

The 26th International Conference on Inductive Logic Programming (ILP'16)

James Cussens & Alessandra Russo

April 21, 2017

The 26th International Conference on Inductive Logic Programming (ILP 2016) (<http://ilp16.doc.ic.ac.uk>) was held in London, during September 4-6 2016, in the pleasant surroundings of the Warren House Conference Centre (<http://www.warrenhouse.com/>) in Kingston-upon-Thames, just next door to Richmond Park. Since its first edition in 1991, the annual ILP conference has served as the premier international forum for learning from structured relational data. Originally focusing on the induction of logic programs, over the years it has expanded its research horizon significantly and welcomed 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 with statistical learning and other probabilistic approaches. Theoretical advances in all these areas have also been accompanied by challenging applications of these techniques to important problems in fields like bioinformatics, medicine, and text mining.

Following the trend of past events, this edition of the conference solicited three types of submissions: (i) long papers describing original mature work containing appropriate experimental evaluation and/or representing a self-contained theoretical contribution; (ii) short papers describing original work in progress, brief accounts of original ideas without conclusive evaluation, and other relevant work of potentially high scientific interest but not yet qualifying for the long paper category; and finally (iii) 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, AAAI, IJCAI, etc. or journal such as MLJ, DMKD, JMLR etc.

The conference received 35 submissions in total: 10 long papers, 19 short papers, and 6 published papers. Each of the long and short paper submissions was reviewed by three Program Committee members. Only 4 of the 10 submitted long papers were accepted for presentation and publication. Short papers were initially evaluated on the basis of their submitted manuscript and their presentation and authors of a subset of these papers were invited to submit an extended version. After a second review process, only 6 extended papers were finally accepted for publication. In summary, together with the 4 long papers,

10 papers were accepted to be included in the present volume. The multiple stage review process, although rather complex, has enabled the selection of high-quality papers for the proceedings. We thank the members of the PC for providing high-quality and timely reviews. Out of all the submitted papers, an additional 14 papers were accepted for publication in the CEUR workshop proceedings series.

The ILP 2016 program included five large technical sessions: Logic and Learning (6 papers), Graphs and Databases (4 papers), Probabilistic Logic and Learning (5 papers), Algorithms, Optimisations and Implementations (9 papers), and Applications (8 papers). The papers presented represent well the current breadth of ILP research topics such as predicate invention, graph-based learning, spatial learning, logical foundations, statistical relational learning, probabilistic ILP, implementation and scalability, and applications in robotics, cyber-security and games. There was an excellent balance across theoretical and practical research.

ILP 2016 also received generous sponsorship by the Machine Learning journal for best student paper awards. The two best student paper awards of ILP 2016 were for the paper *Learning Disjunctive Logic Programs from Interpretation Transition*, by Yi Huang, Yisong Wang, Ying Zhang and Mingyi Zhang, and for the paper *Inductive Logic Programming meets Relational Databases: An Application to Statistical Relational Learning*, by Marcin Malec, Tushar Khot, James Nagy, Erik Blasch and Sriraam Natarajan. The conference also received sponsorship from Springer for a best paper award. This award was given to the paper *Generation of Near-Optimal Solutions Using ILP-Guided Sampling* by Ashwin Srinivasan, Gautam Shroff, Lovekesh Vig and Sarmimala Saikia.

With the intent of stimulating collaborations and discussion between academia and industry, the program also featured three invited talks by academic and industrial distinguished researchers. In the talk *Inferring Causal Models of Complex Relational and Dynamic Systems*, David Jensen, from the University of Massachusetts, presented key ideas, representations and algorithms for causal inference, and highlighted new technical frontiers. Frank Wood, from the University of Oxford, gave a talk entitled *Revolutionising Decision Making, Democratising Data Science, and Automating Machine Learning via Probabilistic Programming*. In his talk, he gave a broad overview of the emerging field of probabilistic programming, from the point of view of both programming (modelling) language and automated inference, and introduced most important challenges facing this field. Finally, Vijay Saraswat, senior research scientist in the Cognitive Computing Research division at the IBM T.J. Watson Center, discussed in his talk *Machine Learning and Logic: the beginnings of a new computer science?* the open challenges of building cognitive assistants in compliance, and the need to bring together researchers in natural language understanding, machine learning, and knowledge representation/reasoning to address them.

The conference featured, for the first time, also an international competition, designed and managed by Mark Law, a member of our Local Committee. The competition was aimed at testing the accuracy, scalability and versatility of the learning systems which were entered. The competition had two main tracks

for probabilistic and non-probabilistic approaches. The winners of the competition were Peter Schüller, from Marmara University, for his non-probabilistic approach and jointly Riccardo Zese, Elena Bellodi and Fabrizio Riguzzi for their probabilistic approach. Results of the competition are publicly available on <http://ilp16.doc.ic.ac.uk/competition>.

The ILP 2016 conference was kindly sponsored by IBM Watson Research, the Association of Logic Programming, Springer Lecture Notes in Artificial Intelligence, Artificial Intelligence journal, and Machine Learning journal. We would like to thank EasyChair for supporting submission handling. We would like to thank the members of the Local Committee of ILP 2016: Krysia Broda, Dalal Alrajeh and Mark Law. Our deep thanks also go to Mark Law for running the competition, and for setting up and maintaining the Website. The conference would not have been possible without their hard work.

Finally, we would like to thank all those involved in making ILP'16 such a success: our invited speakers, our sponsors, the program committee and of course, those who came to ILP'16 to present and discuss their work.