Overview of Linear Program Approximations for Factored Continuous and Hybrid-State Markov Decision Processes

Milos Hauskrecht
Department of Computer Science
University of Pittsburgh
milos@cs.pitt.edu

Abstract

Approximate linear programming (ALP) is as one of the most promising methods for solving complex factored MDPs. The method was applied first to tackle problems with discrete state variables. More recently the ALP methods that can solve MDPs with continuous and hybrid (both continuous and discrete) variables have emerged. This paper briefly reviews the work on ALP methods for such problems.

1 Introduction

Markov decision processes (MDPs) offer an elegant mathematical framework for representing and solving sequential decision making-problems in the presence of uncertainty. Standard MDP solution techniques, such as value and policy iteration [1, 8] were developed for discrete state MDPs. These methods work and scale-up well in terms of the number of states of the MDP. However, the state space of more realistic MDP problems is factorized and the state space becomes exponential in the number of state components. Much of the work in the AI community has focused on factored structured representations of finite-state MDPs and their efficient solutions.

Approximate linear programming (ALP) has emerged as one of the most viable methods for solving complex factored MDPs. The approach uses a linear combination of local feature functions to model the value function. The coefficients of the model are fit using linear program methods. The approach was first applied to solve factored MDPs with discrete state components. A variety of methods for these settings include the work by Guestrin et al [5], de Farias and Van Roy [3, 2], Schuurmans and Patrascu [16], and others [15]. However, the approximate linear programming (LP) methods can be extended also to factored continuous-state MDPs and/or hybrid state MDPs that combine both continuous and discrete state components. These line of work was pioneered by Hauskrecht and Kveton. This work gives a brief review of these methods.

The text is structured as follows. First we review the basics of the MDP and its linear program solutions. After that we introduce ALP and show how one can formulate

factored MDPs as linear program approximation and how one can solve the problem more efficiently. After that we briefly review multiple different solutions related to ALP and its optimizations. This work is based in a number of publications in past years [6, 9, 4, 10, 7, 12, 11, 14, 13].

References

[1] Richard E. Bellman. *Dynamic programming*. Princeton University Press, Princeton, NJ, 1957.

[2] D. P. de Farias and B. Van Roy. On constraint sampling for the linear programming approach to approximate dynamic programming. *Mathematics of Operations Research*, submitted, 2001.

[3] D. P. de Farias and B. Van Roy. The linear programming approach to approximate dynamic programming. *Operations Research*, 51:6, 2003.

[4] Carlos Guestrin, Milos Hauskrecht, and Branislav Kveton. Solving factored MDPs with continuous and discrete variables. In *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence*, pages 235–242, 2004.

[5] Carlos Guestrin, Daphne Koller, and Ronald Parr. Max-norm projections for factored MDPs. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence*, pages 673–682, 2001.

[6] Milos Hauskrecht and Branislav Kveton. Linear program approximations for factored continuous-state Markov decision processes. In *Advances in Neural Information Processing Systems 16*, pages 895–902, 2003.

[7] Milos Hauskrecht and Branislav Kveton. Approximate linear programming for solving hybrid factored MDPs. In *Proceedings of the 9th International Symposium on Artificial Intelligence and Mathematics*, 2006.

[8] Ronald A. Howard. *Dynamic Programming and Markov Processes*. MIT Press, Cambridge, 1960.

[9] Branislav Kveton and Milos Hauskrecht. Heuristic refinements of approximate linear programming for factored continuous-state Markov decision processes. In *Proceedings of the 14th International Conference on Automated Planning and Scheduling*, pages 306–314, 2004.

[10] Branislav Kveton and Milos Hauskrecht. An MCMC approach to solving hybrid factored MDPs. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence*, pages 1346–1351, 2005.

[11] Branislav Kveton and Milos Hauskrecht. Learning basis functions in hybrid domains. In *Proceedings of the 21st National Conference on Artificial Intelligence*, pages 1161–1166, 2006.

[12] Branislav Kveton and Milos Hauskrecht. Solving factored MDPs with exponential-family transition models. In *Proceedings of the 16th International Conference on Automated Planning and Scheduling*, pages 114–120, 2006.

[13] Branislav Kveton and Milos Hauskrecht. Partitioned linear programming approximations for mdps. In *Proceedings of the 24th International Conference on Uncertainty in Artificial Intelligence*, pages 341–348, 2008.

[14] Branislav Kveton, Milos Hauskrecht, and Carlos Guestrin. Solving factored mdps with hybrid state and action variables. *Journal of Artificial Intelligence Research*, 27(1):153–201, 2006.

[15] P. Poupart, C. Boutilier, R. Patrascu, and D. Schuurmans. Piecewise linear value function approximation for factored mdps. In *Proceedings of the Eighteenth National Conference on AI*, pages 292–299, 2002.

[16] Dale Schuurmans and Relu Patrascu. Direct value-approximation for factored mdps. In T. G. Dietterich, S. Becker, and Z. Ghahramani, editors, *Advances in Neural Information Processing Systems 14*, Cambridge, MA, 2002. MIT Press.