

Stochastic Planning and Lifted Inference

Roni Khardon

Department of Computer Science
Tufts University
roni@cs.tufts.edu

Abstract

The paper argues that (1) stochastic planning should be used as a core problem domain for relational probabilistic models providing problems of interest that are challenging for current approaches and significant scope for extending their capabilities, (2) that symbolic dynamic programming solving such problems can be seen as a prime example of lifted inference in relational probabilistic problems, (3) that first order decision diagrams provide a useful tool to drive such lifted computations, and (4) that the resulting lifted inference is qualitatively different from what other approaches are providing. As a result, this relationship can be studied to the benefit of developing foundations for relational probabilistic models and to the benefit of stochastic planning.

Introduction

The well known connection between planning and theorem proving suggests a similar analogy for the stochastic case. Indeed this idea has been explicitly explored by several groups (Attias 2003; Toussaint and Storsky 2006; Lang and Toussaint 2009) although the inference algorithms used do not take advantage of relational structure. On the other hand, interest in lifted inference within SRL has grown dramatically in recent years (Poole 2003; Braz, Amir, and Roth 2006; Milch et al. 2008; Singla and Domingos 2008; Kersting, Ahmadi, and Natarajan 2009; Kisynski and Poole 2009b; Sen, Deshpande, and Getoor 2009; Kisynski and Poole 2009a; Gordon, Hong, and Dudik 2009; Braz et al. 2009; Shavlik and Natarajan 2009). Two obvious questions are how to cast the relationship between lifted stochastic inference and stochastic planning and what are the best algorithmic outcomes of this correspondence. There are, in fact, several ways to look at this question from the perspective of each community and its set of tools, and these provide different insights. In the following we focus on just one perspective arising from seeing stochastic planning as a problem of optimizing Relational Markov Decision Processes (RMDPs). In particular, we argue that one approach, symbolic dynamic programming (SDP) (Boutilier, Reiter, and Price 2001), should be seen as a prime example of this analogy and that it provides an interesting example of lifted inference.

Copyright © 2010, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Markov decision processes (MDP) are models of decision making under uncertainty, and RMDPs are MDPs where the states, actions and transition probabilities can be described compactly by referring to objects and relations among them. Therefore the initial advantage of RMDPs is that they provide a compact description for complex problems. The main motivation for RMDPs is the thought that the compactness can be translated to improvement in run time over approaches that ignore the structure. The first formulation of this idea was given by Boutilier, Reiter, and Price (2001) who introduced the SDP algorithm as a means to perform lifted reasoning in RMDPs. The basic idea is that many states share the same values and transitions and therefore the Bellman equation can be calculated at the abstract level without enumerating states. This requires the algorithm to identify the right grouping of states dynamically while making the updates. It also requires lifted versions of goal regression for stochastic actions, as well as operations over relational value functions and probability functions.

This work was followed by different groups (Großmann, Hölldobler, and Skvortsova 2002; Hölldobler and Skvortsova 2004; Kersting, van Otterlo, and De Raedt 2004; Sanner and Boutilier 2009; Wang, Joshi, and Khardon 2008) proposing and using different representation schemes to support lifted computations for RMDPs. Our own work (Wang, Joshi, and Khardon 2008; Wang and Khardon 2007; Joshi, Kersting, and Khardon 2009; Wang and Khardon 2010) introduced first order decision diagrams (FODD) and showed how they can be used in such algorithms. Two challenges exist in such work, the first showing how a representation scheme can support the abstract calculations correctly and yet maintain compactness, and the second identifying ways to implement such schemes efficiently and deploy them to solve large problems. A solution based on relational linear functions approximation (Sanner and Boutilier 2009) was implemented and shown to solve problems from the International Planning Competition (IPC). The FODD solutions were implemented in the FODD-Planner and similarly applied to problems from the IPC (Joshi and Khardon 2008; Joshi, Kersting, and Khardon 2010). These efforts show that the approach can scale to problems of outside interest.

Notice that FODD-Planner, and in fact any implementation of SDP, performs calculation of values and probabilities “in bulk”. This is a standard desideratum for lifted inference.

A similar analogy can be made between FODDs representing intermediate steps of the computation and parfactors or their generalizations in lifted inference. However, SDP algorithms seem to go a step further in terms of abstraction. Unlike most lifted inference work, in SDP the network is not grounded to start with, and inference is performed as much as possible at the abstract level. In fact, for some problem domains and queries the answer does not depend on the number of objects and can be characterized abstractly, in which case the network is not grounded at all. SDP algorithms have been motivated by these ideal case scenarios but can adapt and ground the domain as needed.

Further insight can be obtained by looking at the details of the algorithm. One can interpret the calculations in SDP as taking advantage of two aspects of the model: the first is (context specific) conditional independence, just as in Bayesian networks. The other might be called “context specific dependence” where many predicates change simultaneously under some conditions. Thus both dependence and independence can be used for speedup and to enable abstract solutions. An interesting question is whether similar observations can be used to improve inference algorithms in networks with hard constraints, especially those arising from relational structure and thus repeating in the network. Our main point, however, is not that SDP is more general but that it provides a different perspective on a related problem and thus provides different tradeoffs and insights into fast calculation with relational models. Other implications from this analogy can be similarly drawn to suggest potential new algorithms.

