



A Testbed for Adaptive Human-Robot Collaboration

Alexandra Kirsch, Yuxiang Chen

► To cite this version:

Alexandra Kirsch, Yuxiang Chen. A Testbed for Adaptive Human-Robot Collaboration. 33rd Annual German Conference on Artificial Intelligence (KI 2010), 2010, Karlsruhe, Germany. hal-01405758

HAL Id: hal-01405758

<https://hal.science/hal-01405758>

Submitted on 30 Nov 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

This work was published in *KI 2010: Advances in Artificial Intelligence 33rd Annual German Conference on AI, Karlsruhe, Germany, September 21-24, 2010. Proceedings*, Dillmann, R.; Beyerer, J.; Hanebeck, U.D.; Schultz, T. (Eds.), Lecture Notes in Computer Science, Vol. 6359

The original publication is available at www.springerlink.com:
<http://www.springerlink.com/content/ex3j621162764366/fulltext.pdf>

A Testbed for Adaptive Human-Robot Collaboration

Alexandra Kirsch, Yuxiang Chen

Technische Universität München*

Abstract. This paper presents a novel method for developing and evaluating intelligent robot behavior for joint human-robot activities. We extended a physical simulation of an autonomous robot to interact with a second, human-controlled agent as in a computer game. We have conducted a user study to demonstrate the viability of the approach for adaptive human-aware planning for collaborative everyday activities. The paper presents the details of our simulation and its control for human subjects as well as results of the user study.

1 Motivation

Robots that are to interact with humans need to be highly adaptive towards their environment. They have to adapt their behavior to different human individuals, have to respect social rules and respond to human expectations, react to unexpected human actions, and show a high degree of safety, even for unanticipated events. In our research we aim to develop adaptive, model-based planning and plan-execution mechanisms for autonomous robots that are to collaborate with a human partner, for example in assistive scenarios like helping elderly people in their daily chores.

With the current state of the art in robotics research on such topics using real robots is extremely limited. There are only a few projects worldwide where complete robotic systems are available with stable state estimation, reliable low-level actions, knowledge processing and high-level planning. Even the best of these systems only support rudimentary additional capabilities needed in human-robot interaction like tracking human poses, and recognizing human activities. Besides, state-of-the art robot hardware is still relatively slow (or is slowed down for safety reasons) compared to human movements. Because of these limitations, researchers have started to use realistic physics-based simulations of robots to facilitate the development of autonomous robots.

In this paper we propose to extend the simulation-based development to robots interacting with humans by transforming a robot simulator into a simple computer game and letting people control an agent on the screen. We have implemented such an environment based on a physical simulation with two robots. One of these robots is autonomous, the other is controlled by a human operator.

* This work was partially funded by the cluster of excellence CoTeSys.

The human’s control of the robot is not prerecorded, but happens at the same time while the autonomous robot is working, allowing all problems of misunderstandings and unexpected activities of the human that could occur in the real world.

For evaluating this simulation, we have performed a user study, where we investigated the usability of our simulation and compared human behavior in real life and in simulation. We had subjects perform a household task (setting and clearing the table) both in reality and simulation.

In the next section we argue for the use of simulation for research on human-robot collaboration. After that we present related work. Then we introduce our simulator and the interface for humans to control an agent. This description is followed by an account of the user study we performed to evaluate the viability of using simulation for HRI research.

2 Simulation-based Robot Development

In robotics, simulation-based development is often rejected with the comment “in the real world everything is different than in simulation”. This statement might be true for low-level control loops, but mostly stems from a time when good physical simulations were unavailable. Meanwhile, some extremely successful researchers use realistic simulation for robot development [9], carrying over the techniques developed in simulation to real robots.

The question is also what you want to develop. If the goal is to program a robot performing a certain well-defined job, you engineer this task to the real hardware, simulation would not help much here. But in our research we are rather interested in *general, adaptive* methods for autonomous robots. Our hypothesis is that if we develop techniques that enable a simulated robot to adapt to several simulated environments, it will also be able to adapt to real-world environments.

Simulation gives us the chance to move the development of adaptive robots from the area of “art” into that of “science”:

- Simulation allows us to *focus* research on a particular aspect of robot behavior. Especially, inaccuracy of the state estimation can be ignored or simulated in varying degrees of complexity.
- Experiments can be *repeated*. When using real robots it is extremely hard to get the same conditions for each experiment (e.g. lighting conditions, battery charging level, object positions).
- Simulation makes available *different environments and robots* at low costs.

