Motion Planning in Non-Gaussian Belief Spaces

Saurav Agarwal, Amirhossein Tamjidi, Suman Chakravorty

Abstract

In environments with information symmetry, uncertain or ambiguous data associations can lead to a multi-modal hypothesis on the robot's state. Thus, a planner cannot simply base actions on the most-likely state. We propose an algorithm that uses a Receding Horizon Planning approach to plan actions that sequentially disambiguate a multi-modal belief to a uni-modal Gaussian and achieve tight localization on the true state of a mobile robot. We call this algorithm Multi-Modal Motion Planner (M3P). We prove that our planner is guaranteed to drive a multi-modal belief to a uni-modal Gaussian under certain assumptions. Simulation results for a 2D ground robot navigation problem are presented that demonstrate our method's performance.

Introduction

Planning under uncertainty is a key requirement for physical systems due to the noisy nature of actuators and sensors. Using a belief space approach, planning solutions tend to generate actions that result in information seeking behavior which reduce state uncertainty. In general, planning for systems under uncertainty belongs to the class of Partially-Observable Markov Decision Process (POMDP) problems which are known to be computationally intractable. In this regard, sampling based methods such as (Prentice and Roy 2009) and (Agha-mohammadi, Chakravorty, and Amato 2014) have shown varying degrees of success. A common thread among these methods is that they rely on a Gaussian representation of the robot's belief.

Issues with Gaussian Belief Representation: Imagine a mobile robot equipped with a laser scanner operating in a world where there are identical rooms as shown in Fig. 1. To the laser, each and every room appears identical. Based on the observations, the robot may believe itself to be in one of multiple states at the same time. Another situation could be that of a robot equipped with only vision sensors, operating in an environment where there are identical looking features in multiple locations (e.g. warehouses). Such a situation implies that the pdf of the robot's belief cannot be represented by a uni-modal Gaussian. Thus, data association can lead to a non-Gaussian belief.

It has been shown in (Dudek, Romanik, and Whitesides 1998) that the problem of eliminating a multi-modal hypothesis of the robot's location using minimum travel distance, is NP-hard. In this work, we propose a computationally feasible, greedy heuristic-based approach to the problem of localization for a multi-modal belief.

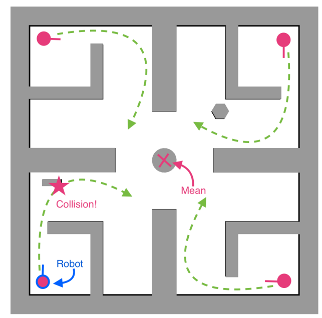

Figure 1: A robot (blue outlined disk) in a world with 4 identical rooms. All hypothesis (red disks) are equally likely. The dashed green arrows show a possible control action based on the hypothesis in the top-right room that can result in collision for the actual robot.

Problem statement

We represent the belief b_k by a Gaussian Mixture Model (GMM) at time k as a weighted linear summation over Gaussian densities,

$$b_k = \sum_{i=1}^{N} w_{i,k} m_{i,k}, \;\; m_{i,k} \sim \mathcal{N}(\mu_{i,k}, \Sigma_{i,k}). \qquad (1)$$

Our goal is to construct a belief space planner $\mu(b_k)$ such that under the belief space planner, given any initial multi-modal belief b_0, the belief state process evolves such that $b_T = m_T$, where $m_T = \mathcal{N}(\mu_T, \Sigma_T)$ for some finite time T.

In other words, our goal is to construct a belief space planner such that it is guaranteed to drive the initial multi-modal belief into a uni-modal belief in finite time.

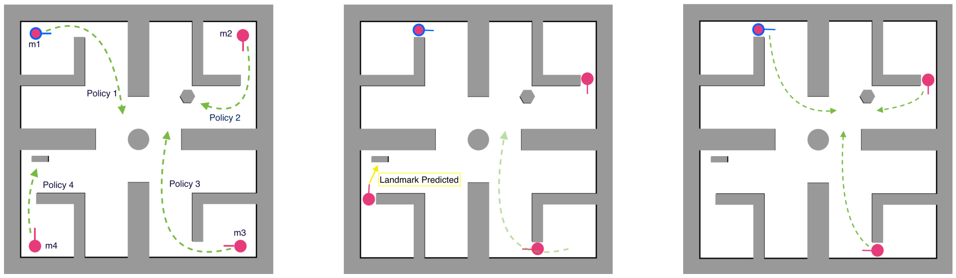

(a) A candidate policy is generated for each mode and we pick the policy with maximum expected gain.

