Guest editorial: special issue on Inductive Logic Programming (ILP 2011)

Stephen Muggleton · Jianzhong Chen

Received: 22 June 2012 / Revised: 10 July 2012 / Accepted: 12 July 2012 / Published online: 26 July 2012
© The Author(s) 2012

This special issue is associated with the 21st International Conference of Inductive Logic Programming (ILP 2011) which was held in Cumberland Lodge, Great Windsor Park and marked 20 years since the first ILP workshop in 1991. During this period the conference has developed into the premier forum for work on logic-based machine learning. The submission procedure followed a similar format to that of previous conferences, and was particularly close to that used in ILP 2006. Submissions were requested in two phases. The first phase involved submission of short papers (6 pages) which were then presented at the conference and posted on the website prior to the event itself. In the second phase, reviewers selected papers for long paper submission (15 pages). These were assessed by the same reviewers, who then decided which papers to include in the journal special issue, and which ones in the conference proceedings. The post-conference proceedings of ILP 2011 have been published as Volume 7207 in the Springer Lecture Notes in Artificial Intelligence. A separate late-breaking paper proceedings is being published by Imperial College Press under the title “*Latest Advances in Inductive Logic Programming*”.

In the first phase there were 66 papers. Each paper was reviewed by 3 reviewers. Out of these 31 were invited as long papers. Out of the long paper submissions 5 were selected for the Machine Learning journal special issue and 23 were accepted for the proceedings. In addition, one paper (“*Machine Learning a Probabilistic Network of Ecological Interactions*”) was nominated by PC referees for the applications prize (the “*Michie ILP Application Prize*” sponsored by Syngenta Ltd) and one (“*k-Optimal: A Novel Approximate Inference Algorithm for ProbLog*”) for the theory prize (the “*Turing ILP Theory Prize*” sponsored by the Machine Learning Journal). The papers in this special issue represent the diversity and vitality in present ILP research.

S. Muggleton (✉) · J. Chen
Computational Bioinformatics Laboratory, Department of Computing, Imperial College, 180 Queens
Gate, London SW7 2BZ, UK
e-mail: s.muggleton@imperial.ac.uk

J. Chen
e-mail: jz.chen@imperial.ac.uk

The theory prize winning paper “*k-Optimal: A Novel Approximate Inference Algorithm for ProbLog*” (Joris Renkens, Guy Van Den Broeck, and Siegfried Nijssen) describes “*k-optimal*”, a new algorithm for approximating the probability of a query given a ProbLog program. The proposed optimizations are based on manipulating the binary decision diagrams (BDDs) and the proof of the approximation lower-bound is an interesting contribution. Another contribution is the analysis of the complexity of the problem of selecting the optimal set of k proofs. In addition, the method has been experimentally evaluated in a decision-theoretic machine-learning case (DTProbLog).

The paper “*Relational Networks of Conditional Preferences*” (Frederic Koriche) presents a generalization of conditional preference networks or CP-nets to the relational case, as well as methods for reasoning and learning the model from data. The proposed theoretical framework is sound and complete based on rigorous running time analysis of algorithms. An example application to movie recommendation is presented with some convincing empirical results.

The paper “*Learning Compact Markov Logic Networks With Decision Trees*” (Hassan Khosravi, Oliver Schulte, Jianfeng Hu, and Tianxiang Gao) presents an improvement to an existing method for learning the structure of a Markov Logic Network (MLN) by revising a learned Bayesian network. It uses a decision tree learner to compress the conditional probabilities in constructing an MLN, which significantly reduces the number of clauses in the revised MLN as shown in the experiments on 4 real world problems.

The paper “*Online Bayesian inference for the parameters of PRISM programs*” (James Cussens) attempts to tackle one of the challenging problems in the area—online parameter estimation of probabilistic logical models, by proposing an online Bayesian approach for PRISM programs. Specifically, assuming a mixture of Dirichlets as prior of the parameters, the paper shows how to compute the resulting Dirichlet mixture of the posterior. The experimental results show the promise of the approach presented in the paper.

The paper “*Learning Directed Relational Models With Recursive Dependencies*” (Oliver Schulte, Hassan Khosravi, and Tong Man) introduces a new approach to extend the parameterized Bayes nets to represent recursive dependencies. The approach is well presented by illustrations of concepts and examples as well as algorithms and proven theorems. Experimental evaluation shows that more accurate and efficient predictions have been achieved.

ILP 2011 was held at Cumberland Lodge in the UK from 31st July to 3rd August 2011 under the auspices of the Department of Computing, Imperial College London. In addition to the many technical paper presentations, the invited talks this year were given by a distinguished group of Artificial Intelligence researchers, namely Hector Geffner, Toby Walsh and Richard Sutton.

We gratefully acknowledge support of Syngenta Ltd for the applications prize and the Machine Learning Journal for the theory prize.