A Framework For Human-Aware Robot Planning

Marcello CIRILLO, Lars KARLSSON and Alessandro SAFFIOTTI
AASS Mobile Robotics Lab, Örebro University, Sweden

Abstract. Robots that share their workspace with humans, like household or service robots, need to take into account the presence of humans when planning their actions. In this paper, we present a framework for human-aware planning in which we consider three kinds of human-robot interaction. We focus in particular on the core module of the framework, a human-aware planner that generates a sequence of actions for a robot, taking into account the status of the environment, the goals of the robot and the forecasted plan of the human. We present a first realization of this planner, together with two simple experiments that demonstrate the feasibility of our approach.

Keywords. Robot task planning, Human-aware planning, Human-robot interaction, Intelligent environments

1. Introduction

Until now, robots have been confined into special working cells under controlled conditions, and have been applied mostly to industrial automation. Now, the interest of the public for home robots is increasing, and people are looking at robots as a new mean to improve the quality of their everyday life. The aging of the population, for instance, could open a wide space for new robotic applications [10]. The robots could then become silent workers, precious butlers and, eventually, friendly helpers in our houses.

The presence of humans introduces other challenges besides how the robots should interact with them. It also has a profound influence on how the robots perform high level reasoning and especially plan their actions. Most AI planning systems use a model in which the world is in a particular state, the robot executes a specific action, and the world changes to another state, and so on. This state may be to some extent unpredictable, but it is the robot's choice of actions that determines what the next states can be. In other words, the robot is in control. However, humans are agents that act independently of the robot, and they are only partially observable. Thus, when planning, the robot needs to consider two different processes affecting the state, both its own actions and the actions of the human.

There has been a number of approaches to planning in partially observable and non-deterministic/stochastic environments for mobile robots, and as mentioned, the presence of human actors introduces partial observability. POMDPs [7], which can deal with sensing and uncertainty, have featured prominently among these, such as the robot Xavier [9]. However, these works tend to be focused on navigation tasks. Considering humans as a

Cite as: M. Cirillo, L. Karlsson, A. Saffiotti. A Framework For Human-Aware Robot Planning. *Proc. of the Scandinavian Conf on Artificial Intelligence* (SCAI). Stockholm, SE, 2008.

Figure 1. An example of human-aware robot planning.

source of uncertainty, one can point to some work on planning with external/exogenous events, which in principle can be caused by humans. An early example is the work of Blythe [4], which used Bayesian nets to compute the probability of success in the presence of external events. In the robotic field, some works have considered human-robot co-habitation. Often these works take a viewpoint which is different from the one adopted here, by focusing on aspects such as security (e.g., within MORPHA [6]) or acceptable motion (e.g., within COGNIRON [1]). The problem of task planning in the presence of humans is currently completely open, although some researchers have started exploring the issue [2,3].

In this paper, we present a framework for human-aware robot task planning in which the human(s) and the robot(s) both have their own goals, but the robots should prepare their own plan taking into account the presence of the humans. We classify the types of human robot interaction that we consider into three categories:

1. Both human and robot have goals, that are different. The robot should try to avoid interference with the human (implicit cooperation)
2. The human has a goal (or a set of goals) and the robot should help to accomplish it. The robot becomes a butler for the human
3. The robot has a goal to accomplish and needs some extra help. In this case, it is the robot that asks for help from the human

Figure 1 shows an example of the intended interaction in case 1.

The three kinds of interaction are ordered by increasing complexity and highlight the steps that are needed in order to bring the two worlds, of the humans and of the intelligent systems, to a closer cooperation. In this paper, we propose a general framework to address the problem of robot task planning and execution within the above categories. The full study and realization of this framework is an ambitious goal, and in this paper we offer a small but concrete first step by proposing a human-aware planner that deals with the first category under restrictive assumptions.

In the next section we present our framework, including an hypothetical example of its intended working. Section 3 presents the first version of our human-aware plan-

Cite as: M. Cirillo, L. Karlsson, A. Saffiotti. A Framework For Human-Aware Robot Planning. *Proc. of the Scandinavian Conf on Artificial Intelligence* (SCAI). Stockholm, SE, 2008.

