Probabilistic, Logical and Relational Learning
Toward a Synthesis
Dagstuhl Seminar 05051, January 30 - February 4, 2005
Executive Summary

Luc De Raedt[1], Tom Dietterich[2], Lise Getoor[3], and Stephen Muggleton[4]

[1] University of Freibrug, DE
deraedt@informatik.uni-freiburg.de
[2] Oregon State University, USA
tgd@cs.orst.edu
[3] University of Maryland, College Park, USA
getoor@cs.umd.edu
[4] Imperial College London, GB
shm@doc.ic.ac.uk

1 Motivation

One of the central open questions of artificial intelligence is concerned with combining expressive knowledge representation formalisms such as relational and first-order logic with principled probabilistic and statistical approaches to inference and learning. This combination is needed in order to face the challenge of real-world learning and data mining problems in which the data are complex and heterogeneous and we are interested in finding useful predictive and/or descriptive patterns.

In this context, the terms probabilistic and statistical refer to the use of probabilistic representations and reasoning mechanisms grounded in probability theory, such as Bayesian networks, hidden Markov models and probabilistic grammars and the use of statistical learning and inference techniques. Such representations have been successfully used across a wide range of applications and have resulted in a number of robust models for reasoning about uncertainty. The primary advantage of using probabilistic representations is that well-understood and principled statistical inference and learning algorithms exist.

The term learning refers to deriving the different aspects of the probabilistic model on the basis of data. Typically, one distinguishes various learning algorithms on the basis of the given data (fully or partially observable variables) or on the aspect being learned (the parameters of the probabilistic representation or the structure of the model). Statistical and Bayesian approaches provide a unified framework for learning a model, whether through model selection or explicitly modeling a distribution over the models.

The terms logical and relational refer to first order logical and relational representations such as those studied within the field of computational logic and database theory. The primary advantage of using such expressive representations is that it allows one to elegantly and naturally represent complex situations involving a variety of objects as

Dagstuhl Seminar Proceedings 05051
Probabilistic, Logical and Relational Learning - Towards a Synthesis
http://drops.dagstuhl.de/opus/volltexte/2006/412

well as relations among the objects, which is not possible using the simpler propositional or feature vector based representations. So, probabilistic, logical and relational learning aims at combining its three underlying constituents: statistical learning and probabilistic reasoning within logical or relational representations.

2 Seminar Goals

The goal of this seminar was to bring together the researchers interested in the area of statistical, logical and relational learning. This allowed the participants to explore the foundations, challenges and research opportunities raised by this important open problem in artificial intelligence.

The first day of the workshop began with a series of introductary talks. Lise Getoor gave an overview talk on 'Tasks and Issues in Probabilistic, Logical and Relational Learning'. Next, David Poole gave a talk on 'Mixes of Logic and Probability' and Benjamin Taskar gave a tutorial on 'Advanced Probabilistic Reasoning'. After lunch, Luc De Raedt gave a tutorial, 'From ILP to Probabilistic ILP' and Stephen Muggleton gave a talk, 'Applications that Require Probabilistic Logics'. These talks all helped to survey some of the existing approaches toward probabilistic, logical and relational learning, and helped introduce some of the different approaches and terminology.

Next came the intruductions to several focus/challenge problems that participants had proposed prior to the start of the workshop. These problems where introduced, and each evening, a group would get together to discuss and debate the challenges. Topics included: bio-molecular network learning, citation analysis, natural language processing, temporal reasoning, relational reinforcement learning and representation issues. These focused breakout sessions worked well and brought together researchers with a common application interest from different research sub-communities.

The second day began with short presentations of the results of each of the focus groups. Then began a series of more focused presentations by the workshop participants. The presentations where organized into different topics, and a discussion leader was chosen to lead a discussion after a set of presentations was given. This helped to ensured a lively and indepth discussion, well beyond that which is easily achievable at a short one-day workshop or meeting. The topics included the following:

- The Big Picture
  - Jensen: Joint modeling of structure and attributes
  - Dietterich: Learning and Reasoning
- The Space of Models
  - Cussens: Relationships between ICL, PRISM, SLP
  - Muggleton: Semantics and Expressive Power of Probabilistic-Logical Models
  - Domingos: Markov Logic
- Relational Models and Sequential Data
  - Fern: Inferring relational states from sequences
  - Kersting: Probabilistic models for logical sequences (LOHMMs)
  - Sato: PRISM: EM learning from negative observations, constrained HMMs, finite PCFGs

