Efficient Planning under Uncertainty for a Target-Tracking Micro-Aerial Vehicle

Abraham Bachrach ABACHRAC@MIT.EDU
Ruijie He RUIJIE@ALUM.MIT.EDU
Nicholas Roy NICKROY@CSAIL.MIT.EDU
Massachusetts Institute of Technology, Computer Science and Artificial Intelligence Laboratory 32 Vassar St., Cambridge, MA 02139 USA

Abstract

A helicopter agent has to plan trajectories to track multiple ground targets from the air. The agent has partial information of each target's pose, and must reason about its uncertainty of the targets' poses when planning subsequent actions.

We present an online, forward-search algorithm for planning under uncertainty by representing the agent's belief of each target's pose as a multi-modal Gaussian belief. We exploit this parametric belief representation to directly compute the distribution of posterior beliefs after actions are taken. This analytic computation not only enables us to plan in problems with continuous observation spaces, but also allows the agent to search deeper by considering policies composed of multi-step action sequences; deeper searches better enable the agent to keep the targets well-localized. We present experimental results in simulation, as well as demonstrate the algorithm on an actual quadrotor helicopter tracking multiple vehicles on a road network constructed indoors.

1. Introduction

This invited abstract summarizes results that appeared at the IEEE International Conference in Robotics and Automation (ICRA) in 2010.

MAVs are increasingly used in military and civilian domains ranging from intelligence, surveillance and reconnaissance operations, border patrol missions, to weather observation and disaster relief coordination efforts. In this paper, we present a target-tracking planning algorithm for a helicopter maintaining surveillance over multiple targets along road networks, such as an autonomous police helicopter tasked with monitoring the activity of several suspicious cars in urban environments.

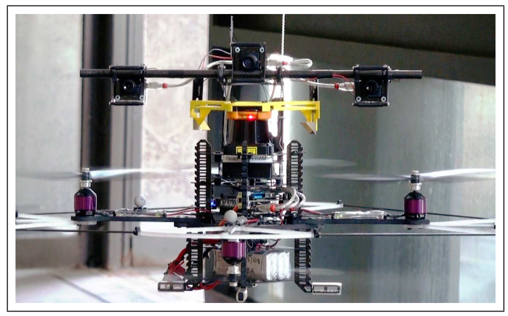

Figure 1. Our quadrotor helicopter, equipped with stereo cameras and a laser range finder, capable of flight in GPS-denied environments such as indoors or the urban canyon.

Target-tracking is a sequential decision-making task that combines target-search – finding targets that are not initially visible, and target-following – maintaining visibility of the discovered targets. As the agent does not have perfect information of the targets' poses and their subsequent actions, it has to reason about its belief of their poses when planning to keep them well-localized. Especially when multiple targets have to be tracked by a single agent, the agent has the additional challenge of reasoning about which target to concentrate on at every timestep.

We approach this planning under uncertainty, target-tracking problem by planning in an online, forward-search manner. Online forward-search methods have demonstrated promising results in problems with large domains (see Ross et al., 2008 for a review), suggesting that planning under uncertainty can be performed efficiently by only considering the belief states that are reachable from the agent's current belief. Unfortunately, the number of reachable belief states grows exponentially with the search depth, making the search quickly intractable for long planning horizons.

We have shown in previous work (He et al., 2011) that for uni-modal Gaussian representations, the sufficient statistics for how the agent's belief are expected to evolve as actions

Appearing in *Proceedings of the 28^{th} International Conference on Machine Learning*, Bellevue, WA, USA, 2011.