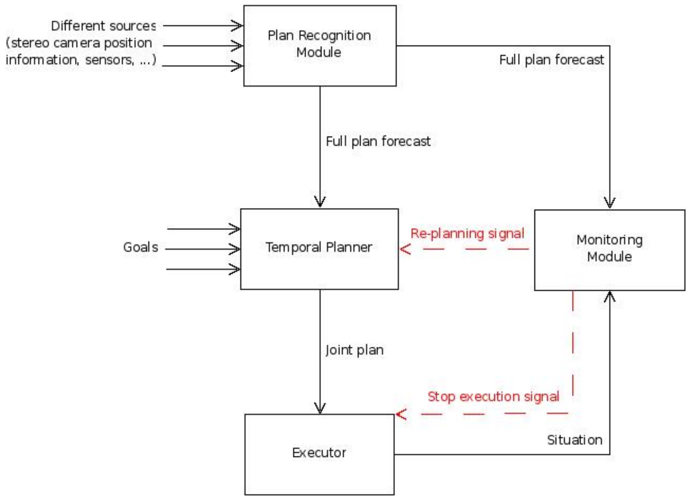

Figure 2. The proposed framework for human-aware robot planning.

ner, while in Section 4 we report two simple experiments that show the validity of this planner. Section 5 concludes.

2. The Overall Framework

To plan safely and effectively in the presence of human, the robot should be aware of the actions performed by the human and should be able to produce a forecast of the human's next actions. Taking into account the forecast, the robot can then make a plan of its own, avoiding any unwanted interference, assisting the human when needed, and, in the most complex case previously described, requesting assistance from the part of the human when this request does not cause discomfort to the inhabitant of the environment.

2.1. Architecture

This idea above is embodied in the architecture sketched in Figure 2. This schema does not describe an implemented system, but rather the proposed organization of functionalities needed to realize human-aware planning. The intent of this schema is to clarify the place and role, in the overall human-aware planning problem, of the planner described in the next section.

Plan Recognition Module The task of this module is to provide the robot with the needed awareness of the actions and intentions of the human(s), by recognizing their actions and forecasting their plans. This module receives as input information from different sources, and it produces, at each time t, a forecast of the possible plans of the human(s) in the environment. Every forecasted plan may be associated with its likelihood. As time passes, the set of plans and their likelihoods are updated as new actions are detected. Human plan recognition is a very active area of research (see, e.g., [5]). In our work, we plan to use existing state of the art systems for this module.

Temporal Planner The temporal planner is the main focus of the research presented in this work. It receives as inputs goals and the human's plan forecast. The output of this module is a temporal, probabilistic joint plan of both robot(s) and human(s). Since it is not certain which plan each human is performing, then the robot plans must be branched to take into account different possibilities. The plan is both temporal (absolute time and time dependencies between actions are taken into account) and probabilistic (actions may fail or lead to different outcomes). The full joint plan of both human(s) and robot(s) is then passed to the executor.

Executor The executor puts the robot plan into action. Information about the current situation is then sent to the monitoring module.

Monitoring Module The monitoring module provides a continuous control over the robot plan execution, ensuring that time constraints are respected and implementing some strong safety policies to protect the human inhabitants. In case inconsistencies are found, the monitor can stop the plan execution and request a new plan generation.

The importance of the monitoring capability lays in the fact that our main objective is to allow the cooperation between two worlds that are, for their nature, strongly different. The goal of synchronizing and keeping them consistent raises a number of open research issues. In every moment the robot should be aware of the overall situation, understanding what is the plan of the human and verifying that the actions of the robot do not interfere with the people present in the environment. The robot should verify as well that its plan is still consistent both with the human's needs and with the overall situation, that is, the robot's plan may have to be changed not only because of possible external events, but also because of the unpredicted actions of humans.

2.2. *Example*

We illustrate the intended working of the above framework by way of a hypothetical example. We assume a household pervaded by sensors and intelligent devices, which provide input to the plan recognition module. For instance, the information that the user is watching TV could be acquired by a stereocamera on the ceiling and then validated by a pressure sensor on the couch and by the fact that the TV is on.

Figure 3 graphically illustrates the example. The four lines at the bottom represent the predicted plans of the inhabitant generated by the plan recognition module. In the example, at time *t0* the user is observed sitting on the couch and the system identifies three possible patterns of actions: the human may want to read a book (**b**), watch television (**c**) or make a phone call (**d**). As time passes and new data are available, some plans that seemed possible are discarded and new possibilities are identified. In the example, at time *t1* the human is observed taking off his glasses, and the system identifies the new possible plan to take a rest (**a**). Then, at time *t2*, the user is observed turning off the phone, and the corresponding plan (**d**) is discarded as not consistent.

