Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI-16)

Inductive Logic Programming: Challenges

Katsumi Inoue
National Institute of Informatics
Tokyo, Japan
inoue@nii.ac.jp

Hayato Ohwada
Tokyo University of Science
Noda, Japan
ohwada@rs.tus.ac.jp

Akihiro Yamamoto
Kyoto University
Kyoto, Japan
akihiro@i.kyoto-u.ac.jp

ILP: Aims and Scopes

Inductive Logic Programming (ILP) is a research area formed at the intersection of Machine Learning and logic-based knowledge representation. ILP has originally used logic programming as a uniform representation language for examples, background knowledge and hypotheses for learning, and then has provided excellent means for *multi-relational learning and data mining* from (non-trivially) structured data. ILP has also explored several connections with statistical learning and other probabilistic approaches, expanding research horizons significantly. A recent survey of ILP can be seen in (Muggleton *et al.* 2012).

The ILP conference series have been the premier international forum on ILP. Papers in ILP conferences address topics in theories, algorithms, representations and languages, systems and applications of ILP, and cover all areas of learning in logic, relational learning, relational data mining, statistical relational learning, multi-relational data mining, relational reinforcement learning, graph mining, and connections with other learning paradigms, among others.

ILP 2015: What's Hot

ILP 2015 is the twenty-fifth edition in the series of the ILP international conferences, and took place in Kyoto University, Kyoto, Japan, from 20th to 22nd of August, 2015. There were 10 technical sessions in ILP 2015, whose topics are: Nonmonotonic Semantics, Logic and Learning, Complexity, Action Learning, Distribution Semantics, Implementation, Kernel Programming, Data and Knowledge Modeling, and Cognitive Modeling. All papers and slides presented in technical sessions have been made open at the website.[1] Two post-conference proceedings will be published in a volume of Springer *Lecture Notes in Artificial Intelligence* series for selected papers and in an electronic volume of *CEUR-WS.org* for late-breaking papers. Moreover, there will be a special issue on ILP in *Machine Learning Journal*.

Three invited talks were given at ILP 2015, which represent three most important aspects of recent ILP: meta-level learning, probabilistic ILP and challenging applications.

[1] `http://www.ilp2015.jp/`

Meta-Level Learning

Stephen Muggleton gave the invited talk "Meta-Interpretive Learning: achievements and challenges". *Meta-Interpretive Learning* (MIL) is an ILP technique aimed at supporting learning of recursive definitions, by automatically introducing sub-definitions that allow decomposition into a hierarchy of reusable parts (Muggleton *et al.* 2014; 2015). MIL is based on an adapted version of a Prolog meta-interpreter, which derives a proof by repeatedly fetching (or abducing) first-order clauses whose heads unify with a given goal. MIL additionally fetches higher-order meta-rules whose heads unify with the goal and saves the resulting meta-substitutions to form a program. MIL can be a powerful method of *inductive programming* (Gulwani *et al.* 2015), which automatically synthesizes a variety of complex computer programs from background knowledge and only a few positive examples. Applications of MIL include learning regular and context-free grammars, learning from visual representations with repeated patterns, learning string transformations for spreadsheet applications, learning and optimizing recursive robot strategies (Cropper and Muggleton 2015) and learning tactics for proving correctness of programs. These applications of MIL were also presented as technical papers of ILP 2015 (Farquhar *et al.* 2015; Cropper *et al.* 2015; Dai *et al.* 2015).

Other approaches of meta-level learning have been proposed in ILP. For example, *meta-level abduction* (Inoue *et al.* 2013) has been applied to completion of FSHR-induced signaling networks (Rougny *et al.* 2015).