The main engine underlying lifted computations is a good representation scheme. This representation must allow for a calculus of complex functions over structured domains; that is, we may want to add, multiply or apply some other operations to functions over structured domains. In our work, FODDs provide this representation and corresponding operations. In more recent work we have shown how to extend these ideas to capture more problems and algorithms and to speed up computations. These include relational policy iteration (Wang and Khardon 2007), relational partially observable MDPs (Wang and Khardon 2010), FODDs capturing more complex structured functions (Joshi, Kersting, and Khardon 2009), and improvements to compactness and run time by using “examples” (Joshi, Kersting, and Khardon 2010). These extensions suggest that the symbolic computation of SDP might be useful in more general contexts.

We have so far argued that there is an interesting analogy that is worth investigating. Our second main point is that a focus on planning might be a useful path for future work on inference algorithms. Current SDP solvers focus on simply-defined domains and they cannot capture action durations, continuous resources, continuous actions, and compositional and hierarchical problem structure. Extending relational probabilistic models and algorithms to capture such features would be an interesting challenge and will provide ample scope for improvements in inference algorithms.

Acknowledgments

This work was partly supported by NSF grant IIS 0936687.

References

- Attias, H. 2003. Planning by probabilistic inference. In *Proc. of the 9th Int. Workshop on Artificial Intelligence and Statistics*.
- Boutilier, C.; Reiter, R.; and Price, B. 2001. Symbolic dynamic programming for first-order MDPs. In *IJCAI*.
- Braz, R.; Amir, E.; and Roth, D. 2006. MPE and partial inversion in lifted probabilistic variable elimination. In *AAAI*.
- Braz, R.; Natarajan, S.; Bui, H.; Shavlik, J.; and Russell, S. 2009. Anytime lifted belief propagation. In *SRL Workshop*.
- Gordon, G.; Hong, S.; and Dudik, M. 2009. First-order mixed integer linear programming. In *UAI*.
- Großmann, A.; Hölldobler, S.; and Skvortsova, O. 2002. Symbolic dynamic programming within the fluent calculus. In *Proceedings of the IASTED International Conference on Artificial and Computational Intelligence*.
- Hölldobler, S., and Skvortsova, O. 2004. A logic-based approach to dynamic programming. In *AAAI-04 workshop on learning and planning in Markov Processes – advances and challenges*.
- Joshi, S., and Khardon, R. 2008. Stochastic planning with first order decision diagrams. In *ICAPS*.
- Joshi, S.; Kersting, K.; and Khardon, R. 2009. Generalized first order decision diagrams for first order markov decision processes. In *IJCAI*.
- Joshi, S.; Kersting, K.; and Khardon, R. 2010. Self-taught decision theoretic planning with first order decision diagrams. In *ICAPS*.
- Kersting, K.; Ahmadi, B.; and Natarajan, S. 2009. Counting belief propagation. In *UAI*.
- Kersting, K.; van Otterlo, M.; and De Raedt, L. 2004. Bellman goes relational. In *ICML*.
- Kisynski, J., and Poole, D. 2009a. Constraint processing in lifted probabilistic inference. In *UAI*.
- Kisynski, J., and Poole, D. 2009b. Lifted aggregation in directed first-order probabilistic models. In *IJCAI*.
- Lang, M., and Toussaint, M. 2009. Approximate inference for planning in stochastic relational worlds. In *ICML*.
- Milch, B.; Zettlemoyer, L.; Kersting, K.; Haimes, M.; and Kaelbling, L. P. 2008. Lifted probabilistic inference with counting formulas. In *AAAI*.
- Poole, D. 2003. First-order probabilistic inference. In *IJCAI*, 985–991.
- Sanner, S., and Boutilier, C. 2009. Practical solution techniques for first order mdps. *AIJ* 173:748–788.
- Sen, P.; Deshpande, A.; and Getoor, L. 2009. Bisimulation-based approximate lifted inference. In *UAI*.
- Shavlik, J., and Natarajan, S. 2009. Speeding up inference in markov logic networks by preprocessing to reduce the size of the ground network. In *IJCAI*.
- Singla, P., and Domingos, P. 2008. Lifted first-order belief propagation. In *AAAI*.
- Toussaint, M., and Storsky, A. 2006. Probabilistic inference for solving discrete and continuous state markov decision processes. In *ICML*.
- Wang, C., and Khardon, R. 2007. Policy iteration for relational MDPs. In *UAI*.
- Wang, C., and Khardon, R. 2010. Relational partially observable mdps. In *AAAI*.
- Wang, C.; Joshi, S.; and Khardon, R. 2008. First order decision diagrams for relational MDPs. *JAIR* 31:431–472.