

# Towards Opportunistic Action Selection in Human-Robot Cooperation

Thibault Kruse, Alexandra Kirsch

Intelligent Autonomous Systems Group  
Department of Informatics  
Technische Universität München

**Abstract.** A robot that is to assist humans in everyday activities should not only be efficient, but also choose actions that are understandable for a person. One characteristic of human task achievement is to recognize and exploit opportunities as they appear in dynamically changing environments. In this paper we explore opportunistic behavior for robots in the context of pick and place tasks with human interaction. As a proof of concept we prototypically embed an opportunistic robot control program, showing that the robot exhibits opportunistic behavior using spatial knowledge, and we validated the feasibility of cooperation in a simulator experiment.

## 1 Introduction

Robots have been identified as a promising approach to enable a more independent life for elderly people. In many cases, elderly people can perform most of their everyday activities on their own, but need assistance in a small number of tasks, that are notwithstanding crucial in the everyday course of life. Our vision is a robot assistant that offers its help to an elderly person to assist in everyday tasks such as preparing meals. In particular, we are interested in planning and plan execution that allow a smooth and natural collaboration.

Human beings are a factor that make the environment of a robot particularly dynamic and difficult to predict. On the other hand, we want to be sure that a human understands what the robot is doing and intending — in other words we want the robot to show not only efficient but also *legible* behavior.

When humans perform tasks, they adapt their behavior to the state of the environment by exploiting opportunities. We think that opportunistic behavior of a robot can benefit its acceptance in two ways: (1) the achievement of goals can be expected to be more efficient, which is probably the underlying reason why humans exploit opportunities and (2) the behavior of the robot will be more human-like and therefore easier to predict for a person. In this paper, we are mostly concerned with the second aspect, although we believe that an efficiency gain can also be expected.

In the context of a robot, we define an opportunity as an alternative or following action to the robot's current action. By this definition, the choice of

the next action also falls in the framework of opportunistic behavior. When considering the next action to take, the different choices can be considered as different opportunities.

The concept of opportunistic planning relies on the detection of opportunities. In principal, opportunities could be defined by a variety of models, which might depend on the execution context. In our work, we have considered one specific model for recognizing opportunities for pick-and-place tasks by using the “reachability” of an object as a measure to rate the desirability of an (alternative) action. The criterion of whether an object can easily be reached is determined with respect to the positions of humans in the world. This means that an object, which is spatially near the robot, but which is blocked by a human, is ranked as less reachable than an object to which a free path can be found. This phenomenon is illustrated in Figure 2a. Although the robot (“Auto”) is nearer to place A, the object lying at place B are ranked as better reachable, because the “human” blocks the way to the objects at A.

In this paper, we introduce a general concept of opportunistic planning and show its application to a realistic robot application. We attempt to validate its usefulness with respect to the legibility and acceptance of humans.

The next section puts our work in perspective with related work. After that we introduce our general concept of opportunities and details of our specific implementation for a robotic household assistant, which is followed by a first validation of our proof-of-concept implementation. We end with conclusions and future work.

## 2 Related Work

The conflicting needs of planning versus reactivity also shows in the classification by Hayes-Roth [5], who describes a continuous space of agent control modes. According to that agent design needs to balance early commitment on actions and sequences with reactivity, to avoid the lack of opportunism that comes with early and persistent commitments.

The same argument is made by Schut and Woolridge [11] who describe a framework that allows agents to make commitment choices at runtime rather than at design time, hence they show this approach to outperform agents with fixed policies. They use the well known BDI framework [1, 4] for their agent design, to embed reactive planning as intention reconsideration in a given control loop of the BDI framework.

They embed the intention reconsideration into a loop, which in sequence reconsiders and executes. However robotics and HRI require even more opportunism than that, because during the execution of an action opportunities can also occur, such as a human giving way to a previously blocked position. The efficiency gain of opportunism mostly be achieved in this kind of plan-execute-loop, however added legibility is achieved by also recognizing opportunities that occur during the execution of a task. The additional efficiency gains to be ex-

pected from that depend on the duration of actions and the frequency of such spontaneous opportunities.

Their computation of the utility of actions is effectively a reactive planner, whose role is to decide which of the currently possible next actions in the plan should currently be pursued. Hayes-Roth has formulated that this way: “At each point in time, many actions may be logically possible, given the current state of the task environment. An intelligent agent must choose among them, either implicitly or explicitly.” [6].