Probabilistic ILP

The *distribution semantics* for *probabilistic logic programming* (PLP) was firstly published by Taisuke Sato (1995). The semantics was proposed for probabilistic abduction, but has much more influenced to the field of *probabilistic ILP* and then a fertile ground for the general AI based on the combination of symbolic and statistical reasoning. ILP 2015 celebrated the 20th anniversary of the distribution semantics in the form of Sato's monumental talk "Distribution semantics and cyclic relational modeling". Since 1995, a lot of PLP languages with the distribution semantics have been developed including PRISM, ICL, LPADs, ProbLog, P-log, CP-logic and PITA to name a few. In his talk, Sato first reviewed Fenstat's representation theorem that mathematically relates probability to first-order logic. Any consistent probability assignment should be based on the possible worlds with a probability distribution, and the distribution semantics gives a way for it. The semantics starts with a simple computable distribution and transforms it to a complex one based on the semantics of logic programs. PRISM is the first implementation of the distribution semantics with the ability of parameter learning for probabilistic modeling (Sato and Kameya 2001). PRISM and other PLP languages are powerful enough to deal with recursively defined infinite models such as Markov chains with infinitely many transitions and probabilistic context free grammars with infinitely many sentence derivations, since these models often require to compute an infinite sum of probabilities. The latest development of PRISM enables us to compute this infinite sum of probabilities using cyclic propositional formulas and the EM algorithm (Sato and Mayer 2014). A promising application of such infinite computation in cyclic relational modeling is plan recognition from partial observations (Sato and Kojima 2015). After Sato's talk in ILP 2015, Fabrizio Riguzzi and his group showed several extended topics based on the distributed semantics, e.g., (Riguzzi *et al.* 2015).

Challenging Applications

Luc De Raedt reported in his invited talk "Applications of Probabilistic Logic Programming" on recent progress in applying PLP to challenging applications. De Raedt firstly looked back on Winograd's SHRDLU and reinterpreted its tasks with modern AI technologies. Now is the time to remove the assumptions of SHRDLU such as perfect knowledge about the world and deterministic actions and to *bridge the gap between low-level perception and high-level reasoning*. The real robot tasks need to deal with structured environments such as objects and their relationships as well as the background knowledge, and cope with uncertainty to learn from data. The details of robot grasping based on *relational kernels* (Frasconi *et al.* 2014) with numerical features was presented by Antanas *et al.* (2015). By extending the distribution semantics to cope with dynamics and continuous distributions, learning multi-relational object affordances has also been developed to specify the conditions under which actions can be applied on some objects (Moldovan *et al.* 2012; Nitti *et al.* 2013). Other recent applications of PLP include the PheNetic system (De Maeyer *et al.* 2015), which extracts the sub-network that best explains genes prioritized through a molecular profiling experiment from an interactome, and the problem of machine reading in CMU's Never-Ending Language Learning, which uses ProbFOIL (De Raedt *et al.* 2015) as an extension of the traditional rule-learning system FOIL with the distribution semantics.

Declarative Modeling

Two best student paper awards of ILP 2015 were given to Golnoosh Farnadi for the paper (Farnadi *et al.* 2015) and Francesco Orsini for the paper (Orsini *et al.* 2015). Both works are related to *declarative modeling*, the concept proposed by De Raedt, and probabilistic ILP. In fact, some declarativeness can be seen everywhere in any domain in the forms of constraints, graphs, actions, kernels, ontologies, etc. Modeling includes learning as a part, in particular in areas such as biology, robotics and cognition. Paramonov *et al.* (2015) have shown that a variety of relational query and graph mining problems can be modeled and solved with a variation of answer set programming (ASP).

Other Directions

Robot learning has been a very active topic in recent ILP. Sammut *et al.* (2015) have shown their ongoing work on a "robot engineer", which enables a closed-loop design, manufacture and testing in the domain of engineering. Learning from time-series data, e.g., *learning from interpretation transition* (Inoue *et al.* 2014), and learning from both discrete and continuous data are also hot topics in biological and other applications, e.g., (Ribeiro *et al.* 2015; Srinivasan *et al.* 2015).

New theories and applications of ILP have also been proposed. For example, both learning logics (Sakama *et al.* 2015) and learning proofs and strategies (Ho *et al.* 2015; Farquhar *et al.* 2015) perform meta-logical learning. These logical works are particularly suitable for ILP applications. Kuželka *et al.* (2015) have shown how to construct a Markov logic network from first-order default rules such that MAP inference from it correspond to default reasoning.

