, randomized search) with a state-of-the-art classification strategy (support vector machines). The role of ILP is thus primarily to infuse diverse sources of background knowledge into constructed features. The encouraging results indicate that ILP can be a factor enabling significant progress in word sense disambiguation.

ILP algorithms pay a high runtime complexity tax for dealing with the expressive languages of first-order and relational logic. ILP researchers have, therefore, been looking for islands of tractability, that is, suitable sublanguages enabling efficient completion of rudimentary tasks such as checking the compliance of a hypothesis with a training example. The paper “Learning Block-Preserving Graph Patterns and its Application to Data Mining” by Hitoshi Yamasaki et al. presents such an island of relational representations. Here, the examples are a special kind of graphs enabling a polynomial-time isomorphism check. The authors introduce a special kind of graph pattern efficiently learnable from such examples in the inductive inference framework, using specially designed refinement operators. To show that this island is large enough to host real-life significant learning tasks, a standard benchmark dataset of chemical compounds is considered, of which a vast majority indeed complies with the chosen representation bias.

We are grateful to the authors, the reviewers, and to the publisher Springer for giving us the opportunity to compile this reflection of current research in inductive logic programming.