Beside the “simulation-is-different” objection, HRI researchers often reject simulation claiming that “you can never replace the real user experience”. In our case, we are more interested in developing adaptive execution mechanisms and testing them with the dynamics and uncertainty of real user interaction than in the realistic perception of the scene by subjects. We rather rely on the ability of humans to interpret the abstracted view in the simulation than to sacrifice realism for the robot. This is also the reason why we use a *physical simulation* rather than a virtual rendering machine.

Taking these requirements into account there are good reasons to use simulation for interaction experiments as opposed to working with real robots:

- In the simulator, humans and robots operate on *the same time scale*, which means both are quite slow. Even the best available robots are very slow in sophisticated tasks like pick-and-place actions. In a realistic scenario, a human would perform a “collaborative” task mostly on her/his own.
- Interaction in simulation is a lot *safer* than that with real robots. We’re not thinking of Wizard of Oz experiments where the robots are controlled by humans, but of experiments with autonomous robots, whose behavior can be hard to predict, even by the developers themselves.

Over time, we expect real robots to become faster and safer. We will then be able to carry over the adaptive methods developed in simulation to real robot applications.

3 Related Work

Ueda et al. [9] report on a realistic simulation of deformable objects for a humanoid robot. The simulation works with the normal robot code and can additionally be used by the robot for predictions and planning of its actions. This work demonstrates how a realistic simulation of the robot and its environment directly enhances the performance of autonomous robots.

For research on social imitation learning, Buchsbaum et al. [1] use two virtual characters, where one character learns motoric movements from the other.

The Restaurant Game [6] is a computer game developed at MIT Media Lab, which people can play over the Internet. The goal of this research was the observation of communication and collaboration patterns to develop a kind of social common sense. Plan networks trained on data acquired from this game could distinguish typical restaurant situations from unusual ones [6]. Moreover, the data was used to generate plans for interaction in the Restaurant Game [7]. The Restaurant Game demonstrates very well, how virtual agents can involve a much larger number of subjects than user studies with robots.

There are other approaches of using virtual reality worlds for research on human-robot interaction [2, 10]. But our goal is to support adaptive plan execution for human-robot collaboration. We are mostly interested in the *robot behavior* than in the human reaction, which is in the focus of the well-established Wizard-of-Oz experiments. But still we want our simulated robot to face the dynamic behavior of a human collaborator.

In a similar line Steinfeld et al. describe the “Oz of Wizard” method [8] for HRI research. It shifts the focus from studying humans as practiced in the Wizard of Oz methodology towards the development of robot skills using models of humans.

The OpenRobots Simulator project¹ follows similar goals as our approach to enable human-robot collaboration. This simulator is based on the Blender

¹ <https://launchpad.net/openrobots-simulator>

software to create virtual worlds including the animation of humans. But it also allows physical interaction between robots, humans and objects to make the robot control realistic. The OpenRobots Simulator will be more versatile than the simulation we have developed for our purposes. In the simulation we introduce, we wanted a low-cost solution using an existing robot simulator.

4 The Simulator

For implementing our HRI simulation platform, we made use of an existing physical simulation of an autonomous household robot [5]. For inserting a human in this simulation, it would be desirable to model a human shape and provide it with human-like movements. But as our simulation is based on realistic physics, we would need a very accurate model and would have to provide accurate motor commands to this human to move realistically. Currently, there is no such simulation available. Therefore, we resorted to using the model of the robot that we already have and let it be controlled by a human operator.

We use the Gazebo simulator, which includes the physical simulation engine ODE. The kitchen is a copy of a real kitchen containing furniture such as cupboards, a table, and a sink. The available objects include cutlery (forks, spoons and knives), dinnerware (plates and cups), and objects necessary for cooking such as pots, colanders and wooden spoons. Beside solid objects, the simulation includes water, which can emerge from the tap and disappear into the sink.

We use the same simulated hardware for the human-controlled robot as for the autonomous one. Our robot is modeled after a B21 robot, which originally comes without arms. We added two arms, which are constructed along the PUMA robot arm, used in industrial environments. To make the arms more agile and expand their operating region, we added four more joints: two joints in the shoulder and two slider joints in the upper and lower arms. The additional joints make the arms more agile and extend the robot's work space.

The user interface for manually controlling one of the robots is a stand-alone program written in C, which uses Player interfaces to control the robot and objects in Gazebo. It uses the GTK library for its graphical user interface and for accessing the keyboard and mouse commands from the user. Figure 1 shows the complete user interface. Currently the graphical interface is only necessary for choosing an object for gripping or a position to put an object down. In the future, we intend to expand it for communication like agreeing on a common goal or giving instructions to the robot.

