

Multi-Agent Planning in Complex Uncertain Environments

Daphne Koller
Computer Science Department
Stanford University
Stanford, CA 94305-9010
koller@cs.stanford.edu

Abstract

Many tasks require a team of agents to act together in a coordinated way in a complex, uncertain environment. Examples include search and rescue, control of a complex system such as a factory, or robot soccer. Such tasks involve many agents, and huge numbers of states and possible actions. Factored Markov decision processes (MDPs) provide a formal foundation for modeling such complex systems naturally and compactly. We propose a framework for multi-agent coordination and planning in factored MDPs based on the use of a factored value function — an approximate decomposition of the team value function as a sum of value functions of small subteams of agents. We show that factored value functions naturally give rise to an optimal distributed algorithm for joint action selection, whose communication structure naturally mirrors the interactions between the subteams. We present an efficient linear-programming-based algorithm for computing a factored value function for a factored MDP. We show how the use of factored value functions can form the basis for inter-domain plan generalization, where we “learn” from plans constructed for some set of problems, and can provide good solutions to other unseen problems of the same type without any need for planning. We describe the application of this approach to the task of multi-agent planning in a strategic computer war game, and the application by Kok et al. for multi-agent coordination in their world-champion RoboSoccer team.

Acknowledgements. Joint work with Carlos Guestrin, Ronald Parr, Chris Gearhart, Neal Kanodia, and Shobha Venkataraman.

References

- [1] C. E. Guestrin, D. Koller, C. Gearhart, and N. Kanodia. Generalizing plans to new environments with relational MDPs. In

Proc. Eighteenth International Joint Conference on Artificial Intelligence (IJCAI-03), pages 1003 – 1010, 2003.

- [2] C. E. Guestrin, D. Koller, and R. Parr. Max-norm projections for factored MDPs. In *Proc. Seventeenth International Joint Conference on Artificial Intelligence (IJCAI-01)*, pages 673 – 680, 2001.
- [3] C. E. Guestrin, D. Koller, and R. Parr. Multiagent planning with factored MDPs. In *Proc. 14th Neural Information Processing Systems (NIPS-14)*, pages 1523–1530, 2001.
- [4] C. E. Guestrin, S. Venkataraman, and D. Koller. Context specific multiagent coordination and planning with factored MDPs. In *Proc. Eighteenth National Conference on Artificial Intelligence (AAAI-2002)*, pages 253–259, 2002.
- [5] J. R. Kok, M. T. J. Spaan, and N. Vlassis. Multi-robot decision making using coordination graphs. In *Proc. 11th International Conference on Advanced Robotics*, 2003.