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 statistical AI. Logical AI has mainly focused on complex representations, and statistical 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 respectively programming languages and databases in several subfields of AI such as reasoning, learning, knowledge representation, planning, databases, NLP, robotics, vision, and so on. Nowadays, we can learn probabilistic 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. Exploiting shared factors can speed up message-passing algorithms for relational inference but also for classical propositional inference such as solving SAT problems. We can even perform lifted probabilistic inference avoiding explicit state enumeration by manipulating first-order state 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. This follows the previously successful editions of this workshop in 2010 at AAAI and 2012 at UAI.

Among a number of high-quality submissions, we selected 18 high quality papers that met our three main criteria: novelty, relevance and impact. Each submission was reviewed by three program committee members. We had three invited speakers, Dan Suciu (University of Washington), Prasad Tadepalli (Oregon State University) and Toby Walsh (NICTA). The selected papers cover a wide range of topics in Statistical Relational AI such as lifted inference, planning, learning, SAT, NLP, data integration, among others, and clearly show the promise of SR techniques for AI.

Lifted inference remained the hot topic of this year's workshop much like its previous edition. As many as 10 papers were in the area of lifted inference with varying focus problems — Van Den Broeck and Taghipour et al. analyzed the complexity of lifted inference. Ahmadi et al. and Noessner et al. explored the problem of parallelizing the inference process. Hadiji et al. and Sarkhel et al. considered the problem of MAP inference and designed approaches based on bootstrapping and lifting local search algorithms respectively. Mladenov and Kersting also considered the problem of MAP inference and explored the connections between MAP LP relaxation and their symmetries. Pu et al. considered the problem of lifted inference in the presence of transitive relations by reducing the models to random graph representation and applying approximation algorithms. Niepert presented an algorithm based on Rao-Blackwellization for post-processing samples in a MCMC inference algorithm when computing marginals. Finally, Van den Broeck et al. presented an invited contribution of using lifted inference for parameter learning in SRL models.

Outside of lifted inference, there were two papers on NLP — Natarajan et al. presented a weakly supervised method for relation extraction while Raghavan and Mooney presented an algorithm for online rule learning from natural language text based on Bayesian Logic Programs. There were two papers on agents — Song and Kautz presented an algorithm for event recognition using Markov Logic. Papai and Kautz combined propositional modal logic in multiagent setting with Markov logic. Webb and Domingos presented a new language called Tractable Probabilistic Knowledge Bases that supports hierarchies and probabilistic relations and uses it to model existence uncertainty. Sato and Kameya presented a position paper on enhancing generative modeling in PRISM. Kuo and Poole presented a discussion paper where they combine ontologies and parameterized random variables and analyzed the properties of this combination. Finally Tadepalli et al. presented a position paper on Relational MDPs.

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.

Vibhav Gogate, Kristian Kersting,
Sriraam Natarajan, and David Poole
Workshop Cochairs