Künstliche Intelligenz manuscript No.
(will be inserted by the editor)

Planning in Inhabited Environments

Human-Aware Task Planning and Activity Recognition

Marcello Cirillo

Received: date / Accepted: date

Abstract Our work addresses issues related to the cohabitation of service robots and people in unstructured environments. We propose new planning techniques to empower robot means-end reasoning with the capability of taking into account human intentions and preferences. We also address the problem of human activity recognition in instrumented environments. We employ a constraint-based approach to realize a continuous inference process to attach a meaning to sensor traces as detected by sensors distributed in the environment.

1 The Dream of a Robot Butler

State-of-the-art commercial robots are finally finding their way into our homes (as the Roomba vacuum cleaners from iRobot) and into our workplaces (e.g., the Magic Shelf from KIVA systems), performing efficiently their tasks. However, we are still far from fulfilling the expectations raised by decades of science fiction. One of the main features that distinguished fictional robots and intelligent environments from the ones we can now deploy in our homes is the capability to closely cooperate with humans. This cooperation is both explicit and implicit, as science fiction intelligent systems can act when directly prompted by the users, but they can also support them by understanding what are their intentions and performing high level reasoning to achieve complex tasks to assist them.

Human-robot cooperation and interaction is a complex research problem, not limited to the identification of suitable communication interfaces. The presence of humans in the space where robots operate has also a profound influence on how the embodied agents should perform high level reasoning and plan their actions. We believe that the interaction and teamwork of robots and humans can be greatly improved if the robot can anticipate the forthcoming actions of the human. The robots should acquire knowledge about people's current and predicted future activities and use such knowledge to plan their actions accordingly.

In the Ph.D. thesis "Planning in Inhabited Environments: Human-Aware Task Planning and Activity Recognition"[1] [1], we propose techniques for human activity recognition in intelligent environments and human-aware robot task planning as means to achieve a better integration between people and service robots.

2 A Motivating Scenario

Malin lives alone in her small apartment. The apartment has been instrumented with a series of sensors and actuators which help her manage some of the physical and cognitive difficulties she has due to age. The home alerts her if she appears to be over-cooking her meals, can recognize when she is sleeping, eating and performing other usual activities at home, and can be easily set up to monitor and respond to the occurrence of specific patterns of behavior, like getting her a drink from the fridge when she watches TV. A robotic vacuum cleaner is also deployed in the apartment. It can plan its cleaning tasks accord-

M. Cirillo
AASS Research Centre, Örebro University, Sweden
Tel: +46-19-303694
E-mail: marcello.cirillo@aass.oru.se

[1] The PDF version of the thesis can be downloaded from `http://oru.diva-portal.org/smash/get/diva2:370049/FULLTEXT03`

ing to Malin's daily schedules, to avoid disturbing the elderly woman.

The scenario described above is just an example of how future robotic systems could improve the quality of our everyday life: the aging of the population in many countries, for instance, will open a wide space for new applications [10]. Cognitively impaired elderly people could live more safely in their homes with a system that monitors their actions and helps them in everyday tasks. Obviously, the same techniques that would help the elderly woman described in the example, could be easily applied to industrial environments: for instance, human operators in dangerous factories could minimize their presence in hazardous areas with the help of a robot assistant. Robots could then become effective workers, precious butlers and, eventually, friendly helpers in houses and factories.

Classical robot task planning [7] typically assumes that the robot is the only acting agent in the world. The main contribution of the Ph.D. thesis is a new approach to robot task planning in the presence of humans, that is, we address the problem of how robots should perform high level reasoning about their own future actions in the presence of people, avoiding interference and helping them in their tasks. We also address the complementary problem emphasized in the scenario above: a robot that aims at cooperating with people should have some knowledge about their current and possible future activities. Therefore, we also propose new techniques for human activity recognition. Our techniques are are both theoretically founded, and experimentally validated on real robots and devices.

3 Human-Aware Planning

In our work, we define human-aware planning, that is, our approach to robot means-end reasoning in the presence of humans. In general terms, human-aware planning can be applied to situations in which there is a controllable agent (the robot) whose actions we can plan, and an uncontrollable agent (the human) whose action we can only try to predict.

A simple example summarizing our approach can be seen in figure 1. A robotic vacuum cleaner has the goal to clean the apartment in which a human lives, and it should do so while not disturbing the human. The robot is provided with a prediction of what the person is going to do in the morning hours (at the bottom of the figure) and it plans accordingly its cleaning operations (the policy at the top of the figure) so as it is never cleaning a room while the human is using it. This is indeed a simple example, in which we have only one prediction for the human and the robot policy generated by the planner is simply a sequence of actions. However, this clarifies a core point of our approach: human actions are predicted, not planned, and they are used as input for the task planner of the robot.

To support human-aware planning, we identified a set of important functionalities that need to be incorporated in a planner:

1. Support for alternative hypotheses of the human plans, where a human plan — or human *agenda* — is a sequence of actions that the human might perform in the future;
2. Temporal duration of actions, both from the robot and from the human side;
3. Possibility to specify interaction constraints, that is, formulas that determine how the robot should or should not interact with the human;
4. Support for partial goal achievement; and
5. Support for observations on effects of human actions.