    - Malerba: relational learning in spatial regression problems
- Aggregation
    - Perlich: Feature construction through aggregation
    - Siebes: Beyond aggregation: Handling M:N relations
    - Tadepalli: Learning with Conditional Influences and Combining Rules
- The Power of Views
    - Page: Improving predictive accuracy via views
    - Scheffer: Multi-view learning and link farm discovery
- Natural Language Processing
    - Roth: Unified representations for NLP: encoding, learning, and reasoning
    - Buntine: Probabilistic approaches to NLP semantics
    - Mooney: Statistical NLP and Relational Information Extraction
- Relational Approaches to Decision Making
    - Lloyd: Bayesian-Logical Agents
    - Khardon: Supervised Relational Reinforcement Learning
    - Passerini: Inductive Logic Programming from Traces
- New Problems/New Algorithms
    - Kern-Isberner: Learning Association Rules for MaxEnt
    - Lachiche: Bayesian Classification of Structured Individuals
    - Neville: Latent variables for discovering hidden structures responsible for generating autocorrelated variables
    - Kruse: Knowledge-based operations on graphical models
- Inference
    - Angelopouous: SLP, MCMC, and Priors
    - Jaeger: Stochastic Sampling Directly in the Relational Language
    - McAllester: A* Instead of Dynamic Programming

3 Conclusion

This workshop brought together a signficant number of researchers from all over the world that are working on all aspects of probabilist, logical and relational learning. It was also the first workshop on this topic where there was sufficent time for indepth discusions, debates and working groups. It was exciting to see the progression through the week. It was clear that some common ground had been identified, yet this was just the start. There was a general feeling that the workshop was a success, and a lot of enthusiasm for a follow on workshop.

4 Participants

Nicos Angelopouos, University of York
Hendrik Blockeel, Katholieke Universiteit Leuven
Wray Buntine, Helsinki Institute for Information Technology
Jianzhong Chen, Imperial College London
James Cussens, University of York
Luc De Raedt, Universität Freiburg

Thomas Dietterich, Oregon State University
Pedro Domingos, University of Washington
Kurt Driessens, University of Waikato
Saso Dzeroski, Jozef Stefan Institute
Alan Fern, Oregon State University
Peter Flach, University of Bristol
Paolo Frasconi, Universita degli Studi di Firenze
Johannes Fürnkranz, Darmstadt, (D)
Lise Getoor, University of Maryland, College Park
Robert Givan, Purdue University
Elias Gyftodimos, University of Bristol
Thomas Grtner, Fraunhofer Institut
Robert Holte, University of Alberta
Manfred Jaeger, Aalborg University
David Jensen, University of Massachusetts
Gabriele Kern-Isberner, Universitt Dortmund
Kristian Kersting, Universität Freiburg
Roni Khardon, Tufts University
Joost Kok, Leiden University
Rudolf Kruse, Universität Magdeburg
Nicolas Lachiche, Université Louis Pasteur
Codrina Lauth, Fraunhofer Institut
John Lloyd, Australian National University
Donato Malerba, Universita degli Studi
Heikki Mannila, University of Helsinki
David McAllester, Toyota Technological Institute at Chicago
Brian Milch, University of California at Berkeley
Ray Mooney, University of Texas at Austin
Stephen H. Muggleton, Imperial College London
Jennifer Neville, University of Massachusetts at Amherst
C. David Page, University of Wisconsin-Madison
Andrea Passerini, Universita degli Studi di Firenze
Hanna Pasula, MIT
Claudia Perlich, New York University
David Poole, University of British Columbia
Foster Provost, New York University
Jan Ramon, Katholieke Universiteit Leuven
Dan Roth, University of Illinois at Urbana-Champagne
Taisuke Sato, Tokyo Institute of Technology
Tobias Scheffer, Humboldt Universitt
Mathieu Serrurier, Université Paul Sabatier
Arno Siebes, Utrecht University
Prasad Tadepalli, Oregon State University
Ben Taskar, University of California at Berkeley
Volker Tresp, Siemens AG

Michael Werman, The Hebrew University of Jerusalem
Stefan Wrobel, Fraunhofer Institut
Martijn van Otterlo, Universiteit Twente