

Guest Editors' introduction

Special issue on inductive logic programming (ILP 2010)

Paolo Frasconi · Francesca A. Lisi

Received: 21 June 2011 / Accepted: 29 June 2011 / Published online: 26 July 2011
© The Author(s) 2011

The 20th International Conference on Inductive Logic Programming (ILP-2010, <http://ilp2010.dsi.unifi.it/>) was held in Florence during June 27–30, 2010. The program of ILP 2010 consisted of invited talks, regular talks, a poster session, a panel session, and featured for the first time a tutorial day. The formal post-conference proceedings of ILP 2010 have been published in the Volume 6489 of the Springer's Lecture Notes in Artificial Intelligence series.

This special issue contains a selection of articles presented at ILP 2010. We invited the authors of eight papers accepted for presentation at the conference to prepare a long version of their work and submit it to the journal as an alternative to having the full conference paper included in the proceedings. Five of these papers were accepted for publication in this issue after two rounds of peer-reviewing according to the journal criteria. One additional paper, "ILP turns 20—Biography and future challenges", (Stephen Muggleton, Luc De Raedt, David Poole, Ivan Bratko, Peter Flach, Katsumi Inoue, Ashwin Srinivasan) was solicited from the participants in the panel session. It offers a stimulating overview of the field, which highlights the major contributions of ILP to the broader areas of machine learning and artificial intelligence, as well as a vision of future research directions. It also provides a bibliography that is particularly valuable for newcomers to the field.

The papers submitted to the conference confirm a recent trend in ILP where techniques from statistics, graphical models, and kernel methods are more and more often incorporated into logic-based learning systems. Indeed, three out of the five papers accepted for publication in this issue deal with statistical relational learning. The other two address theoretical and practical problems, respectively, more in line with the tradition of ILP research.

P. Frasconi (✉)
Dipartimento di Sistemi e Informatica, Università degli Studi di Firenze, Firenze, Italy
e-mail: p-f@dsi.unifi.it

F.A. Lisi
Dipartimento di Informatica, Università degli Studi di Bari "Aldo Moro", Bari, Italy
e-mail: lisi@di.uniba.it

The paper “Gradient-based boosting for statistical relational learning: The relational dependency network case” (Sriram Natarajan, Tushar Khot, Kristian Kersting, Bernd Gutmann, and Jude Shavlik) introduces a novel combination of relational learners where relational regression trees are used in the context of functional gradient boosting of relational dependency networks. The resulting framework can exploit complex features defining the potential functions and perform collective classification. A pseudo-Gibbs sampling technique is employed for inferring latent variables in the case of incomplete data. Gradient boosting is shown to perform well in a set of classic SRL tasks.

The paper “Bridging logic and kernel machines” (Michelangelo Diligenti, Marco Gori, Marco Maggini, and Leonardo Rigutini) introduces a knowledge-based approach to multi-task learning where constraints amongst different tasks are expressed by first-order logic formulae converted into numerical function through the theory of T-norms eventually used as regularizers. The method has a sound theoretical justification based on a variant of the classic representer theorem for reproducing kernel Hilbert spaces. Since the optimization problem resulting from their setting is non-convex, the authors develop an original stage-based technique inspired by developmental psychology, where logic-derived constraints are only enforced in a subsequent stage.

The paper “Applying the information bottleneck to statistical relational learning” (Fabrizio Riguzzi and Nicola Di Mauro) brings Tishby et al.’s information bottleneck (IB) technique to the relational learning setting. After reviewing IB for learning graphical models, the authors show how the technique can be applied to logic programs with annotated disjunctions. The method is compared against some other popular probabilistic programming frameworks like Markov logic, ProbLog, and PRISM. The authors conclude that relational IB performs well in the presence of hidden variables and when the training interpretations share the same Herbrand base.

The paper “Inverse Subsumption for Complete Explanatory Induction” (Yoshitaka Yamamoto, Katsumi Inoue and Koji Iwanuma) is the recipient of the ILP 2010 Best Student Paper Award sponsored by this journal. In it, the authors investigate whether or not inverse subsumption can be embedded in a complete induction procedure and if it can how it is realized. As a result, they show a new form of inverse subsumption that ensures the completeness of generalization. Using the result, inverse entailment—widely used in ILP—can be reduced to inverse subsumption without losing the completeness for finding hypotheses in explanatory induction.

The paper “Data and Task Parallelism in ILP using MapReduce” (Ashwin Srinivasan, Tanveer Faruque and Sachindra Joshi) shows how MapReduce, a popular distributed computing approach, can be used effectively to perform the computationally expensive coverage test that is at the heart of many ILP systems and to perform multiple searches required by a greedy set-covering algorithm used by some popular ILP systems. The authors discuss the principal findings with synthetic and real-world datasets for both data and task parallelism, thus answering questions of implementation for ILP systems when the datasets and/or the search spaces are very large.

Acknowledgements We would like to thank all the authors, the reviewers, and the Journal for making this special issue possible.