

An Integrated Planning and Learning Framework for Human-Robot Interaction

Alexandra Kirsch and Thibault Kruse and Lorenz Mösenlechner

Intelligent Autonomous Systems Group
Department of Informatics
Technische Universität München
Boltzmannstr. 3, D-85748 Garching
{kirsch,kruset,moesenle}@cs.tum.edu

Abstract

Assistive robot systems require a robot to interact closely with people and to perform joint human-robot tasks. Interaction with humans comes with additional challenges to those in other real-world scenarios. Robot plans must be especially flexible and take into account human abilities and preferences. For providing this level of flexibility, we propose a framework for transformational reactive planning that includes the capability to learn models of the human during plan execution. We show how this framework can be extended to the special requirements of human-robot interaction.

Motivation

Planning for autonomous robots is moving more and more into real-world dynamic environments. Especially challenging environments are those where the robot interacts with humans, for example in assistive applications. In this paper, we motivate these application domains as an interesting research for planning technology and present a transformational planning and learning framework to tackle the specific planning problems in human-robot interaction.

People are an extremely dynamic factor in the environment, not just because of their movements and manipulations in the world. Humans can change their mind on their envisaged goal, they can interrupt a task without the robot realizing the reason and they execute their tasks in some feasible, but non-optimal way. Therefore, robots need to form realistic expectations about human behavior and have to react flexibly to unexpected events.

In such a world, it is extremely important for a robot to have a plan about its own actions and that of the human partner. Only with a human-centered plan of the whole task can the robot show behavior that is understandable and can be expected by the human. Equally important is the flexible execution of such a plan considering social rules and personal preferences. We have identified two main issues when planning for human-robot interaction:

Models of human abilities and intentions. For a robot to cooperate with a person, we have to develop novel representation and inference mechanisms for models about a human's abilities and intentions, as well as a robot's own skills. While some general models about human

behavior like social rules might be provided as constraint rules to the planner, individual preferences and habits must be observed and learned by the robot during the interaction with the person.

Planning techniques to produce legible behavior. For assistive tasks, a robot must cooperate closely with a person. This means that it must represent the person's part in the activity in its own plan. When generating joint plans, the robot must consider the person's preferences and abilities. Even with a joint plan humans cannot be ordered or expected to behave in ways predefined in detailed plans. Daily human behavior may at any time be observed to be exceptional such as for humans changing their minds or being diverted. Using its models, a robot must distinguish unusual situations from normal ones. In the latter case, it must reactively adapt its behavior or try to clarify the situation.

To meet these challenges, we propose a framework consisting of TRANER (TRAnsformational planNER) [Müller, Kirsch, & Beetz, 2007] — a transformational planning system, which has optimized reactive plans in the domain of an autonomous household robot — and RoLL (Robot Learning Language) [Kirsch, 2008] — a language extension for continual learning and adaptation of models during program execution. TRANER has optimized plans in the domain of an autonomous household robot. Provided with transformation rules for joint human-robot plans that adapt the task assignment, TRANER makes it possible to adapt plans to human preferences. These preferences will be represented in the form of learned models. The Robot Learning Language (RoLL) [Kirsch, 2008] provides mechanisms to update models continually with the experiences made during plan execution. The combination of these two systems is an extremely flexible planning and plan execution framework, which allows the integration and development of special techniques needed for human-robot interaction. The combination of these two systems is an extremely flexible planning and plan execution framework, which we will extend for human-robot joint activities.

In the following we present a comprehensive scenario and point out the research challenges with respect to the two main research topics we have identified. Then we introduce our transformational planning and learning framework and




Figure 1: Assisting an elderly or handicapped person in loading the dishwasher as an interesting scenario for joint human-robot planning.

how it can be extended for HRI applications. The paper ends with a discussion on related work and a conclusion.

Scenario

Assistive technology is becoming more and more important in aging societies. We propose a robot helping an elderly person in the household as a challenging demonstration and research scenario. Specifically, we illustrate a situation where the robot helps to load the dishwasher as shown in Figure 1. We assume that the robot knows how to load a dishwasher and that it has general knowledge about the abilities of elderly people, but doesn't yet know the person it is working with.