In the thesis [1], we present an implementation of a human-aware planner called HA-PTLplan, which incorporates all of the above functionalities [2]. HA-PTLplan takes a set of possible human agendas, the interaction constraints and a set of goals as input and generates a policy for the robot that is compliant with the inputs. The possible agendas, together with an associated probability distribution, are assumed to be estimated by a separate module. As one does not know in advance which of these agendas is the actual one, the planning problem has the property of partial observability. HA-PTLplan, which is an extension of PTLplan [5], generates policies which are conditional on observations relating to the human's actions.

We also extend the problem definition and the planner itself to deal with situations in which multiple humans are present in the environment, each with his/her own possible agendas. Obviously, considering multiple humans increases the complexity of the planning algorithm, but in our work we show how injecting domain knowledge into the planner can increase its efficiency, without compromising the quality of the solutions.

4 Activity Monitoring

An important requirement for human-robot cooperation and cohabitation is the capability of the robots to have a clear understanding of the state of the human and of his future plans. We argue that the monitoring of a human user is better achieved if data can be collected directly from the environment where he acts, by means

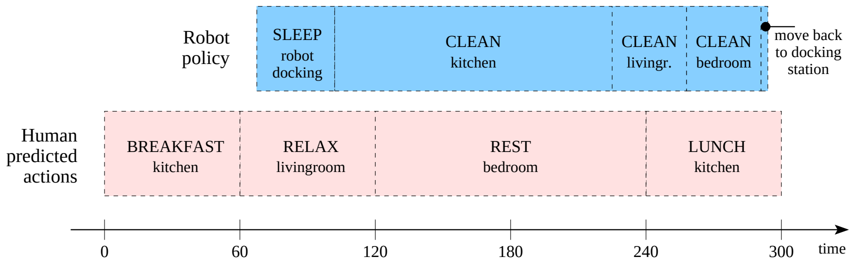

Fig. 1 A simple example of human-aware planning. A robotic vacuum cleaner is equipped with a planner that takes as input the predicted actions that the human user will perform in the following hours (bottom part). The output of the planner is the policy of the robot (top). The moving actions of the robot from one room to the other have been omitted for clarity reasons. Since the robot is required to end its tasks at the docking station, a final *move* action has been specified at the end of the policy.

of distributed, unobtrusive sensors. This research direction is not new, and intelligent environments are quickly evolving to incorporate sophisticated sensing, and even actuation capabilities [6,9,8].

In the thesis, we propose a model-driven approach to activity recognition, in which requirements expressed in the form of temporal constraints are used to specify how sensor readings correlate to particular states of a set of monitored entities in the environment. These entities can represent, for instance, various levels of abstraction of a human user's daily activities.

Our approach has been implemented in SAM, an Activity Management architecture[2] for service providing intelligent environments [3]. SAM continuously refines its knowledge on the state of affairs in the environment through an on-line inference process. This process concurrently performs abductive inference to provide the capability to recognize context from current and past sensor readings. Moreover, thanks to its underlying formalism and inference mechanisms, SAM also provides contextualized service execution, synthesizing proactive plans for actuators distributed in the environment.

SAM is built on top of the Multi-component Planning and Scheduling framework (OMPS) [4] and satisfies a number of important requirements stemming from realistic application settings:

– *Modularity:* it should be possible to add new sensors and actuators with minimal reconfiguration effort;
– *Long temporal horizons:* the system must be capable of recognizing patterns of human behavior in which events are separated by long temporal intervals;
– *On-line recognition and execution:* we require the system to be capable of recognizing activities as soon as they occur;
– *Multiple hypotheses tracking:* the system must be capable of modeling and tracking multiple hypotheses of human behavior, in order to support alternative and/or multiple explanations of sensor readings.

Among the above requirements, the ones that are central for human-aware planning are the last two. In particular, on-line recognition is essential for the execution of policies generated by a human-aware planner, as real time information is required to detect the activity the user is performing. Multiple hypotheses tracking in activity recognition is also a very important feature, as it lays the foundations for multiple hypotheses forecasting. An example of how SAM tracks multiple hypotheses on the user's activities is shown in figure 2. In this particular experiment, fully detailed in the thesis, a human user is performing daily activities in an instrumented environment. Here, identical sensor traces lead to the recognition of two different, yet not mutually exclusive activities (watching TV and having a snack). It could obviously also happen that identical sensor traces would entail mutually exclusive hypotheses. In both circumstances, SAM keeps track of each hypothesis: the continuous inference process will then check the consistency of each one of them as further evidence becomes available from new sensor data, dropping the hypotheses that are no longer supported.

5 Methodology

In the thesis, we validate our approaches to human-aware planning and activity recognition in three different ways. First, when relevant, we present analyses of the formal properties of our algorithms, e.g., complexity analyses of the algorithms.