This kind of parallel reconsidering of choices has also been researched in [10] using goodness, competence and conflict as categories for action selection criteria. The robot Milou in that work had a controller which would review alternative navigation plans (what could be considered actions in action selection) every 100ms with regard to three fuzzy measures: goodness, competence and conflict. Goodness is an a priori priority ranking of actions, competence the degree by which the actions are currently possible, and conflict the degree by which one action conflicts with parallel actions. Their robot however has no notion of humans in the environment or of human comfort.

To get more information about related action selection algorithms, we recommend the survey in [2].

### 3 Approach

The basic procedure of our approach is illustrated in Figure 1. The plan, which is executed by the robot is a reactive plan. For the representation of such plans we use the Reactive Plan Language [9], which allows to reason about and change an existing plan at run-time.

The plan execution is constantly monitored by a parallel process that while the current action can be preempted, generates potential opportunities (i.e. alternative actions) and evaluates them according to a given model. If an opportunity is found to be more desirable in the current situation, a second deliberation process (“opportunity filter” in Figure 1) has to check whether the switch between the actions is worth the gain predicted by the model to evaluate opportunities. This step is important to ensure a certain commitment of the robot to the action it has decided upon. By adding costs of switching, the risk of a livelock is reduced. Although in a highly dynamic environment in which humans constantly change the world livelocks might be unlikely, the legibility of the robot’s behavior would suffer by constantly re-deciding its course of action.

This mechanism of choosing opportunities is also valid for the choice of the next action if that has not been decided yet. In this case, the plan is empty and there are no switching costs for setting the next action. The mechanisms of choosing the next action is a greedy choice of the action with the highest utility value in the situation.

In the procedure as we have described it, and also in the specific implementation that we will describe next, the whole procedure is completely reactive. The robot has a set of goals to achieve, which all have the same a priori utility,




Fig. 1: General procedure for detecting and using opportunities.

but may have different dynamic utility. The robot picks goals to achieve greedily from this set of goals and possibly interrupts some action if a better alternative turns up. In this paper, actions could only be interrupted when the robot was in the state of moving towards an item to pick up. However, our general idea of opportunistic planning is that there are some parts of the robot’s course of actions that are planned (possibly offering alternatives at execution time as the Reactive Plan Language allows) and other parts can be marked as possible candidates to exploit opportunities.

In our current implementation of the procedure of Figure 1, we want a robot to set the table together with a person, where the goal positions of all objects are predetermined. This scenario allows us to reduce the set of actions to “atomic” pick-and-place tasks, which means that action selection narrows down to the choice of the object to take next. In the general case, the granularity for opportunistic selection is a design choice that affects how much the robot can exploit opportunity as well as how much computation is required at run-time to compare potential actions.

Because our actions are narrowed down to the choice of objects, we defined our model for evaluating an opportunity based on the human-aware reachability of the objects. The idea is to calculate the costs for a path from the robot’s current position to a position from which it can grip a certain object. The Human-Aware Navigation Planner (HANP) [12] calculates navigation paths for a robot that take into account human comfort as well as the usual safety criteria to avoid obstacles. We use a variant of HANP [8] that takes into account that the humans in the world are also actively involved in the task and are therefore moving most of the time. We use this navigation planner for the navigation of the robot, but also as a model to predict, which objects are reachable in a human-friendly way, and thus as an evaluation function for opportunities.

Using HANP rather than the Euclidean distance or other simpler heuristics has the benefits of both considering the spatial geometry of the environment (walls and furniture) as well as the comfort of the humans present. Also consider the situation in Figure 2a, where spot A is closer to the robot “Auto”, but obstructed by the other agent. Whether at any given time a special opportunity exists for the robot to save overall time can thus be determined as a function of the minimal navigation costs to move to a location of action.

We don’t claim that the human-aware reachability of an object is the only good model that can be employed in this context. Other ideas are to consider the reachability of the position where the object has to be put down, the feasibility

of subsequent actions (putting an object at the edge of the table can make it more difficult to put other objects into the middle of the table afterwards), as well as preferences and abilities of the human collaborator. But when choosing a model, it is important to ensure that it can be calculated reasonably fast. Our current model is in the order of a second to evaluate all alternatives, which is still acceptable. Another consideration would also be how to filter the possible alternative actions to be evaluated.




Fig. 2: (a): Robot “Auto” decides between picking up knife at A or plate at B. (b): User-study simulation, user perspective, steered robot at bottom.

For identifying opportunities that are worth to abandon the current action, we apply a small threshold equivalent to the costs of moving roughly 50cm, such that the robot will not oscillate between 2 opportunities due to mechanical constraints. When there is no current action, the threshold is 0.

## 4 Validation

The validation of opportunistic action planning in human robot interaction requires a setting where the mechanical, spatial and temporal constraints of robots are influencing results. Real robot prototypes have the disadvantage of being unreliable for complex tasks and a safety risk to human subjects.

