Title

Exploiting Domain Structure in Multiagent Decision-Theoretic Planning and Reasoning

Author

Akshat Kumar, *University of Massachusetts Amherst* Follow

Date of Award

5-2013

Document Type

Open Access Dissertation

Degree Name

Doctor of Philosophy (PhD)

Degree Program

Computer Science

First Advisor

Shlomo Zilberstein

Second Advisor

Daeyoung Kim

Third Advisor

Victor Lesser

Subject Categories

Computer Sciences

Abstract

This thesis focuses on decision-theoretic reasoning and planning problems that arise when a group of collaborative agents are tasked to achieve a goal that requires collective effort. The main contribution of this thesis is the development of effective, scalable and quality-bounded computational approaches for multiagent planning and coordination under uncertainty. This is achieved by a synthesis of techniques from multiple areas of artificial intelligence, machine learning and operations research. Empirically, each algorithmic contribution has been tested rigorously on common benchmark problems and, in many cases, real-world applications from machine learning and operations research literature.

The first part of the thesis addresses multiagent single-step decision making problems where a single joint-decision is required for the plan. We examine these decision-theoretic problems within the broad frameworks of distributed constraint optimization and Markov random fields. Such models succinctly capture the structure of interaction among different decision variables, which is subsequently exploited by algorithms to enhance scalability. The algorithms presented in this thesis are rigorously grounded on concepts from mathematical programming and optimization.

The second part of the thesis addresses multiagent sequential decision making problems under uncertainty and partial observability. We use the decentralized partially observable Markov decision processes (Dec-POMDPs) to formulate multiagent planning problems. To address the challenge of NEXP-Hard complexity and yet push the envelope of scalability, we represent the domain structure in a multiagent system using graphical models such as dynamic Bayesian networks and constraint networks. By exploiting such graphical planning representation in an algorithmic framework composed of techniques from different sub-areas of artificial intelligence, machine learning and operations research, we show impressive gains in increasing the scalability, the range of problems addressed and enabling quality-bounded solutions for multiagent decision theoretic planning.

Our contributions for sequential decision making include a) development of efficient dynamic programming algorithms for finite-horizon decision making, resulting in significantly increased scalability w.r.t. the number of agents and multiple orders-of-magnitude speedup over previous best approaches; b) development of probabilistic inference based algorithms for infinite-horizon decision making, resulting in new insights connecting inference techniques from the machine learning literature to multiagent systems; c) development of mathematical programming based scalable techniques for quality bounded solutions in multiagent systems, which has been considered intractable so far.

Several of our contributions are some of the first for the respective class of problems. For example, we show for the first time how machine learning is closely related to multiagent decision making via a maximum likelihood formulation of the planning problem. We develop new graphical models and machine learning based inference algorithms for large factored planning problems. We also show for the first time how the problem of optimizing agents' policies can be formulated as a compact mixed-integer program, resulting in optimal solution for a range of Dec-POMDP benchmarks.

In summary, we present a synthesis of different techniques from multiple sub-areas of AI, ML and OR to address the scalability and efficiency of algorithms for decision-theoretic reasoning and planning in multiagent systems. Such advances have already shown great promise to bridge the gap between multiagent systems and real-world applications.

Recommended Citation

Kumar, Akshat, "Exploiting Domain Structure in Multiagent Decision-Theoretic Planning and Reasoning" (2013). *Open Access Dissertations*. 734.
https://scholarworks.umass.edu/open_access_dissertations/734