(b) Policy 3 is chosen and executed. Mode m_4 expects to see a landmark outside the door which the robot does not see.

(c) Mode m_4 is rejected and a new set of candidate policies is computed.

Figure 2: Key steps in the Receding Horizon Control methodology for a multi-modal belief.

Methodology

M3P generates a candidate policy for each belief mode to guide it to some neighborhood state such that disambiguating information can be observed. Then, the planner picks the best policy out of the set, based on the expected information gain. This allows us to either prove or disprove a mode in the hypothesis. Fig. 2 shows how the planner disambiguates a multi-modal belief in a sequential manner in a hypothetical situation. Algorithm 1 describes the working of M3P, the key steps in the M3P planner are 3 and 4.

Algorithm 1: M3P: Multi-Modal Motion Planner

1 Input: b (Belief)
2 **while** $b \neq \mathcal{N}(\mu, \Sigma)$ **do**
3 Π = Generate candidate policies for belief modes ;
4 π = Pick policy from Π with maximum expected information gain ;
5 **forall the** $u \in \pi$ **do**
6   b = Apply action u and update belief;
7   **if** *Change in number of modes* $||$ *Expect a belief mode to violate constraints* **then**
8   break;
9 **return** b;

Analysis and Results

We show that the basic receding horizon planner M3P will guarantee that an initial multi-modal belief is driven into a uni-modal belief in finite time under the following assumptions:

Assumption 1 *Given a multi-modal belief* $b_k = \sum_i w_{i,k} m_{i,k}$*, for every mode* $m_{i,k}$*, there exists a disambiguating planner* $\mu_i(.)$ *in the sense that if the robot was really in mode* i*, the planner's actions would confirm that the robot was at mode* i*.*

Assumption 2 *The map does not change during the execution of the planner.*

Further, we present simulation results for a 2D ground robot. The simulation environment is designed such that there are multiple locations with duplicate landmarks (beacons), which leads to ambiguous data associations. The simulations represent two motion planning scenarios wherein the robot is tasked to go from a start to a goal location in an environment where there is information symmetry. Enroute to the goal, the robot is kidnapped to an unknown location such that the sensory observations lead to a multi-modal hypothesis. Thus, it relies on a node-based instantiation of M3P called NBM3P to localize the robot. Once the robot is localized (i.e. the belief converges to a uni-modal Gaussian), its new belief is connected to an existing FIRM (Agha-mohammadi, Chakravorty, and Amato 2014) graph and we switch from M3P to FIRM and find a new feedback policy to complete the task. A video of the simulation results is available at *https://www.youtube.com/watch?v=8kldm34Ya4I*.

Conclusions

Our main contribution in this work is a planner M3P that generates a sequentially disambiguating policy, which leads the belief to converge to a uni-modal Gaussian. We are able to show in simulation that the robot is able to recover from a kidnapped state (multi-modal belief) and execute its task in environments that present multiple uncertain data associations. Future work would involve demonstrating M3P on a physical system and examining closely the difference between the policy generated by M3P and the optimal relocalizing policy.

References

Agha-mohammadi, A.; Chakravorty, S.; and Amato, N. 2014. FIRM: Sampling-based feedback motion planning under motion uncertainty and imperfect measurements. *International Journal of Robotics Research (IJRR)* 33(2):268–304.

Dudek, G.; Romanik, K.; and Whitesides, S. 1998. Localizing a robot with minimum travel. *SIAM Journal on Computing* 27(2):583–604.

Prentice, S., and Roy, N. 2009. The belief roadmap: Efficient planning in belief space by factoring the covariance. *International Journal of Robotics Research* 28(11-12).