

Stochastic Planning in Large Search Spaces

Bilal Kartal

Department of Computer Science and Engineering
University of Minnesota
bilal@cs.umn.edu

Abstract

Multi-agent planning approaches are employed for many problems including task allocation, surveillance and video games. In the first part of my thesis, we study two multi-robot planning problems, i.e. patrolling and task allocation. For the patrolling problem, we present a novel stochastic search technique, Monte Carlo Tree Search with Useful Cycles, that can generate optimal cyclic patrol policies with theoretical convergence guarantees. For the multi-robot task allocation problem, we propose an Monte Carlo Tree Search based satisficing method using branch and bound paradigm along with a novel search parallelization technique. In the second part of my thesis, we develop a stochastic multi-agent narrative planner employing MCTS along with new heuristic and pruning methods applicable for other planning domains as well.

1 Introduction

Large-scale planning has many application areas such as surveillance, patrolling, and games. There has been great improvement in large-scale planning with sampling based methods. One example is Monte Carlo Tree Search (MCTS), an anytime stochastic search algorithm shown to perform well for large planning domains. Particularly, the Go game has been one of the most challenging problems in AI lately and recently an MCTS based approach remarkably beat a top human professional on a full size Go board [Silver *et al.*, 2016].

Throughout my PhD work I have studied various applications of large-scale stochastic planning, including multi-robot patrolling [Kartal *et al.*, 2015], multi-robot task allocation [Kartal *et al.*, 2016], and multi-agent narrative generation [Kartal *et al.*, 2013; 2014]. The multi-robot planning problems, particularly patrolling and task allocation, require high coordination among the robots for efficient behavior. Both patrolling and task allocation problems are \mathcal{NP} -hard and the multi-robot task allocation problem is a much more generalized version of many multi-robot coordination problems. The multi-agent narrative generation problem aims to generate coherent and believable narratives which are potentially applicable in entertainment, education, and training systems.

My thesis will show that sampling-based stochastic search techniques can be adapted to efficiently find approximate solutions to large-scale multi-agent planning problems that arise in a variety of domains including both physical planning for robot patrolling, and task allocation and abstract planning for narrative generation. Sparse sampling based stochastic approaches can be successfully applied to domains with large state spaces, and their sampling-based nature allows flexibility in the details of the problem formulation and supports quick feedback for the user-in-the-loop interaction.

2 Contributions

The stochastic planning techniques developed as part of my thesis has many important applications (Figure 1).

2.1 Multi-Robot Patrolling Problem

We proposed Monte Carlo Tree Search with Useful Cycles (MCTS-UC) [Kartal *et al.*, 2015], a novel stochastic search technique that can generate infinite length cyclic policies with theoretical convergence guarantees for non-adversarial patrolling problem with probabilistic intrusion models. Our approach expands the state space with artificial cyclic nodes to compactly represent infinite length policies. Our approach improves the state of the art in the sense that it can be easily adapted to different intrusion models and arbitrary environments as long as an MDP simulator exists for the formulated problem. A video presenting patrolling policies in different environments can be seen at <http://motion.cs.umn.edu/r/MCTS-UC>.

2.2 Multi-Robot Task Allocation

We proposed an efficient, satisficing and centralized Monte Carlo Tree Search based algorithm using the branch and bound paradigm with a novel search parallelization method to solve the multi-robot task allocation problem. We proposed a new evaluation function to balance between task completion and total traveled distance minimization. We coupled our evaluation function with a novel search parallelization technique, parameterized root parallelization. Here, we create a separate search tree per thread with varying exploration parameters with a globally shared best complete policy distance. The design idea is that highly exploiting threads can find a sub-optimal solution completing all the tasks very early during the search. Since this solution is shared among all threads,