Second, we use simulated scenarios to test the applicability of our approaches in different, realistic situations. The simulations prove to be very useful to highlight how our approaches would adapt to different environments and to test different heuristics to speed up the computation. Finally, we have performed several

[2] SAM stands for "SAM the Activity Manager".

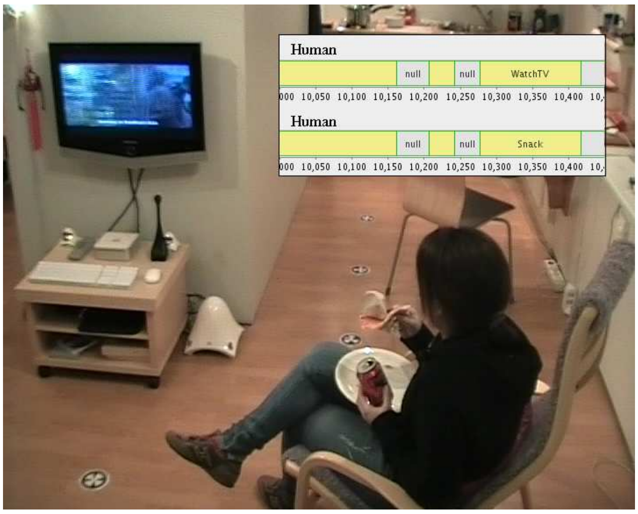

Fig. 2 The system supports two hypotheses on the human activity which are not mutually exclusive: watching TV and having a snack.

experiments in a real environment with people present. These test runs provide a practical demonstration of the applicability of our techniques to a real robotic system.

6 Conclusions

In the thesis, we discuss our approaches to two different, highly interconnected areas, activity recognition and robot task planning with people present. The solutions described for both problems are derived from the same family of methodologies, that is, symbolic reasoning and planning.

The main contributions of our work are the definition of human-aware planning as a new type of planning problem and the identification of the elements which need to be added to classical planning to cope with the presence of humans. Another important contribution of our work is the application of a constraint-based approach and a modular domain description language to realize a decisional framework that operates in a closed loop with physical sensing and actuation components in an intelligent environment.

Our thesis opens the way for interesting future research directions. A challenging but important objective would be to merge our approaches for activity recognition and human-aware planning. Another interesting direction for future investigations would be to integrate SAM with machine learning methods in order to provide on-line model training and adaptation. Finally, we believe that extending other search strategies and different planners to work with the formalization of a human-aware planning problem as described in this thesis could offer new possibilities for future improvements.

Acknowledgements

The work described in the thesis has been supported by CUGS (Swedish national computer science graduate school) and by the Swedish KK foundation. We are grateful to Lars Karlsson, Alessandro Saffiotti and Federico Pecora for their supervision and support. We also thank Dana Nau and Rachid Alami for the useful discussion and insightful advices.

References

1. Cirillo, M.: Planning in inhabited environments: Human-aware task planning and activity recognition. Ph.D. thesis, Örebro University, Sweden (2010)
2. Cirillo, M., Karlsson, L., Saffiotti, A.: Human-aware task planning: an application to mobile robots. ACM Transactions on Intelligent Systems and Technology **Special Issue on Applications of Automated Planning**(2), 15:115:26 (2010)
3. Cirillo, M., Pecora, F., Saffiotti, A.: Proactive assistance in peis-ecologies: a constraint-based approach. In: Handbook of Research on Ambient Intelligence and Smart Environments: Trends and Perspectives (in print). IGI Global (2011)
4. Fratini, S., Pecora, F., Cesta, A.: Unifying planning and scheduling as timelines in a component-based perspective. Archives of Control Sciences **18**(2), 231–271 (2008)
5. Karlsson, L.: Conditional progressive planning under uncertainty. In: Proc. of IJCAI (2001)
6. Lee, J.H., Hashimoto, H.: Intelligent space - concept and contents. Advanced Robotics **16**(3), 265–280 (2002)
7. Nau, D., Ghallab, M., Traverso, P.: Automated Planning: Theory & Practice. Morgan Kaufmann Publishers Inc. (2004)
8. Saffiotti, A., Broxvall, M., Gritti, M., LeBlanc, K., Lundh, R., Rashid, J., Seo, B.S., Cho, Y.J.: The PEIS-Ecology project: Vision and results. In: Proc. of IROS (2008)
9. Sgorbissa, A., Zaccaria, R.: The artificial ecosystem: a distributed approach to service robotics. In: Proc. of ICRA (2004)
10. Tapus, A., Mataric, M.J., Scassellati, B.: The grand challenges in socially assistive robotics. *IEEE* Robotics & Automation Magazine **14**(1), 35–42 (2007)

Marcello Cirillo received his M.Sc. degree in Computer Science Engineering from Politecnico of Milano, Italy, in 2005. In 2010, he successfully defended his Ph.D. thesis at the Center for Applied Autonomous Sensor Systems (AASS) in Örebro, Sweden. His research interests include robot task planning with people present and human activity recognition.