As a compromise we chose to use a 3D simulator already used in robotics and augment it with a virtual agent to be steered by a human.

In order to show the applicability of our approach, we set up a simulation environment where a human can interact with a robot in a household task of setting the table [7]. In the simulation, two robot-shaped agents are able of picking and placing household items in a 3D kitchen environment, one of them would be steered by a human, the other by an autonomous robot controller.

The human perspective when steering the robot is seen in Figure 2b, it is somewhat behind the robot as suggested to be optimal in [3].

The maximal movement speed for both robots was set to 0.5m/second. Grasping and dropping items in the simulator took about 15 seconds. Those times are constrained by the CPU load for simulation and the controllers of the virtual agents. The autonomous robot would take 25 seconds with most of the time spent on turns on the spot, and slow average velocities while following waypoints. The virtual robot steered by a human would travel similar distances in about 10 seconds.

One of our own trials, shown in Figure 3 illustrates the opportunistic behavior of the robot. The initial situation that was used was one where the robot R would be closer to items at A than items at B. Immediately after launching the robot controller, the Human H was moved towards the Items A. As expected, the Robot R changed its intention when the human presence to A made this paths more costly than a path to B. The independent looping cycles calculating the costs for all the available intentions (10 in this case) too 2 seconds.

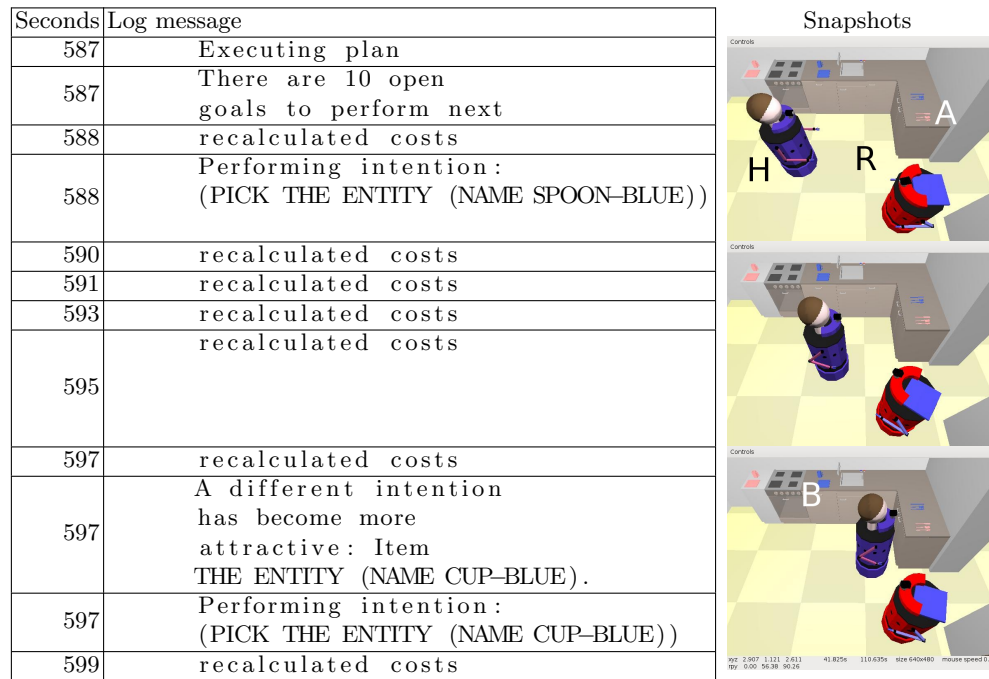


Fig. 3: Logfile and snapshots showing opportunistic change of intention due to spatial circumstance. At time 588 the robot R decided to next pick item SPOON-BLUE shown as location A, and at time 597, the robot instead decided to pick item CUP-BLUE at location B, EVEN though SPOON-BLUE was still closer given Euclidean distance, but has higher path costs due to the human H having moved.

We then performed tests with four different human subjects, male and female students of different disciplines, having to set the virtual table with six items (plates, cups, knife). Due to minor failures such as virtual agents dropping items, some trials became invalid and had to be repeated.

In the experiments the robot calculated the attractivity of open tasks in 1-3 seconds, which was acceptable to the users. When moving to the target, recalculating the attractiveness in parallel took up to 4 seconds, as the motion planner was then also used for actual motion planning, not just for cost estimation. However the robot continuing to its currently committed goal was not adversely impacted by that calculation in parallel.

This shows the approach to use the motion planner for opportunistic action is possible given real-time constraints.