The robot's position can be controlled by the arrow keys to move the robot forward and backward and turn it. Simultaneous rotation and movement is possible, too.

The user can manipulate objects by using preprogrammed arm routines. Gazebo is not designed for identifying objects by clicking on them on the 2D rendering of the scene. Adding this functionality would have been extremely laborious. Therefore, we present all available objects in a list from which the user can select an object to grasp. When the automatic gripping routine fails to find




Fig. 1. The simulation and the control GUI. The entity choice window is opened upon command by the user.

a solution, the user is informed about the error. This happens mostly when the robot is positioned too far away from the object.

For the put-down process, the user can choose among some predefined areas where to put an object. In each area, there are predefined positions for a plate, a cup, a knife, a spoon and a fork, where the robot will automatically put them. Also for putting down, the robot has to be navigated to an appropriate position where the simulation control can find waypoints for the arm.

Although these predefined routines have nothing to do with real human movements, they provide a usable interface. The alternative of letting the users control the arm joints would have resulted in a very difficult user interface, which would rather hinder than enhance the interaction with the robot.

5 User Study

When developing adaptive robot behavior in simulation, we must be sure that the (high-level) behavior of humans is comparable in simulation and reality. We also verify that results are not corrupted by different skills in controlling the simulated agent. More specifically, we conducted a pilot study [4] to answer the following questions: (1) Is our simulation control usable for subjects who were previously untrained in handling the simulation? (2) Do subjects show comparable behavior in simulation and the real world when performing everyday activities?

For the second question, we have to keep in mind the kind of research we want to conduct with this simulation. In this user study, we concentrated on the topic, which is most interesting to us: model-based planning and plan execution of joint human-robot tasks in everyday activities. As real-world tasks for this study, we used table setting and table clearing. The actions allowed by the simulation give the users the freedom to move to any spot in the kitchen at any time and to grasp and put down objects. Although the positions for putting down objects are restricted to a few predefined positions, we believe they did not influence the humans' behavior significantly assuming the subjects complied with the task they were given.

In our experiment we only had one agent, the one controlled by the human. The goal is to use this simulation for studies on human-robot collaboration

(which in fact we have done for proof-of-concept trials). However, for a human-robot scenario, we have no chance to compare the simulated behavior with that in the real world for the reasons given in Section 2. This is why in this user study (simulating robots with the Gazebo simulator is a generally accepted method).

Nine subjects participated in the study, four female and five male. Eight participants work with the computer regularly, six of which are IT professionals (including students of computer science or related subjects).

The study consisted of two parts for each subject: tasks in the real world and control of the robot in simulation. Five subjects performed the real-world tasks first, the others started with the simulation part. To make the data comparable the subjects were asked in the real world to carry only one item per hand, which corresponds to the capabilities of the simulated agent.

We defined six tasks for the subjects to perform (both in reality and in simulation), three consisted in setting the table, three in clearing it. The order was randomized. The available objects (marked with different colors) were plates, cups, knives and in some tasks also spoons and forks. In total, there were three complete sets available (red, blue and yellow), but not all objects were used in every task.

The trials in the real world were recorded with four cameras from different angles for later evaluation. In the video data we can measure durations and observe which objects were used and placed at which positions. In the simulation, we made use of the data acquisition capabilities of the Robot Learning Language [3] to record the robot’s position and orientation 20 times per second, the commands given by the person, failures in the manipulation actions and the original and end positions of objects in the grasping and put down tasks. Beside the quantitative measures, the participants filled out a questionnaire.

We evaluated the ability of the subjects to control the simulation along two criteria: speed and failure rate. The speed of fulfilling the tasks can only be evaluated in comparison to the other subjects. For each scenario we calculated the ratio of the time t_n it took for subject n divided by the average time \bar{t} needed for this scenario by all subjects (t_n/\bar{t}). The maximum deviation from the average for a single scenario is 48%. The average score for each participant was up to 16% higher (i.e. slower) than the average and 22% lower than average.

This deviation seems to be high, but we compared the values to those observed in the real world. The values for the real-world tasks are on a different time scale than those from simulation. Therefore, we compared the deviation for each task normalized by the average duration of all tasks in reality and simulation respectively. This normalized deviation ranged from 16% to 25% in reality and was between 14% and 21% in simulation. The average of this value over all six tasks was 19% both for reality and simulation. So the deviation of times from the average time taken in simulation can be explained by the variance with which people execute real-world tasks.