First of all, the robot has to generate a joint human-robot plan considering the abilities and needs of elderly people. The robot might propose a plan in which the robot brings the dirty dishes to the dishwasher and the person loads it. The robot might propose this plan explicitly to the person and then they start executing it.

One elderly user might have a severe walking impairment and would prefer to sit down while putting the dishes into the dishwasher. Such a user might also want the robot to hand him objects directly (instead of putting them somewhere to be picked up). It would also be important for this user that the robot knows in which order the person likes to fill the dishwasher.

Another person might be able to walk, but cannot carry heavy objects. This person might be bored by the robot bringing the dishes when the person is well able to get some of them herself. Here a better way of collaboration would be the robot and the human bringing dishes to the dishwasher. As current robots are hardly capable of putting things into a dishwasher, the robot would better leave the dishes on the worktop for the person to pick them up.

To meet these personal preferences, the robot must remember wishes that the person expresses and observe the person's reactions when executing the joint plan, as well as possible failures. It should use them for learning models such as the person's ability to walk and carry objects or the time a person is willing to wait for the robot to do something.

With this new knowledge, the robot must adapt its plans with respect to the task assignment to each partner, the order of the tasks, the mode in which tasks are executed (e.g. handing over objects versus leaving them for the other to

be picked up), the closeness of the interaction (some people might prefer to avoid close contact with the robot), etc. The plans must optimize the joint human-robot activity with respect to legibility and meeting human preferences. Just like a human carer, the robot should also take care not to take over activities that the person can do on her/his own.

Even when a plan is well-adapted to the needs of a person, it can happen that the person doesn't do what the robot expects. It could happen that a person doesn't keep to the proposed plan while obviously still pursuing the same goal, e.g. when the person gets dishes himself instead of waiting for the robot to bring them. Or a person could abandon the plan temporarily, for example when answering the phone. There is also the possibility that a person makes a mistake like dropping a dish. Using its learned knowledge of the person's behavior, a robot should be able to recognize such unusual situations, classify them and react accordingly. Even the reaction of the robot might depend on human preferences — some people might prefer the robot to take over their task while they are distracted, others would like the robot to wait for their return.

Models for Human-Robot Interaction. The scenario gives examples of models that the robot needs to have of its individual user: the person's abilities and preferences for gripping something, the person's habits in which order to place objects in the dishwasher, the mobility of the person and implied constraints for the joint plan. Beside these individual abilities and preferences, the robot also needs knowledge about general human preferences, for example the way to approach a human or general expectations of people like committing to a task and performing actions in a certain order. In short, the robot needs representations of a person's abilities and skills, personal preferences and social rules.

These models fall in two classes: 1) models that should be derived from social and psychological studies and can be provided by hand-coded decision rules or constraints on plans (see for example [Koay *et al.*, 2007]), and 2) individual preferences and abilities of a user that can change over time and should be acquired and updated constantly by the robot in the interaction with the human. While both aspects are important for successful human-robot cooperation, the framework we propose for human-robot planning is especially apt for the second class of models.

Adaptive Planning. As illustrated in the scenario, the robot needs a plan that represents the human as well as the robot. Such a plan cannot be generated or even evaluated without knowledge of the personal preferences and abilities of a user. The adaptation to the user is necessary on several layers of abstraction: on a high-level planning level when assigning tasks to the partners, on an intermediate level when deciding to hand over an object or placing it somewhere, on a lower level when deciding how to approach the person.

For plan execution in real-world scenarios, failure monitoring, detection and repair are essential. When collaborating with a person, the robot must also monitor the failures of the person or other unexpected events and act appropriately. The robot must use its models to infer if its help is necessary or if the person can handle the failure on his/her own.

