

TrustMeIA Internship Proposal

Assumption-based planning with incomplete information about the environment: How to act safely ?

Advisors:

Alexandre Albore, alexandre.albore@onera.fr
Caroline Chanel, caroline.chanel@isae-supaero.fr

Description

Real world applications of Artificial Intelligence impose to deal with uncertainty and partially knowledge about the environment. In Robotics, these difficulties are increased by sensors failure, exogenous events, and, more generally, by the impossibility to access the intelligent agent's state completely/perfectly.

However, producing plans in such environments is central for the development of intelligent agents (e.g. autonomous robots). Often this hindrance is overcome by imposing strong assumptions on robot sensing capabilities and on the dynamic behaviour of the environment, with the aim of reducing the uncertainty and eventually producing a plan that solves the task at hand. Safety can be guaranteed by searching for plans that allow to observe if the initial assumptions still hold before executing a risky maneuver (Albore and Bertoli, 2006), or by enriching the problem with observations about the validity of actions preconditions (Ponzoni Carvalho Chanel and Teichteil-Königsbuch, 2013).

Even if these approaches have encountered some success, on one hand, they lack the ability of quantifying the risk or even to automatically provide initial safe assumptions on the dynamics of the environment, and on the other hand, they assume to have available *ad hoc* sensing actions to check whether actions preconditions hold at run-time, which is not always possible.

Nevertheless, a safe approach to robotic planning would impose to have a clear and explainable risk evaluation, with an execution strategy supported by a wide variety of verification and validation tools.

Objective

Our aim is to extend current approaches to AI safe automated planning with algorithms that can guarantee that a plan is executable in an uncertain environment, given a risk threshold applied to the (automatically-generated) assumptions on the environment. Assessing the risk for an agent to fail to reach its goal would provide more flexibility to employ intelligent robotic agents in various missions, and to provide to the human user an explanation of the AI's decision process, and, eventually helping validating the adopted strategy.

The problem of planning in partially observable environments can be regarded as a search problem in belief space where beliefs express the collection of states that are deemed possible.

Said that, it is common, when planning with incomplete information, that the goal of a mission can be simply unreachable for all the possible situations. Or, that there is a risk on failing (dead-end) when applying a given action depending on the current context. An alternative is to apply contingent planning under partial observability, where the agent can condition its choices after applying sensing actions, in order to observe some properties to leverage

ambiguities before acting. However, contingent planning doesn't help in ensuring goal reachability, and neither does the introduction of probabilities to consider risk expectation.

For example, if a wheeled robot has to move to a target while sensing if adjacent cells are free or not, and the map is not (initially) known, the (hidden) state where there is no free path to the target is indeed possible and can be considered as a dead end (a terminal state with infinite cost). Indeed, the problem can be cast as a Partial Observable Markov Decision Process¹ (POMDP) by filling in the probabilities that each cell is free and maintaining the goal of reaching the target with certainty (Cassandra et al., 1994). Yet even then, the expected cost of reaching this target belief will be infinite as there is a nonzero probability that the paths to the goal are blocked. Even by bounding the expected costs and rewards or by discounting them, meaningful safe policies are difficult to produce. Of course, the robot shouldn't get paralysed in such cases; it should move toward the target and give up only when there is the certainty that it is unreachable.

Inspired by previous approaches, such as (Koenig and Smirnov, 1997; Albore and Bertoli, 2006; Albore and Geffner, 2009), we proposed to extend the assumption-based approach to planning under uncertainty and partial observability, by 1) automatically generating and evaluating the assumptions to impose to the planning process, if the goal appears not to be 100% reachable initially; 2) reevaluating the risk of such assumptions at run-time, 3) guaranteeing for a solution-plan to reach the goal under assumption within a well defined "safety envelope".

A solution-plan will then be for us either a plan for which : (i) it is possible to observe its safe executability under assumptions; (ii) or a plan under the safest assumption that guarantees its execution withing a "safety envelope" defined by the user for instance.

Moreover, monitoring the robot execution of the plan would allow to faithfully trace success-endangering assumption failures, providing the necessary traceability to the initial assumptions and design.

This approach is close to the continuous proactive planning with multiple hypotheses (T'Hooft et al., 2016), where multiple solution-plans are generated so to have a database from which an applicable action can potentially be selected when required.

Candidate's profile:

M2 student with programming skills and a background in robotics and AI is suitable. The candidate will study/learn the automated planning theory, assess what is a "safety envelope" for execution in various application contexts, and develop a robotic navigation application implementing the algorithms in our robotic lab.

References

- Albore, A. and Bertoli, P. (2006). Safe LTL assumption-based planning. In *Int. Conf. Aut. Planning and Scheduling (ICAPS)*, pages 193–202.
- Albore, A. and Geffner, H. (2009). Acting in partially observable environments when achievement of the goal cannot be guaranteed. In *Proc. of ICAPS Workshop on Planning and Plan Execution for Real-World Systems*.
- Cassandra, A. R., Kaelbling, L. P., and Littman, M. L. (1994). Acting optimally in partially observable stochastic domains. In *AAAI*, volume 94, pages 1023–1028.
- Koenig, S. and Smirnov, Y. (1997). Sensor-based planning with the freespace assumption. In *Proceedings of International Conference on Robotics and Automation*, volume 4, pages 3540–3545. IEEE.
- Ponzoni Carvalho Chanel, C. and Teichteil-Königsbuch, F. (2013). Properly acting under partial observability with action feasibility constraints. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 145–161. Springer.
- T'Hooft, J., Lesire, C., and Ponzoni Carvalho Chanel, C. (2016). Online proactive planning with multiple hypotheses. In *STAIRS*.

¹POMDP is a framework for sequential decision-making under uncertainty about action effects and partial observability on states