

Preface

This volume contains the regular papers of the 28th International Conference on Inductive Logic Programming (ILP 2018) held in Ferrara, Italy, during September 2–4, 2018.

Inductive Logic Programming (ILP) is a subfield of machine learning, which relies on logic programming as a uniform representation language for expressing examples, background knowledge, and hypotheses. Due to its strong representation formalism, based on first-order logic, ILP provides an excellent means for multi-relational learning and data mining. The ILP conference series, which started in 1991, is the premier international forum for learning from structured or semi-structured relational data. Originally focusing on the induction of logic programs, it has expanded its research horizon significantly over the years and welcomes contributions on all aspects of learning in logic, multi-relational data mining, statistical relational learning, graph and tree mining, learning in other (non-propositional) logic-based knowledge representation frameworks, exploring intersections to statistical learning, and other probabilistic approaches.

The conference will be co-located with two events:

- ACAI 2018: Advanced Course on AI, a summer school on Statistical Relational Artificial Intelligence, August 27–31, 2018
- PLP 2018: 5th Workshop on Probabilistic Logic Programming, September 1, 2018

This year we changed the submission model, trying to encourage participation, simplify the publication process, and attract high quality submissions. Two different tracks were organized, defining five kinds of submissions:

1. The journal track, whose accepted papers were published in the *Machine Learning Journal's* special issue on Inductive Logic Programming – ILP 2017 and 2018, accepting both new submissions and the best papers from ILP 2017.
2. The conference track, allowing four types of submissions:
 - a. Long papers describing original mature work, containing appropriate experimental evaluation and/or representing a self-contained theoretical contribution.
 - b. Short papers describing original work in progress, presenting preliminary results, brief accounts of original ideas, and other relevant work of potentially high scientific interest that does not yet qualify for the long paper category. Accepted papers appear in CEUR proceedings.
 - c. Works in progress papers describing ideas and proposals that the author(s) would like to present at the conference.
 - d. Papers relevant to the conference topics and recently published or accepted for publication by a first-class conference or journal.

The conference had two proceedings: the present LNAI proceedings for accepted long papers (submission category a), and the CEUR proceedings for up-and-coming

papers (submission category a, describing promising but less mature works) and short papers (submission category b). Submissions from categories c and d were presented at the conference but not included in any proceedings.

There were 24 submissions in total for categories a and b: 18 long papers and 6 short papers. We accepted 14 long papers split into 10 regular papers, published in these proceedings, and 4 up-and-coming papers, published in the CEUR proceedings. We accepted 4 out of the 6 short papers and they appear in the CEUR proceedings. All papers received 2.83 reviews on average by members of the Program Committee. Each accepted paper was presented at ILP 2018.

Submissions covering a wide range of topics are included in these proceedings, spacing from learning theories and rules to connections with deep learning, from the exploitation of knowledge graphs to applications of ILP to diagnostic systems to minimize the maintenance cost and downtime of equipment.

We had the pleasure of welcoming three invited speakers at ILP 2018:

- William Cohen, Professor at Carnegie Mellon University, USA:
“Using Deep Learning Platforms to Perform Inference over Large Knowledge Bases”
- Marco Gori, Professor at the University of Siena, Italy:
“Learning and Inference with Constraints”
- Maximilian Nickel, Research Scientist at Facebook AI Research:
“Hierarchical Representation Learning on Relational Data”

Three prizes were awarded:

- Best paper (supported by Springer);
- Best student paper among regular papers (supported by the *Machine Learning Journal*);
- Best student paper among up-and-coming papers (supported by the *Machine Learning Journal*).

The winners were announced during the conference and published on the conference website at <http://ilp2018.unife.it/>.

We would like to really thank all the people who contributed to the success of ILP 2018: the members of the Organizing Committee, the members of the Program Committee, the additional reviewers that have been solicited, and the sponsors.

July 2018

Fabrizio Riguzzi
Elena Bellodi
Riccardo Zese