

# Learning Ability Models for Human-Robot Collaboration

Alexandra Kirsch and Fan Cheng

**Abstract**—Our vision is a pro-active robot that assists elderly or disabled people in everyday activities. Such a robot needs knowledge in the form of prediction models about a person’s abilities, preferences and expectations in order to decide on the best way to assist. We are interested in learning such models from observation. We report on a first approach to learn ability models for manipulation tasks and identify some general challenges for the acquisition of human models.

## I. PLAN-BASED CONTROL FOR HUMAN-ROBOT COLLABORATION

We are interested in planning and plan execution mechanisms that allow a robot to actively participate in joint tasks with a human partner, especially in assistive scenarios. We assume that in such a task there will only be limited explicit communication similar to the way when humans participate in a shared task.

One crucial factor in the development of such joint planning abilities is knowledge about the capabilities and preferences of the partners involved. When helping a person, a robot should take over those tasks that are difficult to perform for the person, but feasible for the robot. For example, an elderly person with a walking impairment can be supported by bringing items she needs to prepare a meal. But the robot should not attempt to cut the ingredients: first, current robots are not capable of doing such sophisticated manipulation tasks reliably, so the probability of failing would be very high and second, the robot should leave work to do for the person to keep her active and healthy and not to intrude into her private life more than necessary.

There is a wide variety of models that are interesting for a robot to interact with a human. Models describing the workspace of robots and humans [4] can be used to coordinate actions in joint workspaces or to choose high-level actions that avoid spatial conflicts. Another approach is to model criteria of social comfort into planning algorithms [1]. In our work, we try to describe capabilities on an action level, for example if a specific manipulation task would succeed in a given situation. We are interested in criteria such as success probability, effort for a person, efficiency and social acceptability (for example it might not be appropriate for a robot to touch food).

Because capabilities and preferences vary strongly among individuals, we would like a robot to learn models about itself and its human partner from experience. In the following we present our observations from a first attempt to learn the capabilities of agents to pick up and put down objects. We report on the challenges that we found in this work and summarize them in a general way to identify problems for similar learning problems.

## II. LEARNING ABILITY MODELS

For our research on joint human-robot activities we use a physical simulation of two agents (that are both displayed as robots): one is acting as an autonomous robot, the other one is controlled by a human via the keyboard. A person can move such an agent freely in the world and give commands for gripping and putting down objects. The gripping and put down actions are executed autonomously based on heuristics. These manipulation actions are implemented in different ways for the autonomous robot and the human-controlled one, so that the capabilities are not identical.

In a user study [2] we acquired data from nine subjects who had the task to set and clear the table in two simulated kitchen environments. In total, we observed about 60 gripping and put-down tasks respectively from the execution of complete plans for each participant. For learning capability models, we assumed that all participants were equally skilled in the manipulation tasks, because those were executed autonomously. This is not completely true, because the success of the task also depends on the position where the agent is standing while gripping. But we doubted that any learning algorithm would succeed with the small number of samples we had for each participant and with this simplification we had around 700 examples in total (This includes gripping and put-down tasks that were performed in incomplete runs that were aborted for some reason. This data was not used for evaluation in the user study, but can be used for our purposes here). Beside the data from the user study, we also collected analogous data for an autonomous robot.

As a first approach, we tried to learn prediction models of when such tasks succeed or fail by using decision trees (using the Weka J48 algorithm), for example this function for putting down an object:  $object-goal-position \times object-type \rightarrow success/failure$ . The object positions were given relative to the piece of furniture they were standing on or had to be put on, which was general enough given the samples from predefined scenarios from our user study. The result of these learning attempts was not surprising, but still disappointing: the decision tree judged that both the robot and the human-controlled agent will succeed in all cases. Looking at the learning experience, the reason for this result was very obvious: only about 4% of the gripping tasks from the user study were identified as failures and 7% of put-down task. The rates for the autonomous robot are similar.

Beside the low failure rate overall, there was no obvious structure when manipulation tasks succeed and fail. One hypothesis was that the wall behind the worktop would be




Fig. 1. Possible results of successful put down actions.

a crucial factor. Since we used only one kind of kitchen furniture, the way we specified positions should have accounted for the closeness to the wall. But it seems there were not enough samples in the training set to identify this in the learning process.

