Asking Human Help in Contingent Planning

(Doctoral Consortium)

Ignasi Andrés Franch
Supervisor: Prof. Leliane Nunes de Barros
Department of Computer Science
IME-USP, Brazil
ignasi@ime.usp.br

ABSTRACT

Contingent planning models a robot in a partially observable environment and (non)deterministic actions. In a contingent planning problem, a solution can be found by doing a search in a space of *belief states*, where a belief state is represented by a set of possible states. However, in the presence of *dead-end belief states*, a situation where a robot may fail to complete its task, the only way to accomplish it is to ask for human help. In this work, rather than limiting a contingent planning task to those that only include actions and observations that an agent can perform autonomously, agents can instead reason about asking humans in the environment for help in order to complete tasks that would be unsolvable otherwise. We are interested in develop agents with *symbiotic autonomy*: an agent that proactively and autonomously ask for human help. We make the assumption that humans may only be interrupted when it is extremely necessary, i.e. when the planner can not solve a task or solving it implies in a too high cost. To solve this problem we propose an extension of an existing translation technique [2] that translates a contingent planning problem into a non-deterministic *fully observable* planning problem. The proposed translation can deal with different types of *dead-end belief states*, including a pure dead-end (e.g. a broken robot) and a dead-end that is due to uncertainty about the environment.

Keywords

Artificial Intelligence; Automated Planning; Robot Planning; Human-Robot Collaboration

1. CONTINGENT PLANNING

In a contingent planning problem, where the state is partially observable, a solution can be found searching in a space of *belief states*. A belief state represents a set of possible states of the world. In this search, an action (or observation) is applicable in a belief state b if it is applicable for every state $s \in b$. One of the approaches for contingent planning relies on epistemic translations (named *k-translation*) of a contingent planning problem into a FOND (*Fully Observable Non Deterministic*) planning problem, and uses a state-of-the-art FOND planner to find a solution.

Appears in: *Proc. of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2017), S. Das, E. Durfee, K. Larson, M. Winikoff (eds.), May 8–12, 2017, São Paulo, Brazil.*

DEFINITION 1 (CONTINGENT PLANNING PROBLEM). *A Contingent Planning Problem $\mathcal{P}$ is a tuple $\langle F, I, A, O, G\rangle$ where [1]:*

- *F is a set of fluent symbols of the problem;*
- *I is a set of clauses over F that defines the initial state b_0, partially known;*
- *A is a set of deterministic or non-deterministic actions specified in terms of preconditions and effects (used to induce the set of successor belief states of a belief state b);*
- *O is a set of observations (also specified in terms of preconditions and effects) and*
- *G is a conjunction of atoms over F that defines the planning goal.*

The solution of a contingent planning problem $\mathcal{P}$ is a contingent plan, i.e. a policy π that maps belief states b into actions $a \in A$. Notice that this policy induces a hipertree where nodes are labeled with belief states and hiperedges with actions, a branch for each outcome of a nondeterministic action or observation.

2. DEAD-END BELIEF STATES

A dead-end state is a state from which the agent can not achieve the goal. Thus, a contingent planning problem can have three types of solutions: a *weak* plan, with no guarantees to achieve the goal; a *strong* plan, which guarantees to achieve the goal and; a *strong cyclic* plan that eventually achieves the goal, despite the cycles. We say that π is a strong (cyclic) solution of a contingent problem $\mathcal{P}$, iff every execution of π is applicable and finishes in a belief state b' where for all states in b', the goal G holds. We say that π is a weak solution of a contingent problem $\mathcal{P}$, iff there is at least one execution of π that finishes in a dead end belief state b'. In this proposal we assume that a robot can only accomplish its task if it can find a strong (cyclic) plan, that can include human help for observation and actuation when necessary.

Thus, in a contingent planning problem $\mathcal{P}$ with a unavoidable belief state b that is a *dead-end*, there is not a strong (cyclic) plan starting from b that is a solution for P. A *dead-end belief state* b can be of three types: (DEB1) the states that belong to b are all pure dead-end states, i.e. states from which there is no weak or strong (cyclic) plan; (DEB2) b contains only *some* pure dead-end states and there is a *weak contingent plan*, i.e., a plan starting from b that can reach the goal but with no guarantees; (DEB3) b contains none or only *some* pure dead-end states, but there is none action that can be applied in b, i.e. the uncertainty in b is such that the agent can not come up with a strong or weak plan, unless it can ask for external help. E.g., a service robot in any of the three dead-end situations can proactively ask human to help [6, 7]: in cases (DEB1) and (DEB2) a human could be asked to modify the environment so the agent can continue to plan (for instance, to open a door when the robot is locked in an aisle); in case (DEB3) a human can help to reduce the agent uncertainty in problems by performing an observation the agent can not do.