At time *t2* the temporal planner receives a goal, and it generates the probabilistic temporal robot plan represented by the upper part of Figure 3. This plan contains conditional branches: actions that are applicable if the human real plan is (**a**) may be not an adequate choice in case the plan is in fact (**b**) or (**c**). Each action of the robot plan (represented in the figure by the capital letter **A**) is applicable to a subset of the possible human plans. The robot will then start to execute the actions that correspond to the

Cite as: M. Cirillo, L. Karlsson, A. Saffiotti. A Framework For Human-Aware Robot Planning. *Proc. of the Scandinavian Conf on Artificial Intelligence* (SCAI). Stockholm, SE, 2008.

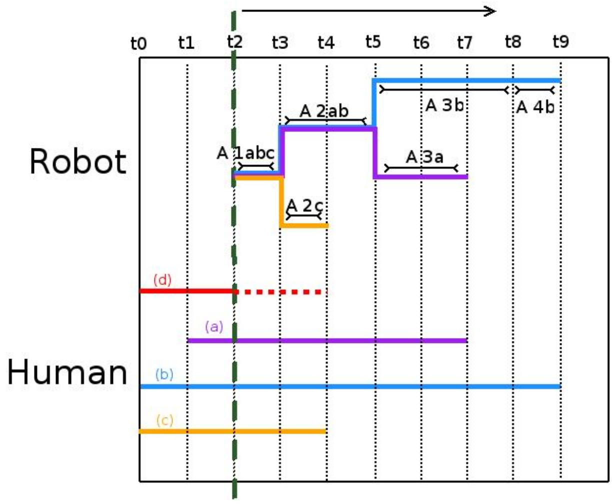

Figure 3. Graphical representation of the output of human-aware planning. The *x* axis represents time, quantized in slots. The lower part of the picture represents the possible human plans generated by the plan recognition module. The upper part represents a conditional robot plan generated by the human-aware planner. Branches in the robot plan specify courses of robot's actions corresponding to different human plans.

human's predicted plan with the highest likelihood. The monitoring module will ensure both the safety of the human and the soundness of the execution.

3. A Human-Aware Planner

In this paper, we focus on the human-aware planning module (that is, the "temporal planner" box in Figure 2), its design and implementation. As we said, human-aware planning involves a number of totally open research issues. Therefore, we adopt an incremental approach to the problem, starting our work with a very simplified scenario and, with each step of our research, relaxing some constraints to tackle more general cases. The complexity of the research problem requires simple, even unrealistic assumptions as starting point. For instance, in the first step we present in this paper, we assume the presence of a single human and of a single robot in the environment. Moreover, we assume that every action of the human ends in a deterministic way and we exclude the possibility for the human to change his intentions in the middle of an action: every activity that is started must be concluded without any interruption.

The input for the planner is a single plan for the human, that is, a linear sequence of actions split in discrete time slots (the expansion of the planner to handle multiple human plan with different levels of confidence is part of our future work). The planner produces a plan according to the goals of the robot, the state of the environment, the constraints imposed by the sequence of actions of the human and the cost of the actions that the robot can perform. From a more technical prospective, we use as a starting point PTLPlan [8], a probabilistic conditional planner that we extend to perform our first experiments on human-aware planning.

The extended planner works as follows. It takes as input:

1. A planning domain description specifying both the actions the robot can perform, and those of the human. Action are specified in terms of preconditions and

effects; the latter may be context dependent and/or stochastic, and may involve sensing (for the robot). In other words, the actions are of the type commonly found in POMDPs.

2. The robot's current belief state, including the topology of the environment and the estimated positions of the robot and the human.
3. A hypothesis about the plan of the human, as a sequence of actions.
4. A goal

In this first step, we make the following assumptions on the forecasted plan for the human:

- The forecast received by the planner is limited to one single plan.
- In the human plan there are no conditional branches.
- The forecast does not contain parallel actions.
- The duration of the actions is fixed.

Notice, however, that the actions of the human can have non-deterministic effects with associated probabilities.

In the original planning algorithm, the planner starts from the current belief state and then tries different actions, resulting in new belief states from which the algorithm can continue its search. If an action involves sensing and can result in different observations, then that action will result in several different belief states (one for each potential observation) and the planner will consider them all. The algorithm takes advantage of user-specified control rules (specified in temporal logic) to prune certain sequences of belief states. The algorithm continues until a plan (or policy) is found for which a goal state is reached with sufficient probability. That plan is conditional on the observations that the robot may do. For instance, it may contain branches like "if the door is observed to be open, then move to the bedroom".

