

Preface

Much has been achieved in the field of AI, yet much remains to be done if we are to reach the goals we all imagine. One of the key challenges with moving ahead is closing the gap between logical and probabilistic AI. Logical AI has mainly focused on complex representations, and probabilistic AI on uncertainty. Intelligent agents, however, must be able to handle both the complexity and the uncertainty of the real world.

Recent years have seen an explosion of successes in combining probability and (subsets of) first-order logic, programming languages and relational databases in several subfields of AI such as reasoning, learning, knowledge representation, planning, databases, NLP, robotics, vision, and others. Nowadays, we can learn statistical relational models automatically from millions of inter-related objects. We can generate optimal plans and learn to act optimally in uncertain environments involving millions of objects and relations among them. We can even perform lifted probabilistic inference, avoiding explicit state enumeration by manipulating first-order representations directly. So far, however, the researchers combining logic and probability in each of these subfields have been working mostly independently.

This workshop was designed for attempts at synthesis, forming a common core of problems and ideas, cross-pollinating across subareas, and ultimately starting to explore what might be called Statistical Relational AI: the study and design of intelligent agents that act in noisy worlds composed of objects and relations among the objects.

In total, the workshop received 30 submissions. Each submission was reviewed by at least two program committee members. The committee decided to accept 14 full technical papers and 13 short position papers and abstracts. Additionally, we had three invited speakers, Henry Kautz (University of Rochester, USA), Vibhav Gogate (University of Texas, Dallas, USA), and Luc De Raedt (KU Leuven, Belgium). The selected papers cover a wide range of topics in Statistical Relational AI such as parameter and structure learning, lifted inference, probabilistic programming, planning, information extraction, constraint optimization, NLP, among others, and clearly show the promise of statistical relational techniques for AI.

More precisely, Apse et al. present a framework for lifted MAP-LP inference with a new symmetry detection technique. Beltagy and Mooney develop a new inference algorithm for Markov logic that computes probabilities of complete formulae rather than ground atoms. Ceylan and Penalosa study the problem of reasoning in the probabilistic description logic BEL by reduction to Bayesian network inference. The paper by Choi et al. introduces a parameter learning algorithm for relational Kalman filters. Dirks et al. apply relational influence diagrams to a challenging real-world geological planning problem. Farnadi et al. propose to extend probabilistic soft logic with fuzzy quantifiers. Gribkoff et al. analyze the complexity of lifted inference in probabilistic databases and statistical relational learning. Kazemi et al. introduce relational logistic regression, a directed analogue of Markov logic networks. The paper by Khot et al. presents a one-class classification algorithm for relational domains. Lee et al. use marginal MAP inference algorithms for solving conformant planning problems. Lowd et al. argue that existing statistical relational techniques can be readily adapted to do adversarial reasoning. Nath and Domingos investigate tractable learning of relational sum-product networks, as well as automated debugging techniques based on probabilistic programs. Niepert and Domingos present a framework for tractable probabilistic knowledge bases that feature a rich query language. Niepert and Van den Broeck show that exchangeability is a statistical property that renders probabilistic inference tractable. Nishino et al. propose a new parameter learning algorithm for ProbLog that reduces the number of para-

meters. Poyrekar et al. perform an empirical analysis of counting belief propagation, identifying grounding as the bottleneck. Renkens et al. contribute a new bounded approximation method for weighted model counting based on probabilistic logic programming principles. The paper by Riedel et al. introduces a probabilistic programming language for declarative machine learning models. Rocktaschel et al. propose to learn low-dimensional embeddings that simulate the behaviour of first-order logic. Ruttenberg et al. introduce an approach to representing and reasoning about hierarchies in probabilistic programming. The paper by Schulte and Riahi applies Bayesian networks to relational outlier detection. Venugopal and Gogate propose a general method for addressing the evidence problem in lifted inference by reducing the domain size. Verman et al. present an algorithm that models distributed constraint optimization as a graph processing problem. Vlasselaer et al. study efficient inference algorithms for dynamic relational models with local structure. Wang et al. introduce a scalable probabilistic logic that enables efficient learning and inference on graphs. Finally, Zhengya et al. develop a scalable MLN structure learner that combines random walks with subgraph pattern mining.

We believe that this compilation of papers provides useful perspectives on statistical relational AI and hope it may inspire the reader to contribute to this challenging and exciting new research area.

– *Guy Van den Broeck, Kristian Kersting, Sriraam Natarajan, David Poole*