Figure 1: Overview of our contributions for story generation [Kartal *et al.*, 2014], patrolling [Kartal *et al.*, 2015], and task allocation [Kartal *et al.*, 2016]: **(Left)** The GUI for our story generator is shown which enables the user-in-the-loop to alter story domain parameters interactively to interact with the planner. **(Center)** The robots employ an optimal patrol policy obtained with our MCTS-UC approach. **(Right)** A near-optimal task allocation policy found by our parallelized approach with an approximation rate of 1.03 to an optimal solution is shown.

both this will better calibrate all the threads' search directions and bounding many branches guaranteed to be not better than the best found will improve memory efficiency due to pruning. Unlike previous heuristics tailored for specific problems, our approach is general, it has asymptotic convergence guarantees, and it finds near-optimal solutions for non-trivial problems in an existing test benchmarks within an hour.

2.3 Multi-Agent Narrative Generation

We modeled the narrative generation problem as a multi-agent planning problem and proposed an MCTS based user-driven story planner along with several heuristics based on search biasing and tree pruning. Our approach [Kartal *et al.*, 2013; 2014] aims to generate more *believable* narratives given user defined story goals. We assumed an expert user has authored a PDDL based logic relations for the story actions along with *believability* of different actions. We present an example output of our planner for the *Detective* story domain in Figure 2. The main novelty of our approach improving the state-of-the art is its ability to allow the user to interact with the planner by changing believability of actions and interact with the planner in real-time as MCTS is an anytime approach. A video showing our story planner running can be found at: <http://motion.cs.umn.edu/tr/StoryMCTS>.

Alice picked up a vase from her house. Bob picked up a rifle from his house. Bob went to Alices house. While there, greed got the better of him and Bob stole Alices vase! This made Alice furious. Alice pilfered Bobs vase! This made Bob furious. Bob slayed Alice with a rifle! Bob fled to downtown. Bob executed Inspector Lestrade with a rifle! Charlie took a baseball bat from Bobs house. Sherlock went to Alices house. Sherlock searched Alices house and found a clue about the recent crime. Bob fled to Alices house. Sherlock wrestled the rifle from Bob! This made Bob furious. Sherlock performed a citizens arrest of Bob with his rifle and took Bob to jail.

Figure 2: A story generated in less than a minute.

3 Conclusions and Future Work

The main contributions of my thesis is to show how sampling-based stochastic search methods can be efficiently adapted for multi-agent problems both in physical domains for the multi-robot patrolling and task allocation, and in abstract domain for the computer narrative generation problem. We show that the sampling based nature proves to be very useful for planning problems as simulation-driven mechanisms permit modification to planning problem details with no changes to the algorithm, e.g. different intruder models can be used for the patrolling problem, and action believabilities can be altered for the narrative generation problem. This ability renders the approach more general and easily adaptable to different problem variants. Also, the MCTS algorithm is a complete search method guaranteeing optimality along with the anytime property for flexible computation cost. There are many avenues for future work such as exploring new parallelization techniques and integration of MCTS with control theory.

References

[Kartal *et al.*, 2013] Bilal Kartal, John Koenig, and Stephen J Guy. Generating believable stories in large domains. In *Intelligent Narrative Technologies 6, AIIDE*, 2013.

[Kartal *et al.*, 2014] Bilal Kartal, John Koenig, and Stephen J Guy. User-driven narrative variation in large story domains using Monte Carlo Tree Search. In *AAMAS*, pages 69–76, 2014.

[Kartal *et al.*, 2015] Bilal Kartal, Julio Godoy, Ioannis Karamouzas, and Stephen J Guy. Stochastic tree search with useful cycles for patrolling problems. In *ICRA*, pages 1289–1294, 2015.

[Kartal *et al.*, 2016] Bilal Kartal, Ernesto Nunes, Julio Godoy, and Maria Gini. Monte Carlo Tree Search for multi-robot task allocation. In *AAAI*, 2016.

[Silver *et al.*, 2016] David Silver, Aja Huang, Chris J Maddison, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.