The extended, human-aware, version of the algorithm simply applies the next action in the human's plan after the robot tries to apply its own action. Hence, each application of a robot action yields one or more new belief states which include the effects of both the robot's and the human's action, and a new human plan where the last applied action has been removed. Non-determinism of the human's actions is handled thanks to the ability of PTLPlan to consider actions with non-deterministic effects. The output of the planner is a joint plan, a list of human actions coupled with the ones that the robot performs at the same time.

4. An Illustrative Experiment

We have performed some simple experiments with the human-aware planner, in which we simulate the input from the plan recognition module using synthetic data and we analyze the output without actually executing it. The goal is to test the behaviour of the planner when human constraints are provided. The human plan is composed by a sequence of linear actions, quantized in time slots of 5 minutes each. Therefore, an action may take several time steps to complete (e.g., if the human cooks for 18 minutes, the system will split his action in 4 basic units). In this experiments, all actions are deterministic. The human's actions are passed to the planner as a list and then a joint plan is generated, as we will see in the following.

Time	Human action	Robot action
0	(MOVE KITCHEN)	(MOVE LIVINGROOM)
1	(COOK)	(STAY LIVINGROOM)
4	(EAT)	(STAY LIVINGROOM)
7	(WASHDISHES)	(STAY LIVINGROOM)
9	(MOVE LIVINGROOM)	(CLEAN LIVINGROOM)
10	(READ LIVINGROOM)	(MOVE BEDROOM)
11	(READ LIVINGROOM)	(STAY BEDROOM)
24	(WATCHTV)	(STAY BEDROOM)
42	(MOVE KITCHEN)	(STAY BEDROOM)
43	(COOK)	(STAY BEDROOM)
48	(EAT)	(STAY BEDROOM)
53	(WASHDISHES)	(CLEAN BEDROOM)
54	(MOVE BEDROOM)	(MOVE KITCHEN)
55	(REST)	(CLEAN KITCHEN)
65	(REST)	(CLEAN KITCHEN)

Time	Human action	Robot action
0	(MOVE KITCHEN)	(MOVE LIVINGROOM)
1	(COOK)	(STAY LIVINGROOM)
3	(EAT)	(STAY LIVINGROOM)
5	(WASHDISHES)	(STAY LIVINGROOM)
6	(OUTSIDE)	(STAY LIVINGROOM)
10	(OUTSIDE)	(STAY LIVINGROOM)
24	(MOVE KITCHEN)	(STAY LIVINGROOM)
25	(COOK)	(STAY LIVINGROOM)
34	(EAT)	(STAY LIVINGROOM)
38	(EAT)	(CLEAN LIVINGROOM)
40	(EAT)	(MOVE BEDROOM)
41	(EAT)	(CLEAN BEDROOM)
42	(MOVE LIVINGROOM)	(MOVE KITCHEN)
43	(WATCHTV)	(CLEAN KITCHEN)
53	(WATCHTV)	(CLEAN KITCHEN)

Table 1. Left: full joint plan of Johanna and the cleaning robot, during a typical week day. Right: full joint plan of Johanna and the cleaning robot, during a Sunday. Since actions take usually more than a single time step, only time steps when the human or the robot start a new action are displayed.

The scenario we consider involves one human, an elderly person that lives alone, and one robot, a floor cleaner that can perform three basic actions: stay idle in a room, move to another room, and clean. The apartment is composed by three rooms: the bedroom, the livingroom and the kitchen. The human, Johanna, can perform different actions that we assume can be recognized by the system: she can cook, eat, wash the dishes, watch TV, read (both in the bedroom and in the livingroom) or she can go out for a walk. Some of her actions affect the cleanliness of the apartment (e.g., cooking produces more dirt on the kitchen floor) and this information is known to the planner. To further simplify the scenario, we assume that the days of Johanna are driven by a quite strict routine, which means that the plan recognition module can safely guess the activities of Johanna by observing a few cues, e.g., the day of the week, the wake-up time, the clothes she wears, and so on.

The cleaner is not allowed to start its activity until Johanna is awake and she leaves her bedroom. After that time, it can start performing its job with some strict limitations: the robot can never clean a room where Johanna is present in order to avoid any disturbances, and it must wait in rooms where no human is present to avoid her to stumble into it by accident. Every action of the robot has a cost in terms of battery consumption. The most expensive action is cleaning. Moving from one room to the other is less expensive, while the cheapest action is to wait still in a room.