ILP 25 Years: What's Next

To celebrate the 25th anniversary of ILP conference series, ILP 2015 organized a panel discussion on past and future progress of ILP. The panelists were Stephen Muggleton, Fabrizio Riguzzi, Filip Zelezny, Gerson Zaverucha, Jesse Davis, Katsumi Inoue, who are all chairs of the last five years of ILP conferences (2011–2015), and Taisuke Sato. The discussion at the last panel held at ILP 2010 has been summarized as the survey paper (Muggleton *et al.* 2012), in which several future perspectives at that time were shown. Since then, the areas related to Machine Learning and AI have been rapidly growing and changing. Recent trends include learning from big data, from statistical learning to deep learning, integration of neural and symbolic learning, general intelligence, etc. The panel "ILP 25 Years" then considered what ILP could contribute to these recent changes. Panelists gave their own views of recent trends of ILP for these five years and perspective on "What next for ILP".

Muggleton claimed that integration of learning, perception and action are important, and social skills should be learned in human-like computing. Zelezny argued inspiration from deep learning, in the context of predicate invention and lifted neural networks. Riguzzi predicted a stronger connection to the Semantic Web and Linked Open Data. Zaverucha considered that big data may need more techniques for propositionalization and parallelism. Davis noticed a future trend that learns from very few examples and lots of (commonsense) knowledge. Inoue mentioned learning from state transitions with lots of fluents and learning with general intelligence. Finally, Sato told us the importance of ILP that should exploit millions of propositions learned from big data. All believe that advances of innovative techniques for ILP as well as new challenges will give significant impacts on many fields of AI and Machine Learning.

References

Laura Antanas, Plinio Moreno, and Luc De Raedt. 2015. Relational kernel-based robot grasping with numerical features. *ILP 2015*.

Andrew Cropper and Stephen Muggleton. 2015. Learning efficient logical robot strategies involving composable objects. In *Proceedings of IJCAI-15*, pp.3423–3429. AAAI Press.

Andrew Cropper, Alireza Tamaddoni-Nezhad, and Stephen Muggleton. 2015. Meta-interpretive learning of data transformation programs. *ILP 2015*.

Wang-Zhou Dai, Stephen Muggleton, and Zhi-Hua Zhou. 2015. Logical vision: Meta-interpretive learning for simple geometrical concepts. *ILP 2015*.

Dries De Maeyer, Bram Weytjens, Joris Renkens, Luc De Raedt, and Kathleen Marchal. 2015. PheNetic: Network-based interpretation of molecular profiling data. *Nucleic Acids Research*, 43: W244–W250.

Luc De Raedt, Anton Dries, Ingo Thon, Guy Van den Broeck, and Mathias Verbeke. 2015. Inducing probabilistic relational rules from probabilistic examples. In *Proceedings of IJCAI-15*, pp.1835–1843. AAAI Press.

Golnoosh Farnadi, Stephen H. Bach, Marjon Blondeel, Marie-Francine Moens, Martine De Cock, and Lise Getoor. 2015. Statistical relational learning with soft quantifiers. *ILP 2015*.

Colin Farquhar, Gudmund Grov, Andrew Cropper, Stephen Muggleton, and Alan Bundy. 2015. Typed meta-interpretive learning for proof strategies. *ILP 2015*.

Paolo Frasconi, Fabrizio Costa, Luc De Raedt, and Kurt De Grave. 2014. kLog: A language for logical and relational learning with kernels. *Artificial Intelligence* 217: 117–143.

Sumit Gulwani, José Hernández-Orallo, Emanuel Kitzelmann, Stephen Muggleton, Ute Schmid, and Benjamin Zorn. 2015. Inductive programming meets the real world. *Communications of the ACM*, 58(11): 90–99.

Tuan Dung Ho, Min Zhang, and Kazuhiro Ogata. 2015. A case study on extracting the characteristics of the reachable states of a state machine formalizing a communication protocol with inductive logic programing. *ILP 2015*.