The variance in the failure rates among subjects is a lot higher than for the speed. We plan to give users more advice on the physical properties of the robot

to help them evaluate the chances of success. But the different failure rates didn't cause a strong effect in the times people needed for completing the tasks.

As a counter check for the measured data, we asked users how they felt about the simulation using a five-level Likert scale². On average, the participants "partially agreed" that they could handle the simulation well and that they could achieve the tasks as quickly as the simulation allows. However, participants were not too happy about the absolute time scale. The average opinion for "the simulated robot executes the actions fast" was rated with an average value of 3.7. Indeed, the scaling factor between the times in reality and simulation is about 20. The relatively high failure rate was not perceived too bad by the participants. The statement "The simulated robot executes the actions without errors." was answered with an average score of 2.7.

For planning joint human-robot activities, we are mostly interested in the actions and the manipulated objects. The actions for table setting as such are not too interesting and don't differ in reality and simulation (moving, grasping, putting down objects); more interesting are the manipulated objects.

We grouped the execution of a task into *carrying tasks*, which involves the gripping of two objects, carrying them and putting them down. We identified preferences of people to carry the same or similar kinds of objects at a time (similar means cups and plates form one class and cutlery another) or objects of the same color and compared them in reality and simulation (Figure 2).




Fig. 2. Comparison of preference types over all subjects. The last line shows the percentage of actions in which any kind of preference (same or similar object or color preference) was applied.

Taken together, these preferences account for 76% of all carry tasks observed in reality and 75% of carry tasks in simulation. Interestingly, when asked about their strategies, the subjects answered that in reality they try to carry objects of the same type or put the plates and cups before they get the cutlery, but that they sacrificed these preferences for the sake of efficiency in simulation. In contrast, our observations suggest that these preferences are also applied in simulation. However, the frequency of each preference is slightly different. In simulation, the subjects carried similar object types slightly more often than in reality and the preference to carry objects of the same type was only visible about half as frequently as in real-world execution. This is compensated by a higher color preference in simulation as compared to reality.

² 1: fully agree, 2: partially agree, 3: don't know, 4: partially disagree, 5: fully disagree

Overall, the way people set and clear the table in the simulator is not very different from how they do it in reality. For example, our simulated robot could have learned from observing the subjects in the experiments that people like to carry objects of the same color. Although this preference is not as strong in reality, it is still a valid observation. As the models of humans have to be adapted for each individual anyway, the development of adaptive planning techniques can very well be done in our simulation testbed.

6 Conclusion

We have presented an implemented approach of using a physical simulation to integrate human behavior in the development cycle for research on human-robot collaboration. The user study we have performed has shown that on an abstract level people show similar behavior patterns when they control a simulated robot as when they perform an activity on their own. Taking into account the current state of the art in robotics and the limited capabilities of complete robot systems, there are good reasons to develop high-level methods for human-robot interaction in such a simulation.

References

1. Buchsbaum, D., Blumberg, B., Breazeal, C., Meltzoff, A.N.: A simulation-theory inspired social learning system for interactive characters. In: IEEE International Workshop on Robots and Human Interactive Communication (2005)
2. Howard, A.M., Remy, S.: Utilizing virtual environments to enable learning in human-robot interaction scenarios. *The International Journal of Virtual Reality* (2008)
3. Kirsch, A.: Robot learning language — integrating programming and learning for cognitive systems. *Robotics and Autonomous Systems Journal* 57(9) (2009)
4. Kirsch, A.: Be a robot — a study on everyday activities performed in real and virtual worlds. Tech. Rep. TUM-I1006, Technische Universität München (2010)
5. Müller, A., Beetz, M.: Designing and implementing a plan library for a simulated household robot. In: *Cognitive Robotics: Papers from the AAAI Workshop*. (2006)
6. Orkin, J., Roy, D.: The restaurant game: Learning social behavior and language from thousands of players online. *Journal of Game Development (JOGD)* 3(1), 39–60 (December 2007)
7. Orkin, J., Roy, D.: Automatic learning and generation of social behavior from collective human gameplay. In: *Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (2009)
8. Steinfeld, A., Jenkins, O.C., Scassellati, B.: The oz of wizard: Simulating the human for interaction research. In: *Proceedings of the 4th ACM/IEEE international conference on Human robot interaction* (2009)
9. Ueda, R., Ogura, T., Okada, K., Inaba, M.: Design and implementation of humanoid programming system powered by deformable objects simulation. In: *10th International Conference on Intelligent Autonomous Systems*. (2008)
10. Xin, M., Sharlin, E.: Exploring human-robot interaction through telepresence board games. In: *Advances in Artificial Reality and Tele-Existence*. Springer (2006)