Using MCMC in Probabilistic Relational Models

Fabian Kaelin
School of Computer Science
McGill University
Montreal, Canada
fabian.kaelin@mail.mcgill.ca

Doina Precup
School of Computer Science
McGill University
Montreal, Canada
dprecup@cs.mcgill.ca

In the active field of statistical relational learning, various frameworks are publicly available; e.g. Alchemy [1] for Markov Logic [Richardson and Domingos(2006)], Primula [2] for Relational Bayesian Networks and the Bayesian Logic (BLOG) Inference Engine [3]. We present a framework for Probabilistic Relational Models [Getoor(2000)], a directed graphical model that defines a language which can be used to describe the relationships - structural and probabilistic - between classes and variables, and thus allows the representation of dependencies between sets of objects.

Inference in PRMs is a hard problem. Recent work, e.g. [Milch et al. (2008), Kisynski & Poole (2009)], has proposed algorithms for exact inference in lifted (i.e. non-grounded) models using aggregation; Structured variable elimination introduced by [Pfeffer and Koller(2000)] is an extension of the original variable elimination method [Zhang and Poole(1996)] to the relational domain. However, exact inference is computationally expensive, and in the Bayesian networks community, it is often replaced by approximate inference methods, which can handle large models much better.

In [Kaelin and Precup(2010)] we presented an inference algorithm based on MCMC that leverages the relational structure to find an efficient sampling order. The algorithm is implemented as part of a framework for modelling PRMs called **P**rob**R**e**M** (`www.cs.mcgill.ca/~fkaeli/probrem/`; available now, public release planned for the workshop). **P**rob**R**e**M** allows the specification of PRM models for discrete variables, learns the model parameters from a relational database, and inference is made using MCMC.

Current (unpublished) work is based on [Pasula and Russell(2001)] and generalizes the presented approach for reference uncertainty in PRMs. Reference uncertainty introduces uncertainty about the structure of the data itself, e.g. the entries of a relationship table of an ER diagram, and thus the state space of the Markov Chain increases considerably. Our approach builds on the DAPER model [Heckerman and Koller(2004)], we associate a binary 'exist' variable with every possible entry in uncertain relationship tables. As the number of 'exist' attributes grows exponentially with the size of the tables, inference becomes intractable. We avoid the explosion of the state space by introducing a 'constraint' attribute that enforces certain structural properties, e.g. a *1:n* relationship. However, this results in complex probabilistic dependencies among the 'exist' objects. A more involved Metropolis-Hastings algorithm is required that samples 'exist' objects using an appropriate proposal distribution. A proposal is an assignment to all 'exist' objects associated with one 'constraint' object, which allows us to introduce probabilistic dependencies that would not be allowed in a traditional PRM. The 'constraint' can be seen as a fixed-parameter tractability approach which in general enforces a *k:n* relationship, for $k = 1$ this reduces to a *1:n* relationship which is common in relational databases. This new algorithm depends on reliable convergence diagnostics as one inference run can be seen as a sub-step of the full algorithm - similar to the Expectation Maximization algorithm used for missing data and latent variables.

In summary, we present the novel framework **P**rob**R**e**M** for modelling PRMs which relies on MCMC algorithms for inference.

[1] `http://alchemy.cs.washington.edu`
[2] `http://www.cs.aau.dk/~jaeger/Primula/`
[3] `http://people.csail.mit.edu/milch/blog/`

References

[Getoor(2000)] Lise Getoor. Learning probabilistic relational models. In *SARA '02: Proceedings of the 4th International Symposium on Abstraction, Reformulation, and Approximation*, pages 322–323. Springer-Verlag, 2000.

[Heckerman and Koller(2004)] D. Heckerman and D. Koller. Probabilistic models for relational data. Technical Report MSR-TR-2004-30, Microsoft Research, 2004.

[Kaelin and Precup(2010)] Fabian Kaelin and Doina Precup. An approach to inference in probabilistic relational models using block sampling. In *Asian Conference on Machine Learning (ACML)*, pages 325–340, 2010.

[Kisynski and Poole(2009)] J. Kisynski and D. Poole. Lifted aggregation in directed first-order probabilistic models. In *IJCAI*, pages 1922–1929, 2009.

[Milch et al.(2008)Milch, Zettlemoyer, Kersting, Haimes, and Kaelbling] B. Milch, L. Zettlemoyer, K. Kersting, M. Haimes, and L.P. Kaelbling. Lifted probabilistic inference with counting formulas. In *AAAI*, pages 1062–1067, 2008.

[Pasula and Russell(2001)] H. Pasula and S. Russell. Approximate inference for first-order probabilistic languages. In *IJCAI*, pages 741–748, 2001.

[Pfeffer and Koller(2000)] A.J. Pfeffer and D. Koller. Semantics and Inference for Recursive Probability Models. In *AAAI*, pages 538–544, 2000.

[Richardson and Domingos(2006)] M. Richardson and P. Domingos. Markov logic networks. *Machine Learning*, 62(1):107–136, 2006.

[Zhang and Poole(1996)] N.L. Zhang and D. Poole. Exploiting causal independence in bayesian network inference. *JAIR*, 5:301–328, 1996.