

Mach Learn

DOI 10.1007/s10994-006-9573-z

1 Guest editorial

2 **Rui Camacho · Ross King · Ashwin Srinivasan**

3 Published online: xxx

4 Springer Science + Business Media, LLC 2006

5 The papers in this issue are snapshots of ILP at the age of fourteen. Adolescence, as we all
6 remember, was a time when change was normal and we knew everything. We are reminded
7 here of Mark Twain's recollection: "When I was a boy of fourteen, my father was so ignorant
8 I could hardly stand to have the old man around. But when I got to be twenty-one, I was
9 astonished at how much he had learned in seven years". And so it is with ILP. Visiting friends
10 and relatives would be surprised at the changes: some of which is captured in the papers
11 here. In them we see a face of ILP in which systematic search for rules has been replaced
12 by randomised ones; first-order trees have been replaced by random forests of first-order
13 trees; observations no longer refer to the target concept being learnt; theories are constructed
14 using very large ensembles of clauses, and agent programs are constructed from observing
15 an expert. Compare this against the ILP that constructed a small number of definite clauses,
16 each with about 4 literals, using a data set of no more than a few hundred examples of the
17 target concept. And, as with Twain, we can expect to be astonished over the next few years.

18 As the character of ILP develops over the next few years, there are some traits that
19 are intriguing, and some that require a pause for thought. Amongst the former is a trend
20 towards the combination of statistical and relational information. We expect this will allow
21 ILP to be applied across a much wider range of problems than previously attempted. Also
22 interesting is the use of ILP as a general-purpose technique for constructing new features.
23 This has repeatedly been shown to be a very good way to use an ILP system. There are two
24 characteristics which may be viewed as a cause for concern by some. First, comprehensibility
25 of theories constructed appears to be no longer a feature of all ILP-constructed theories. This

R. Camacho (✉)

LIACC/FEUP University of Porto, Portugal

e-mail: rcamacho@fe.up.pt

R. King

Department of Computer Science University of Wales, Aberystwyth Wales, United Kingdom

e-mail: rdk@aber.ac.uk

A. Srinivasan

IBM India Research Laboratory Indian Institute of Technology New Delhi, India

e-mail: ashwin.srinivasan@in.ibm.com




Fig. 1 A rough categorisation of ILP papers published from 1991. These are estimates obtained from a manual classification by one of us (A.S.) of the ILPNet2 bibliographic entries. The average values over the period are: 38 (Conceptual), 9 (Implementation) and 17 (Application)

is specially so with ILP systems capable of constructing and using very large numbers of clauses, or combining probabilistic and logical information in complex ways. Second, there appears to be a trend towards fewer conceptual and application-oriented papers (see Fig. 1). If these estimates are reflective of reality, then the problem needs addressing. We are pleased to see that the ILP application presented in this issue has some substantially different features to many in the recent past and it could take the place of the “mutagenesis” problem as the new *Drosophila* for ILP.

Now for the clerical details. This volume contains five papers based on original ones presented at the 14th International Conference on Inductive Logic Programming. All the papers were significantly re-worked and extended with respect to their original version that appear in the ILP 2004 Proceedings. ILP 2004 was held in Porto, Portugal, in September 2004. This annual meeting of ILP continues to be the premier forum for presenting the most recent and exciting work in the field. Six invited talks—three from fields outside ILP, but nevertheless highly relevant to it—and 20 full presentations formed the nucleus of the conference.

Van Assche et al. show how the concept of random forests can improve relational decision trees. They address and propose solutions for the problem of introducing aggregate functions. The resulting first order random forest induction algorithm was implemented and experimentally evaluated on a variety of datasets. The results indicate that first order random forests with complex aggregates are an efficient and effective approach towards learning relational classifiers that involve aggregates over complex selections.

Mach Learn

47 Developing efficient methods to search the hypothesis space continues to be an important
48 issue in ILP. Železný et al. demonstrate the benefits of randomising and restarting ILP search
49 procedure. They show that significant speedups can be obtained using randomised restarts
50 by comparing five search algorithms operating on two principally different classes of ILP
51 problems.

52 In the paper by Tamaddoni-Nezhad et al. a problem that is increasingly important in bio-
53 logical applications is addressed. Observations do not necessarily refer to the target predicate
54 and background knowledge cannot be assumed to be complete. Using a combination of ab-
55 duction and induction an ILP program is used to construct a model by completing definitions
56 for background predicates.

57 The paper by Goadrich et al. addresses the problem of dealing with large, highly-skewed
58 datasets in ILP. They propose a randomized search method that collects good clauses from a
59 broad spectrum of points along the recall dimension in recall-precision curves. The method is
60 evaluated on information extraction tasks that typically involve many more negative examples
61 than positive examples. Those at Porto in September 2004 may recall that the conference
62 submission of this paper won the “Best Student Paper” award that year.

63 The paper by Könik et al. describes a method of relational learning by observation. In this,
64 a learner automatically creates cognitive agent programs that model expert task performance
65 in complex dynamic domains. Using observed behavior and goal annotations of an expert
66 as the primary input, their method interprets them in the context of background knowledge,
67 and returns an agent program that behaves similar to the expert.

68 ILP is still a young field, and one that will change much more before it achieves its
69 true potential. But we are amongst those that believe that this potential is great: in the
70 end, few other machine learning techniques have inherited the combination found in ILP of
71 representational power, mathematical foundations and explicit use of background knowledge.
72 The papers published in this special issue are all steps in the journey of growing up.