

Guest editors introduction: special issue on inductive logic programming

Jesse Davis¹ · Jan Ramon¹

Received: 22 February 2016 / Accepted: 24 February 2016 / Published online: 13 May 2016
© The Author(s) 2016

This special issue focuses on the field of inductive logic programming (ILP), which is a subfield of machine learning that uses logic as a uniform representation language for examples, background knowledge and hypotheses. From these roots, ILP's scope has grown to encompass different approaches that address learning from structured relational data. One notable example is the area of statistical relational learning which focuses on extending ILP to model uncertainty.

The special issue is also in conjunction with the 24th International Conference on Inductive Logic Programming (ILP), which was held from September 14th to 16th, 2014, in Nancy, France in co-location with ECML/PKDD-2014. To avoid the redundancy between the conference proceedings and the special issue, authors with an accepted paper at ILP were asked to either have their paper appear in the conference proceedings or submit an extended version of the paper to the special issue. While associated with the ILP conference, there was an open call for submissions to this special issue.

The special issue received six submissions of which three were originally submitted to the ILP conference. Ultimately, four were accepted to appear in the special issue. The papers offer a nice reflection on the strengths of ILP and relational learning and where the field is headed. Namely, the articles build off ILP's established track record of being particularly well suited to addressing important applications and the vibrant recent work that focuses on modeling uncertainty in relational data.

One of the first application areas where ILP gained significant traction was in analyzing molecules, in particular for the task of drug design. Drugs are small molecules that affect disease by binding to a target protein in the human body. Approaches to drug design depend on whether properties and/or structure of the drug or target are known. Previously, ILP has been successfully applied to identify properties of a drug molecule responsible for it binding

✉ Jesse Davis
Jesse.Davis@cs.kuleuven.be

Jan Ramon
Jan.Ramon@cs.kuleuven.be

¹ Department of Computer Science, KU Leuven, Celestijnenlaan 200A, 3001 Heverlee, Belgium

to the target protein. In contrast, the article “ILP-assisted De Novo Drug Design” by Kaalia, Srinivasan, Kumar, and Ghosh addresses this problem from a new perspective. Instead of analyzing the properties of potential drug molecules, it attempts to infer information about the active site of a protein. It uses ILP to identify common interaction patterns between a series of probes (small molecules) and a set of proteins related to the target. Technically, this paper exploits several well-known strengths of ILP. First, it exploits the ability to incorporate domain-knowledge in order to guide and customize the search process. Second, ILP returns learned models that are easy for experts to interpret.

The second paper, “Transductive Hyperspectral Image Classification: Towards Integrating Spectral and Relational Features via an Iterative Ensemble System” by Appice, Guccione, and Malerba combines several ideas from (statistical) relational learning with other machine learning techniques to tackle the practical problem of labeling pixels in hyperspectral images. The paper exploits two important ideas from relational learning to capture the spatial correlations in the data. First, it makes use of domain-specific relational features to capture information about a pixel’s spatial neighborhood. Second, it employs an iterative collective classification scheme to jointly label the pixels to exploit the fact that neighboring pixels are more likely to have the same label. The authors find that their proposed framework results in improved accuracy on the standard benchmark data sets for this problem. Furthermore, the authors carefully assess the contribution that each subcomponent of the learning system makes to its overall performance.

The third paper, “Fast Learning of Relational Dependency Networks” by Schulte, Qian, Kirkpatrick, Yin, and Sun, falls into the area of statistical relational learning. It tackles the problem of learning the structure of a relational dependency network (RDN) from data. The proposed approach first learns a Bayesian network (BN), which is then converted into a relational dependency network. The paper provides a closed-form method for deriving the RDN’s parameters from the learned BN. Theoretically, the authors show sufficient and necessary conditions for this transformation to result in a consistent RDN. The main empirical advantage is that the proposed approach provides improved scalability.

The final paper, “Probabilistic Logic Programming for Hybrid Relational Domains” by Nitti, De Laet, and De Raedt, combines the move towards incorporating uncertainty into ILP and focus on applications. The paper studies probabilistic logic programming, which incorporates probabilistic constructs into logic programming. Technically, it shows how to model both discrete and continuous variables within probabilistic logical programs for both static and dynamic settings. The paper then proposes and establishes the theoretical correctness of a particle-filter based inference algorithm for the language. Among other empirical results, the paper shows how the framework can successfully model real-world tracking problems that are common in robotics and vision.

We hope that the readers enjoy the papers in this issue. We want to thank all the authors and reviewers for the time spent in putting this special issue together. Finally, we want to thank Filip Zelezny for serving as the action editor for “Probabilistic Logic Programming for Hybrid Relational Domains.”