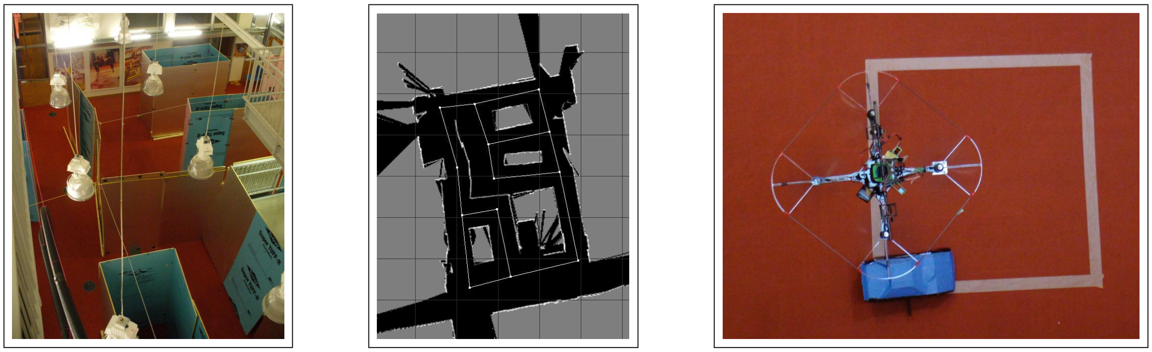

(a) (b) (c)

Figure 2. (a) Mock-up of road network constructed indoors. (b) Graph structure extracted from SLAM map. (c) Helicopter tracking car through an area of interest.

are taken can be directly computed without having to enumerate the possible observations. This property allows the planner to search deeper by considering plans composed of multi-step action sequences, commonly referred to as macro-actions. We extend our previous algorithm to tracking multiple targets by representing the agent's belief of each target's pose as a multi-modal Gaussian belief. For this belief representation, we show that we can approximate the distribution of posterior beliefs at the end of a macro-action in an efficient manner, thereby enabling the planner to search deeper.

Table 1 summarizes the performance of our strategy (MMPBD) compared to POMDP forward search (FS) and a greedy information-gathering strategy over 10 runs in simulation. Although our algorithm results in the longest distance traveled, the agent is better able to localize the targets, ensuring that its belief of each target is uni-modal most of the time. The agent anticipates when modes will split and can search deep enough to recognize that if it does not localize the target before this split, it will be harder to localize the targets subsequently. In contrast, the naive forward search algorithm causes the agent to travel the shortest distance because it focuses on trying to localize a single target.

	Dist. Traveled	Ave. Modes	Total Cost
MMPBD	138.76	1.0810	-51.7698
FS	112.20	1.7747	-85.1101
Greedy	133.52	1.5240	-61.2492

Table 1. Performance of different planners over 10 runs.

We also demonstrate our algorithm on an actual quadrotor helicopter (Figure 1) with a downward pointing camera for tracking ground vehicles from the air. Our helicopter is capable of autonomous flight in unstructured and unknown indoor environments (Bachrach et al., 2009). This real-world experiment validates the simulation results, showing that the vehicle can keep the different targets well-localized. We set up a mock-up of a road network indoors (Figure 2a), with two autonomous cars driven at approximately constant speeds in the environment. The helicopter exhibited behaviors similar to those seen in simulation, moving between the different targets to keep them well-localized.

This work relies on domain knowledge to generate macro-actions, however, we have presented initial results for a domain-independent macro-action algorithm (PUMA; He et al., 2010) that automatically generates macro-actions for planning in partially-observable domains. Borrowing the notion of sub-goal states from the fully-observable planning literature (McGovern, 1998), PUMA uses a heuristic that macro-actions can be designed to take the agent, under the fully-observable model, from a possible start state under the current belief to a sub-goal state. We plan to integrate the PUMA algorithm with the forward search strategies in the future.

Acknowledgments

We would like to thank the organizers of ICML for the invitation to present this work. Ruijie He, Abraham Bachrach and Nicholas Roy were supported by NSF IIS under grant # 0546467, the ARO MAST CTA and ONR under MURI N00014-07-1-0749.

References

Bachrach, A., He, R., and Roy, N. Autonomous flight in unstructured and unknown indoor environments. In *Proc. EMAV*, 2009.

He, R., Brunskill, E., and Roy, N. PUMA: planning under uncertainty with macro-actions. In *Proceedings of the Association for the Advancement of Artificial Intelligence (AAAI)*, 2010.

He, R., Brunskill, E., and Roy, N. Efficient planning under uncertainty with macro-actions. *Journal of Artificial Intelligence Research*, 40:523–570, February 2011.

McGovern, A. acQuire-macros: An algorithm for automatically learning macro-actions. In *NIPS 98 Workshop on Abstraction and Hierarchy in Reinf. Learning*, 1998.

Ross, S., Pineau, J., Paquet, S., and Chaib-draa, B. Online planning algorithms for POMDPs. *JAIR*, 32:663, 2008.