We consider here two sample runs corresponding to two typical days in Johanna's life. In the first one, a normal week day, she wakes up early in the morning, she cooks her breakfast and eat it. She washes the dishes and then move to the living room, to read and to watch TV. Then, at lunch time, she moves back to the kitchen, prepares and consumes her lunch, washes again the dishes and then goes to the bedroom to have a one-hour rest. The second run considers a typical Sunday. Johanna wakes up later than usual, she has a quick breakfast (cooking, eating and washing the dishes) and then she goes out for a walk. At about half past eleven, she comes back home, she spends more time than usual to prepare a Sunday meal and to consume it. Finally, she moves into the living room to watch TV.

For each run, the input to the planner is the guessed plan of Johanna, and the goal is to have the apartment clean at the end of the day. The corresponding outputs from the planner are shown in Tables 1. Each row in the tables lists the time step and the actions of both the human and the robot that start at that time. Time 0 is the time when Johanna first leaves her bedroom in the morning. As it can be seen, the robot adapts its plan to the behaviour of the human, performing its tasks without disturbing the inhabitant or threatening her safety, according to the above constraints. It remains inactive as much as possible, in order to avoid useless power consumption and then performing the necessary actions to leave the apartment clean at the end of the day.

5. Conclusions

In this paper we have presented a framework to cope with the issues that arise in human-aware planning. Considering the complexity of the problem, we focused our research on the core module of the framework, a planner that takes into account the actions of the human(s) in the environment. Although we made strong simplifying assumptions regarding the plan of the human in our first step, our experiments confirm the feasibility of our approach. The simulated robot performs its tasks without interfering with the human, considering her actions and plans.

We explored so far only the first kind of human-robot interaction defined in the introduction. Our future work will focus first on relaxing the above assumptions and then on increasing the complexity of the interaction level between human and robot. We also intend to introduce into the environment more agents, both robotic and human.

6. Acknowledgments

This work has been supported by CUGS (Swedish national computer science graduate school) and by the Swedish KK foundation.

References

[1] E. Akin Sisbot, A. Clodic, L.F. Marin, M. Fontmarty, L. Brèthes, and R. Alami. Implementing a human-aware robot system. In *Proc. of the ROMAN Symposium*, pages 727–732, 2006.

[2] R. Alami, A. Clodic, V. Montreuil, E.A. Sisbot, and R. Chatila. Task planning for human-robot interaction. In *sOc-EUSAI '05: Proc. of the 2005 joint conference on Smart objects and ambient intelligence*, pages 81–85, New York, NY, USA, 2005.

[3] R. Alami, A. Clodic, V. Montreuil, E.A. Sisbot, and R. Chatila. Toward human-aware robot task planning. In *AAAI Spring Symp 'To boldly go where no human-robot team has gone before'*, pages 39–46, 2006.

[4] J. Blythe. Planning with external events. In *UAI '94: Proc. of the Tenth Annual Conference on Uncertainty in Artificial Intelligence*, pages 94–101, Seattle, Washington,USA, 1994.

[5] B. Bouchard, S. Giroux, and A. Bouzouane. A smart home agent for plan recognition of cognitively-impaired patients. *Journal of Computers*, 1(5):53–62, 2006.

[6] B. Graf, M. Hans, and R.D. Schraft. Mobile robot assistants: issues for dependable operation in direct cooperation with humans. IEEE *Robotics and Automation Magazine*, 11(2):67–77, 2004.

[7] L.P. Kaebling, M.L. Littman, and A.R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1–2):99–134, 1998.

[8] L. Karlsson. Conditional progressive planning under uncertainty. In *Proc. of the 17th Int. Joint Conference on Artificial Intelligence (IJCAI)*, pages 431–438, 2001.

[9] S. Koenig and R. Simmons. Xavier: A robot navigation system based on partially observable markov decision process models. In Kortenkamp, Bonasso, and Murphy, editors, *Artificial intelligence based mobile robotics: case studies of successful robot systems*, pages 91–122. MIT Press, 1998.

[10] A. Tapus, M.J. Mataric, and B. Scassellati. The grand challenges in socially assistive robotics. IEEE *Robotics & Automation Magazine*, 14(1):35–42, Mar 2007.

Cite as: M. Cirillo, L. Karlsson, A. Saffiotti. A Framework For Human-Aware Robot Planning. *Proc. of the Scandinavian Conf on Artificial Intelligence* (SCAI). Stockholm, SE, 2008.