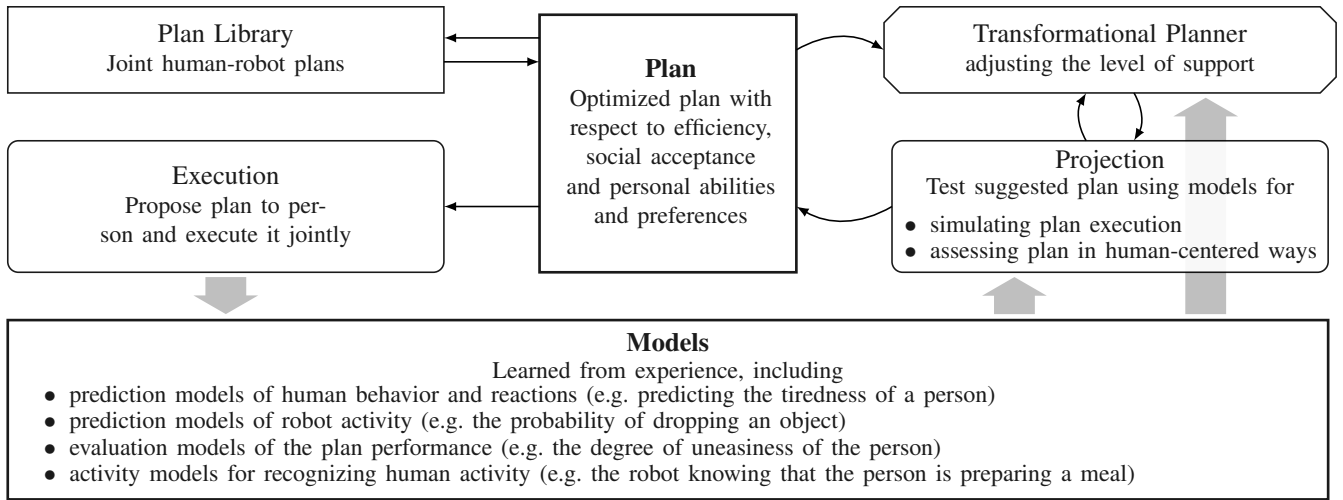


Figure 2: Plan-based adaptive control for human-robot interaction. The upper part shows the components provided by the transformational reactive planning framework of TRANER, the lower part describes the learning mechanism of RoLL, which is used to provide the models.

The Planning and Learning Framework

The framework we use for adaptive transformational planning is shown in Figure 2. It makes use of two existing systems: TRANER for transformational planning and RoLL for continuous adaptation through learning. We first explain the overall framework as we use it for HRI and then present the two subsystems and necessary adaptations of the current systems to a human-robot context.

If a robot is to achieve a goal — possibly together with a human — it chooses a plan from its plan library that can achieve this goal in the given situation. The current plan library of [Müller, 2008] includes activities for setting the table, clearing the table and preparing pasta, as well as all lower-level plans that are necessary to fulfill these activities like carrying objects and navigating.

Before using the plan, its usefulness in the present situation is assessed by predicting the outcomes of the plan. In a human-robot setting this also includes the person’s reaction and acceptance of the behavior. If the plan shows potential for improvement, it is transformed by generic plan transformation rules.

Once a plan has been transformed sufficiently to be expected to work well in the current situation, the robot executes the plan. During execution, it makes new experiences about the person’s behavior and reactions and also gets data about its own actions. With these experiences the robot updates the models in its program, which it needs for reliable plan execution as well as for plan generation.

The strength of the system doesn’t lie in individual planning techniques or learning algorithms. It rather provides an integrated framework that allows the robot to perform better over time by adapting its behavior in the form of plans and its knowledge in the form of models. These adaptations happen while the robot is active and thus enables a continuous improvement.

TRANER— Transformational Planner

TRANER [Müller, 2008] is a transformational planning system based on the Reactive Plan Language (RPL) [McDermott, 1993]. In RPL, robots do not only execute their programs but also reason about and modify them. TRANER realizes planning through a generate-test cycle in which plan revision rules propose alternative plans and new plans are simulated in order to test and evaluate them. The unique features of TRANER are that it can realize very general and abstract plan revisions such as “stack objects before carrying them instead of handling them one by one” and that it successfully operates on plans in a way that they generate reliable, flexible, and efficient robot behavior.

TRANER plans contain declarative specifications for monitoring the robot’s behavior, signaling failures, catching failures, and recovering from them.

