Ann Math Artif Intell (2008) 54:1–2
DOI 10.1007/s10472-009-9131-z

Preface

Manfred Jaeger · Lise Getoor · Kristian Kersting

Published online: 14 April 2009

Probabilistic Relational Learning is a research area that brings together several previously separated lines of research. One line of research originates with early approaches to combine probabilistic graphical models with higher-level, logic-based representation languages. Originally mostly pursued for the purpose of knowledge representation and reasoning (and then called *knowledge-based model construction*), this line of research gained significant momentum when its focus shifted to machine learning, where it was found that these new representation languages are well-suited to provide statistical models for relational data. A second line of research is represented by *inductive logic programming*, which for a long time had been concerned with learning purely logical models from logical or relational data, and over the past decade has increasingly turned towards probabilistic-logic models as well. Finally, over the past few years a further integration has taken place with sub-areas of machine learning that are concerned with learning from structured data, notably *graph mining*.

The vitality of the emerging field is manifested by numerous international workshops and seminars, notably the sequence of workshops on *Statistical Relational*

M. Jaeger (✉)
Institut for Datalogi, Aalborg Universitet,
Selma Lagerlöfs Vej 300, 9220 Aalborg, Denmark
e-mail: jaeger@cs.aau.dk

L. Getoor
Department of Computer Science &
University of Maryland Institute for Advanced Computer Studies,
University of Maryland, College Park, MD, USA
e-mail: getoor@cs.umd.edu

K. Kersting
Knowledge Discovery Department,
Fraunhofer IAIS, Schloss Birlinghoven,
53754 Sankt Augustin, Germany
e-mail: kristian.kersting@iais.fraunhofer.de

Learning (2000, 2003, 2004, 2006, 2009), on *Mining and Learning with Graphs* (held annually since 2003), on *Multi-Relational Data Mining* (since 2001), and Dagstuhl seminars on Probabilistic, Logical and Relational Learning (2005, 2007).

Contributions for this issue of the *Annals of Mathematics and Artificial Intelligence* were solicited by an open call for papers, and by special invitations to participants of the 2007 Dagstuhl seminar to submit papers related to their seminar contributions. The seven papers that were selected for inclusion in this issue cover a wide range of topics in Probabilistic Relational Learning: the paper by Mateescu and Dechter provides an introduction to mixed networks, which combine probabilistic and constraint networks. Angelopoulos and Cussens address Bayesian learning of Bayesian network structures by using stochastic logic programs to declaratively and flexibly specify the prior over the structures, while using MCMC sampling is used to approximate the posterior. Fierens et al. describe how ordering search, a technique originally developed for learning standard Bayesian networks, can be lifted to first-order probabilistic-logical representations. Sato et al. present a variational learning approach for the symbolic-statistical modeling language PRISM. The paper by Ng et al. presents a novel, highly expressive framework for probabilistic-logical models based on higher-order logic. Costa et al. present an approach to structured input-output prediction that reduces the prediction of a structured output to the prediction of a set of multinomial parameters, and show how this reduced problem can be solved by kernel logistic regression. The paper by Natarajan et al. develops gradient descent and expectation maximization algorithms for learning parameters in relational models for multiple and nested causal mechanisms.