3. PRELIMINARY RESULTS

3.1 Human help for observation

We first propose a way to solve contingent planning problems with dead-end belief state b of type DEB3. We assume the human can perform any observation that the agent is designed to do in the actual planning environment, but with a relaxation of its preconditions and with high cost; then the planner must find a strong (cyclic) plan with minimal cost. (We could also restrict the set of observations the human can perform, which will reduce the search space).

3.2 Creating new actions to avoid dead-ends

The idea is based on the work of [3], that explains why a deterministic planning problem $\mathcal{P}$ has no solution by generating the closest state from the initial state of $\mathcal{P}$, from which a planner can find a plan solution. In this work, a *fictitious action* is a new action that only changes the value of a single state variable (a literal).

We use the concept of *fictitious actions* to create new human actions. We do not allow the planner automatically choose what literal will be transformed into a *fictitious action*, since this will result in a very large number of actions for the planner to reason about. Instead, we first consider only the literals appearing in the preconditions of the actions and observations. A better idea is to compute the set of relevant literals to achieve the goal, $Rel(G)$, from a *causal graph* [4], that represents the relations between the literals and the actions. In order to do this, we had to define the concept of *causal graph* for contingent planning problems.

3.3 Planning with Bound Cost

The proposed approach can also be easily adapted to solve problems with no dead-end belief states but with a cost bound. In this case, the agent may prefer to rely on human help instead of come up with a solution that can be too costly. For instance, if a robot is taking too long to look for a key, it may decide to ask a human: - *Where is the key?*, even when there exists a strong (cyclic) plan but it is too costly.

3.4 Evaluation

To test our ideas, we have developed: (1) a parser to read a contingent planning problem written in PDDL, deriving all of its invariants, i.e. variables that never change its value, to reduce the size of the original problem; (2) the basic k-translation, compiling the original contingent problem to a FOND problem; (3) a modified FF-planner to be able to deal with the non-determinism and feed it the translated problem to solve it; (4) we implemented the contingent Causal Graph, to extract information about the relevant literals that will be considered as fictitious actions. The preliminary results show that problems with dead-ends or weak solutions are solved by our planner, that is able to compute strong plans considering the human help and the actions that modify the literals from the problem.

Currently we are studying the Causal Graph to see how we can obtain more information about the clauses that define dead-ends in a planning domain.

4. PHD OBJECTIVE

The objective of this PhD proposal is to develop a contingent planning agent that, on the presence of dead-end belief states, relies on human help to: (1) reduce agent uncertainty in problems with a dead-end belief state with no pure dead-ends but with no applicable actions (by asking the human to make an observation with a relaxation of its preconditions); (2) create new actions (which were not originally designed to the agent) in problems that have no strong (cyclic) solution, which corresponds to a modification of the environment (performed by the human) to allow the agent to continue its search and; (3) find a cheaper solution when the solution cost without external help is too high. Notice that, in our approach, the agent does not see a human as a decision maker, instead, the agent plans to achieve a goal deciding what observations or actions he can ask the human to perform, in order to complete a task. Another important aspect of our approach is to: (a) guarantee certain conditions for human help, e.g., human safety and competence; and (b) minimize human help action in the plan.

To solve this problem, we first translate the original problem into a FOND problem using a *modified k-translation* (to deal with the three types of dead-end belief states) that includes human actions and observations. Then we use an out-of-the-shelf FOND planner [2] [1] [5]

Acknowledgment

The student is supported by funding from the CAPES Foundation within the Ministry of Education, Brazil.

REFERENCES

[1] A. Albore, H. Palacios, and H. Geffner. A Translation Based Approach to Contingent Planning. In *IJCAI*, 2009.

[2] B. Bonet and H. Geffner. Belief tracking for planning with sensing: Width, Complexity and Approximations. *JAIR*, 50:923–970, 2014.

[3] M. Göbelbecker, T. Keller, P. Eyerich, M. Brenner, and B. Nebel. Coming up with good excuses: what to do when no plan can be found. In *ICAPS*, 2010.

[4] M. Helmert. The Fast Downward Planning System. *JAIR*, 26:191–246, 2006.

[5] C. J. Muise, V. Belle, and S. A. McIlraith. Computing contingent plans via fully observable non-deterministic planning. In *AAAI*, 2014.

[6] S. Rosenthal, J. Biswas, and M. Veloso. An effective personal mobile robot agent through symbiotic human-robot interaction. In *AAMAS*, 2010.

[7] S. Rosenthal and M. Veloso. Mobile robot planning to seek help with spatially-situated tasks. In *AAAI*, 2012.