The strength of the approach has been demonstrated in household tasks in a realistic physical simulation. The gain in efficiency of a transformed plan was up to 45% compared to the original plan. The plans monitored eight kind of failures and the robot always recognized if a failure occurred. 86 % of the failures could be recovered from, otherwise the robot could at least explain the type of failure.

For using TRANER in a human-robot setting, we need to add 1) a representation in the plan language to include human activities in the plan, 2) additional plan transformation rules and 3) a more general mechanism for evaluating plans.

1) We have complemented the plans with an additional parameter assigning a plan either to be achieved by the robot or the human. If the plan is assigned to the human, its execution consists in waiting for the person to perform it. With this mechanism, the TRANER failure mechanism can also be used to observe the human and detect unusual behavior.

2) Current transformation rules include changing the order of plan steps, combining plan steps and adding new intermediate steps. For human-robot interaction we need addi-

tional rules for redefining the assignment of actions to partners. This might also make it necessary to include additional plan steps such as hand over activities.

3) The evaluation of joint human-robot plans faces two difficulties compared with a non-HRI setting: Since there are no reliable simulations of humans, we cannot test the plans in the simulation, and the objective functions to decide if a plan is appropriate has to take into account human preferences. For both problems, we need good models of humans. For replacing the simulation, we will use the more general concept of plan projection, where the expected outcome of the plan is estimated with learned models (e.g. the exhaustion of the person after an activity, the duration of plan steps). For assessing the quality of a plan, we also need models of human preferences (e.g. the distance the robot kept during plan execution, the legibility of the behavior).

RoLL— Robot Learning Language

The Robot Learning Language (RoLL) [Kirsch, 2008] makes learning tasks executable in the control program. It allows for the specification of complete learning processes including the acquisition of experience, the execution of learning algorithms and the integration of learning results into the program. RoLL's design allows the inclusion of arbitrary experience-based learning algorithms.

Learning problems solved using RoLL include time prediction models of activities and decision functions for parameters of actions. However, the power of RoLL cannot be demonstrated by a specific learning problem, but lies in its versatility to use arbitrary learning mechanisms and to improve the robot's knowledge and behavior over time.

RoLL can be used directly for learning models of humans. One big challenge for learning predictive models — such as the time a person needs for an activity, the probability of needing help, the chance of succeeding in a task — is the reliable perception and classification of human activities [Tenorth & Beetz, 2008]. For more advanced models — like predicting the tiredness of a person, the acceptance of a plan, the space a person will cover during an action — the question of how to represent and learn those models with few experiences will have to be solved. RoLL offers the unique possibility to test and develop new methods for learning human models. It also facilitates the use of meta-learning, which might enable the robot to learn adequate state space representations.

Related Work

Alami *et al.* (2006) develop the Human-Aware Task Planner (HATP), which generates plans using additional constraints of social interactions. To allow legible movements of the robot Sisbot (2008) presents a method for human-aware motion planning including navigation and manipulation tasks where the robot respects social rules. A reactive planning approach for robots in human environments is described by Shiraishi and Anzai [Shiraishi & Anzai, 1996]. Here, the robot receives a command and has to generate a plan quickly in order to answer if it is able to fulfill the plan. During the execution phase the plan is transformed to perform better. The described approach only works on very abstract

plans with unparameterized actions and the interaction with humans is restricted to receiving commands. The human-aware planning framework of Cirillo, Karlsson, & Saffiotti (2008) uses predictions of the human plan to coordinate human and robot activities, but don't support direct human-robot interaction. In the domain of assistive technology, Pollack *et al.* (2003) describe a flexible reminder system for people with cognitive impairment using temporal planning techniques.

Rogers, Peng, & Zein-Sabatto (2005) report on models about human-robot interaction. The robot keeps track of its interaction state, which ranges from solitude, i.e. no interaction to active engagement with a human partner. In the scenario, the interaction consists of commands given by a person to the robot instead of performing a collaborative task. Representing spatial knowledge has been examined in the domain of human-robot communication [Moratz *et al.*, 2003; Levelt, 1996]. For natural language understanding, a common model of spatial relations is necessary. Similar to spatial knowledge is the idea of the robot taking the human's perspective and disambiguating commands [Trafton *et al.*, 2005].

