

Guest editors' introduction: special issue on Inductive Logic Programming and on Multi-Relational Learning

Gerson Zaverucha¹ · Vítor Santos Costa^{2,3}

Published online: 16 July 2015
© The Author(s) 2015

We are delighted to introduce the 2013 special issue on Inductive Logic Programming and Multi-Relational Learning. This issue focuses on the problems in Machine Learning that are specific to data organised as multiple tables. Often this data is structured, and often items in the data-base are interconnected and dependent. Originally based on the induction of logic programs, the area broadened its scope and attracted a lot of attention and interest in recent years. Active topics include learning in logic, multi-relational learning and data mining, statistical relational learning, graph and tree mining, relational reinforcement learning, and other forms of learning from structured data.

The special issue is on the occasion of the 23rd International Conference on Inductive Logic Programming ILP'13, held on August 28–30, 2014 in Rio de Janeiro, Brazil. The ILP conference series, started in 1991, is the premier international forum on learning from structured data. The format of the 2013 conference followed the format of previous editions, with one invited talk per day, followed by long and short oral presentations. This edition of the conference accepted three types of contributions:

1. long papers (12 pages) describing original mature work containing appropriate experimental evaluation and/or representing a self-contained theoretical contribution.
2. short papers (6 pages) describing original work in progress, brief accounts of original ideas without conclusive experimental evaluation, and other relevant work of potentially high scientific interest but not yet qualifying for the above category.
3. papers relevant to the conference topics and recently published or accepted for publication by a first-class conference such as ECML/PKDD, ICML, KDD, ICDM, etc., or journals such as MLJ, DMKD, JMLR, etc.

ILP'13 received 42 submissions, 18 long, 21 short submissions, and 3 previously published papers. Each submission was reviewed by at least 3 program committee members. The short

✉ Gerson Zaverucha
gerson@cos.ufrj.br

¹ COPPE Universidade Federal do Rio de Janeiro, Rio de Janeiro, Brazil

² Center for Research in Advanced Computing Systems (CRACS/INESC-TEC), Porto, Portugal

³ Computer Science Dep, Fac of Sciences (DCC-FCUP), Universidade do Porto, Porto, Portugal

papers were evaluated on the basis of both the submitted manuscript and the presentation at the conference. Accepted papers were published in Springer Verlag's Lecture Notes in Computer Science as volume 8812, "Inductive Logic Programming - 23rd International Conference, ILP 2013, Rio de Janeiro, Brazil, August 28–30, 2013, Revised Selected Papers" with ISBN 978-3-662-44922-6. Work in progress, i.e., reports on ongoing research were also collected as "Late Breaking Papers of the 23rd international conference on Inductive Logic Programming, Rio de Janeiro, Brazil, August 28th - to - 30th, 2013", and are available as CEUR Workshop Proceedings 1187, CEUR-WS.org 2014.

After the conference, we had the opportunity to issue an open call for a *Machine Learning Journal* special issue on Inductive Logic Programming and Multi-Relational Learning. Although we did invite the authors of the four most high-evaluated ILP contributions to submit to the special number, the call was open to all ILP-13 and to unrelated papers. After several rounds of reviewing and revision, five of the submitted papers were selected for inclusion in this special number. We will now briefly describe the contribution that each paper makes; the papers are ordered in reverse alphabetic order on title.

'Probabilistic (logic) programming concepts', by Angelika Kimmig and Luc De Raedt surveys the main issues in combining logic and programming. The paper draws on the expertise of its two authors to provide an insightful and principled discussion of the strengths and pitfalls found when trying to work in this area, and it will surely become a reference work in this field.

'Meta-Interpretive Learning of Higher-Order Dyadic Datalog: Predicate Invention revisited' by S. H. Muggleton, Dianhuan Lin and Alireza Tamaddon-Nezhad. By studying the connection between induction and abduction, the authors have been able to propose an elegant and exciting new framework for Predicate Invention that opens new doors to ILP.

'Gradient-based Boosting for Statistical Relational Learning: The Markov Logic Network and Missing Data Cases' by Tushar Khot, Sriraam Natarajan, Kristian Kersting, and Jude Shavlik continues a highly successful line of work on applying gradient-based boosting to Statistical Relational Learning. The authors show how these techniques apply well to Markov Logic Networks. They first address the full-data case by optimising for pseudo-likelihood. In the followup, they use an EM based algorithm to address missing data.

'Efficient Inference and Learning in a Large Knowledge Base: Reasoning with Extracted Information using a Locally Groundable First-Order Probabilistic Logic' by William Yang Wang, Kathryn Mazaitis, Ni Lao, and William W. Cohen addresses the foundations of representation and learning in the NELL system. The authors use a representation close to Stochastic Logic Programs. Learning clause weights is possible by adapting personalised Page Rank. The end result is a highly scalable system that can learn recursive programs on large-scale datasets.

'Bandit-Based Monte-Carlo Structure Learning of Probabilistic Logic Programs' by Elena Bellodi, Nicola di Mauro, and Fabrizio Riguzzi addresses a major problem when learning pure and probabilistic Logic Programs: the cost one pays for the expressiveness of logic programs is a huge search space, that in practice cannot be completely explored. The authors introduce to this area a form of randomized search, single-player Bandit-Based Monte-Carlo tree search. The results are highly promising.

This special issue wouldn't have been possible without the cooperation and help of many people. We first would like to thank the authors for choosing the *Machine Learning* special number to publish their work; the reviewers for their insightful comments which often significantly improved the papers; the ILP organisers, Aline Paes, Kate Revoredo, Roosevelt Sardinha, and Fernanda Baião, for making it all possible. Last, but not least, we have to say

thanks to the Machine Learning journal staff and editors for their support throughout. Muito Obrigado!