Katsumi Inoue, Andrei Doncescu, and Hidetomo Nabeshima. 2013. Completing causal networks by meta-level abduction. *Machine Learning* 91: 239–277.

Katsumi Inoue, Tony Ribeiro, and Chiaki Sakama. 2014. Learning from interpretation transition. *Machine Learning* 94: 51–79.

Ondřej Kuželka, Jesse Davis, and Steven Schockaert. 2015. Constructing Markov logic networks from first-order default rules. *ILP 2015*.

Bogdan Moldovan, Plinio Moreno, Martijn van Otterlo, José Santos-Victor, and Luc De Raedt. 2012. Learning relational affordance models for robots in multi-object manipulation tasks. In *Proceedings of 2012 IEEE International Conference on Robotics and Automation* (ICRA 2012), pp.4373–4378. IEEE.

Stephen Muggleton, Luc De Raedt, David Poole, Ivan Bratko, Peter Flach, Katsumi Inoue, and Ashwin Srinivasan. 2012. ILP turns 20—biography and future challenges. *Machine Learning* 86: 3–22.

Stephen Muggleton, Dianhuan Lin, Niels Pahlavi, and Alireza Tamaddoni-Nezhad. 2014. Meta-interpretive learning: application to grammatical inference. *Machine Learning* 94: 25–49.

Stephen Muggleton, Dianhuan Lin, and Alireza Tamaddoni-Nezhad. 2015. Meta-interpretive learning of higher-order dyadic datalog: predicate invention revisited. *Machine Learning* 100: 49–73.

Davide Nitti, Tinne De Laet, and Luc De Raedt. 2013. A particle filter for hybrid relational domains. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems* (IROS 2013), pp.2764–2771. IEEE.

Francesco Orsini, Paolo Frasconi, and Luc De Raedt. 2015. kProlog: An algebraic Prolog for kernel programming. *ILP 2015*.

Sergey Paramonov, Matthijs van Leeuwen, Marc Denecker, and Luc De Raedt. 2015. An exercise in declarative modeling for relational query mining. *ILP 2015*.

Tony Ribeiro, Morgan Magnin, Katsumi Inoue, and Chiaki Sakama. 2015. Learning multi-valued biological models with delayed influence from time-series observations. *ILP 2015*.

Fabrizio Riguzzi, Elena Bellodi, Evelina Lamma, Riccardo Zese, and Giuseppe Cota. 2015. Structure learning of probabilistic logic programs by MapReduce. *ILP 2015*.

Adrien Rougny, Yoshitaka Yamamoto, Hidetomo Nabeshima, Gauvain Bourgne, Anne Poupon, Katsumi Inoue, and Christine Froidevaux. 2015. Completing signaling networks by abductive reasoning with perturbation experiments. *ILP 2015*.

Chiaki Sakama, Tony Ribeiro, and Katsumi Inoue. 2015. Learning deduction rules by induction. *ILP 2015*.

Ashwin Srinivasan, Michael Bain, Deepika Vatsa, and Sumeet Agarwal. 2015. Identification of transition models of biological systems in the presence of transition noise. *ILP 2015*.

Claude Sammut, Raymond Sheh, Adam Haber, and Handy Wicaksono. 2015. The robot engineer. *ILP 2015*.

Taisuke Sato. 1995. A statistical learning method for logic programs with distribution semantics. In *Proceedings of the 12th International Conference on Logic Programming* (ICLP 1995), pp.715–729. MIT Press.

Taisuke Sato and Yoshitaka Kameya. 2001. Parameter learning of logic programs for symbolic-statistical modeling. *Journal of Artificial Intelligence Research* 15: 391–454.

Taisuke Sato and Ryosuke Kojima. 2015. Goal and plan recognition via parse trees using prefix and infix probability computation. In *Proceedings of ILP 2014*, LNAI 9046. Springer.

Taisuke Sato and Philipp Meyer. 2014. Infinite probability computation by cyclic explanation graphs. *Theory and Practice of Logic Programming* 14(6): 909–937.