Only in 3 of those 48 moves did the robot change its choice of what item to move next. This is easily explained by the speed of the simulator. In most cases when the robot made an initial choice for an item to pick, the human was already busy in a location of the kitchen and did not move during the robot approach of its target, only in 3 instances did the human move and thus made a different item more attractive.

The subject's rating of the assisting robot was captured on a scale from 1 to 5, 1 meaning not helpful and 5 meaning very helpful, and the average score was 4.2. The subjects were also given the opportunity to give general comments, and the only comments given were about the simulator slowness and of failures due to dropped items, not about the robot behavior.

The experiments could not serve as a basis for further statements about the algorithms, but the mere lack of complaints about the robot behavior may serve as an indication of successful cooperative behavior. The subjects did also not seem to be distracted by the fact that no explicit communication was possible between the autonomous robot controller and the human.

The performance of the plans in our scenario varied a lot due to arbitrary factors, so that even when the robot changed plans, we could not measure the benefit in terms of efficiency.

## 5 Conclusion and Future Work

We have introduced a general framework for detecting and using opportunities in collaborative human-robot tasks. We have focused on a specific model to identify opportunities — the reachability of places for manipulation considering human comfort in navigation tasks.

We combined opportunistic action selection with a model of human-centric motion planning, using a model of opportunities that allows to suppress actions during their execution in favor of other actions, improving the legibility of the set of robot goals.

We conducted a feasibility study with humans, in a realistic joint table setting scenario in simulation, in which the greedy action selection and opportunistic action switching worked as expected.

One thing we could see from preliminary trials is that when in that table setting scenario an agent had to wait once for the other, from then on the human and robot agent did not compete for the same spots, but visited shared spots of interest alternately. This is also a consequence of the rather similar individual pick and place tasks taking the same amount of time. How often such rhythmic patterns occur in real-world cooperation is yet unknown, but could serve as a model for temporal alignment of actions.

We also intend to generalize opportunistic action selection to a finer granularity of tasks, such as combining and interleaving actions, such as picking up one item per robot arm. So far in experiments we only allowed for opportunistic selection of pick and place actions as a whole. A higher impact on robot efficiency (and thus hopefully also on human satisfaction) may be achieved by reactively planning a bit more into the future, such as not just considering the effort to a pickup task, but also the expected effort of placing it in the immediate future. In our example this could help the robot to decide setting the table for a different person, and thus conflicting less at the table.

As research exists on stochastic prediction of human space occupancy for trajectory planning[13], a combination of those works with ours would be promising.

## References

1. M. Bratman. *Intention, Plans, and Practical Reason*. Harvard University Press, Cambridge, Massachusetts, 1987.
2. C. Brom and J. Bryson. Action selection for intelligent systems. *European Network for the Advancement of Artificial Cognitive Systems*, 2006.
3. D. J. Bruemmer, D. A. Few, and C. W. Nielsen. Spatial reasoning for human-robot teams. In B. N. Hilton, editor, *Emerging Spatial Information Systems And Applications*, pages 350–372, Hershey, PA, USA, 2006. IGI Publishing.
4. M. P. Georgeff, B. Pell, M. E. Pollack, M. Tambe, and M. Wooldridge. The belief-desire-intention model of agency. In *ATAL*, pages 1–10, 1998.
5. B. Hayes-Roth. Opportunistic control of actions in intelligent agents, 1992.
6. B. Hayes-Roth. Intelligent control. *Artificial Intelligence*, 59:213–220, 1993.
7. A. Kirsch and Y. Chen. A testbed for adaptive human-robot collaboration. In *Proceedings of the 33rd Annual German Conference on Artificial Intelligence*, 2010.
8. T. Kruse, A. Kirsch, E. A. Sisbot, and R. Alami. Dynamic generation and execution of human aware navigation plans. In *Proceedings of the Ninth International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2010.
9. D. McDermott. A reactive plan language. Technical report, Yale University, Computer Science Dept., 1993.
10. P. P. Saffiotti, S. Parsons, O. Pettersson, A. Saffiotti, and M. Wooldridge. Robots with the best of intentions, 1999.
11. M. Schut and M. Wooldridge. Principles of intention reconsideration. In *AGENTS '01: Proceedings of the fifth international conference on Autonomous agents*, 2001.
12. E. A. Sisbot, L. F. Marin-Urias, R. Alami, and T. Simeon. A human aware mobile robot motion planner. *IEEE Transactions on Robotics*, 23:874–883, 2007.
13. B. Ziebart, N. Ratliff, G. Gallagher, C. Mertz, K. Peterson, J. A. D. Bagnell, M. Hebert, A. Dey, and S. Srinivasa. Planning-based prediction for pedestrians. In *Proc. IROS 2009*, October 2009.