We got slightly better results when including the distance of the agent to the robot as an input variable. This doesn't give a model about the capability of gripping the object, because the agent might move, but it could be interesting as a comfort model to predict if a person would have to move and if so this might be an indication of low comfort and a good opportunity for assistance.

Another drawback of our approach is the way we define success and failure of gripping. We used a local decision after each grip and put down task to decide if the object was in the agent's gripper or the object was put near enough the goal position. However, when observing longer tasks, it becomes obvious that this is not the only source of failures that can occur. Figure 1 shows the end position of a scenario in which two plates, cups and knives had to be brought from the table to the worktop. In both runs all the put down actions were rated as successful. In the left picture, this is true, but in the right picture, the agent involuntarily changed the positions of other objects and the result was anything but satisfactory.

This leads to the question of how to model the state space of such activities. Which other objects apart from the one to be manipulated might be affected by the task and which parameters of these objects are important for the action execution? We are not aware of learning algorithms dealing with a varying number of input parameters. Therefore, the existence of a varying number of objects would have to be coded into a fixed number of parameters.

We continue our work on how to learn prediction models in several ways:

- We are currently exploring situations where two objects are affected by a manipulation task (one is an obstacle, the other one has to be moved) and try to work towards general representations of the state space that is suitable for learning and could be generalized to multi-object cases. One problem here is that the state space will be bigger, which means that more experience is needed for learning. As long as we can use a simulation to generate examples, this is fine, but acquiring enough information in real-world environments where the observations can be expected to be more noisy, this will be a serious problem.
- The Boolean outcome of success or failure seems to be inappropriate for describing the outcome of actions. A continuous value might be more appropriate both from the viewpoint of how to learn it and from the decision-

making process of the robot. The robot can then decide at runtime, which level of quality is needed for a task in a certain situation and can trade off different possibilities for distributing the tasks of joint plans.

### III. CHALLENGES AND FURTHER RESEARCH

From our first approach to learning predictions models, we identify some general problems of representing and learning models about human and robot activities. By tackling these issues we hope to develop general mechanisms for action modeling that can equally be applied to other predictions like social acceptability, the effort needed for a task, or personal comfort and for all kinds of actions on different levels of the action hierarchy.

These are the main challenges that we have identified:

- The structure of our models is currently a simple function. However, we think that the temporal component of actions (e.g. modeled with hybrid automata) and their effects in the world must also be represented. In our example we realized that the action as such is not the crucial question, but the situation in the world plays an important role. And it is not sufficient to find a clever representation for one specific action: there are too many to be modeled individually. We rather need general representation methods and possibly ways to acquire them automatically by the robot using and observing those actions.
- We will have to consider other forms of adaptation than currently used learning algorithms. Standard machine learning methods have two drawbacks for learning the models we are interested in: 1) they need a large number of samples, which might be difficult to obtain from observing humans, 2) they work on functional representations, which is not directly compatible with the temporal representations we would like to use.
- In our attempt to learn models of actions, we relied only on learning. However, there are explicit reasoning methods, e.g. for spatial reasoning [3], that could help to make predictions. When reasoning explicitly about the geometry in the world, the individual differences in capabilities and personal preferences of humans are hard to include. Besides, there are the same questions of how to model the relevant part of the world to make geometric reasoning feasible. We think that a combination of explicit reasoning and learned predictions would be a promising way to model human and robot actions.

### REFERENCES

- [1] Samir Alili, Rachid Alami, and Vincent Montreuil. A task planner for an autonomous social robot. In *Distributed Autonomous Robotic Systems 8*. Springer Verlag, 2009.
- [2] Alexandra Kirsch. Be a robot — a study on everyday activities performed in real and virtual worlds. Technical Report TUM-I1006, Technische Universität München, 2010.
- [3] Emrah Akin Sisbot, Luis F. Marin-Urias, Rachid Alami, and Thierry Simeon. A human aware mobile robot motion planner. *IEEE Transactions on Robotics*, 23:874–883, 2007.
- [4] F. Zacharias, Ch. Borst, and G. Hirzinger. Capturing robot workspace structure: representing robot capabilities. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3229–3236, 2007.