Conclusion

When working in close interaction with a human, robots must plan both for the human and themselves. In this paper, we have pointed out specific challenges for plan-based control in the domain of human-robot collaboration, with special attention to assistive scenarios.

We have put those challenges in two categories that are closely related to one another: 1) the representation, acquisition and use of models of humans as well as their constant maintenance and update; and 2) requirements for plan-based robot control in joint human-robot tasks.

We propose an existing framework that combines transformations of reactive plans with experience-based learning, which produces flexible and adaptive plans. We have started to extend this framework to be applied in human-robot assistive scenarios.

References

- Alami, R.; Chatila, R.; Clodic, A.; Fleury, S.; Herrb, M.; Montreuil, V.; and Sisbot, E. A. 2006. Towards human-aware cognitive robots. In *AAAI-06, Stanford Spring Symposium*.
- Cirillo, M.; Karlsson, L.; and Saffiotti, A. 2008. A framework for human-aware robot planning. In *Proc. of the Scandinavian Conf on Artificial Intelligence (SCAI)*.
- Kirsch, A. 2008. *Integration of Programming and Learning in a Control Language for Autonomous Robots Performing Everyday Activities*. Ph.D. Dissertation, Technische Universität München.
- Koay, K.; Sisbot, E.; Syrdal, D.; Walters, M.; Dautenhahn, K.; and Alami, R. 2007. Exploratory studies of a robot approaching a person in the context of handing over an object. In *Proc. AAAI - Spring Symposium 2007: Multidisciplinary Collaboration for Socially Assistive Robotics*.

- Levelt, W. J. 1996. Perspective taking and ellipsis in spatial descriptions. In Bloom, P.; Peterson, M. A.; Nadel, L.; and Garrett, M. F., eds., *Language and Space*. MIT Press. 77–109.
- McDermott, D. 1993. A reactive plan language. Technical report, Yale University, Computer Science Dept.
- Moratz, R.; Tenbrink, T.; Bateman, J.; and Fischer, K. 2003. Spatial knowledge representation for human-robot interaction. In et al., C. F., ed., *Spatial Cognition III*, volume 2685 of *LNAI*. Springer-Verlag. 263–286.
- Müller, A.; Kirsch, A.; and Beetz, M. 2007. Transformational planning for everyday activity. In *Proceedings of the 17th International Conference on Automated Planning and Scheduling (ICAPS'07)*, 248–255.
- Müller, A. 2008. *Transformational Planning for Autonomous Household Robots using Libraries of Robust and Flexible Plans*. Ph.D. Dissertation, Technische Universität München.
- Pollack, M. E.; Brown, L.; Colbry, D.; McCarthy, C. E.; Orosz, C.; Peintner, B.; Ramakrishnan, S.; and Tsamardinos, I. 2003. Autominder: An intelligent cognitive orthotic system for people with memory impairment. *Robotics and Autonomous Systems* 44:273–282.
- Rogers, T. E.; Peng, J.; and Zein-Sabatto, S. 2005. Modeling human-robot interaction for intelligent mobile robots. In *IEEE International Workshop on Robots and Human Interactive Communications*.
- Shiraishi, Y., and Anzai, Y. 1996. Task planning based on human-robot interaction for autonomous mobile robots. In *IEEE International Workshop on Robot and Human Communication*.
- Sisbot, E. A. 2008. *Towards Human-Aware Robot Motion*. Ph.D. Dissertation, Université Paul Sabatier, Toulouse.
- Tenorth, M., and Beetz, M. 2008. Towards practical and grounded knowledge representation systems for autonomous household robots. In *Proceedings of the 1st International Workshop on Cognition for Technical Systems, München, Germany, 6-8 October*.
- Trafton, J. G.; Schultz, A. C.; Bugajska, M.; and Mintz, F. 2005. Perspective-taking with robots: Experiments and models. In *IEEE International Workshop on Robots and Human Interactive Communications*.