



HAL
open science

Social acceptability of opportunistic behaviour of an assistant robot in human–robot everyday collaboration

Nataliia Kushnirenko, Alexandra Kirsch

► **To cite this version:**

Nataliia Kushnirenko, Alexandra Kirsch. Social acceptability of opportunistic behaviour of an assistant robot in human–robot everyday collaboration. Workshop “Planen, Scheduling und Konfigurieren, Entwerfen”, 2014, Stuttgart, Germany. ⟨hal-01691689⟩

HAL Id: hal-01691689

<https://hal.science/hal-01691689v1>

Submitted on 24 Jan 2018

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.



HAL Authorization

Social acceptability of opportunistic behaviour of an assistant robot in human–robot everyday collaboration

Nataliia Kushnirenko, Alexandra Kirsch

Department of Computer Science
University of Tuebingen
Sand 14
72076 Tuebingen
nataliia.kushnirenko@uni-tuebingen.de
alexandra.kirsch@uni-tuebingen.de

Abstract: In dynamically changing environment, an autonomous robot and human, performing a joint task together, should take decisions depending on many factors. A fundamental feature of human decision–taking is opportunism, and in this paper we explore how socially acceptable is opportunistic behaviour of robot in every day life situation with human interaction. We conducted a pilot study to assess whether a robot using opportunities is regarded as a better companion. Showing videos of a simulated robot, we asked for personal evaluations of subjective factors as helpfulness, politeness and competency.

1 Introduction

Making autonomous robots act understandably for humans is an important prerequisite for the human–robot collaboration in everyday life. We are interested in the design and implementation of flexible planning and plan execution strategies for robots. One aspect that characterizes human activity is that actions are not all planned in advance, but that humans react to opportunities for efficient action selection. For example, when we tidy up a room, we store away the things as we see them lying around, rather than making a list of all things to be tidied up and then storing the objects respectively. According to Botvinick [Bot08] there is a “[...] research that suggests the existence of two systems for action control: a habit system based on context–response associations and a goal–directed system operating through the anticipation of action outcomes.” This means that it would be unnatural for an assistant robot to suggest a complete plan to a person to execute jointly, because humans decide their next action from the perceptual context rather than a premeditated plan [SWB97].

In the context of our work on a robot assistant we are specifically focusing on opportunities that appear from dynamics caused by humans and how to detect and exploit them to behave

in a legible¹, human-friendly way.

It is very difficult to evaluate robot behaviour efficiently with regard to user acceptance. The very nature of opportunities requires a variety of situations with different kinds of opportunities arising in different contexts of task execution. Also, one might argue that a robot does not need to be opportunistic, but it should rather be predictable [Nor07]. Thus, before investing time in a specific strategy, we conducted a pilot study to assess whether a robot with some opportunistic behaviour is better accepted than one following a strict plan.

We measure the performance of opportunistic planning along the social acceptability dimension, which includes objective measures such as waiting effort of the human during the task, the qualities of the robot being safe, pleasant, friendly, polite, competent, knowledgeable, responsible, intelligent, efficient, helpful, able to predict human intention, and the level of social willingness to buy this kind of robot. This self-assessment is complemented by open questions.

We present a video-based study in which the participants watch a scene without being occupied with a task, and rate the video according to proposed criteria.

2 Background and related work

Kruse and Kirsch [KK10] have proposed an architecture for opportunistic action selection as shown in Figure 1. A parallel process is monitoring the plan execution, generates possible actions and evaluates them according to a given model. In case of positive evaluation result an opportunity filter is checking if a change to these alternative actions would bring predicted benefit according to the model. In our study, we use this architecture




Figure 1: Detecting and using opportunities.

and examine possible opportunity models. The used architecture builds on concepts from goal-driven (or goal-directed) autonomy [MKA10] with its most prominent architecture, the Belief-Desire-Intention (BDI) framework [Bra87, GPP⁺98]. Similar to behaviour-based approaches an agent has abstract desires, which can be compared to motivations of Sevin and Thalmann [dST05], but more often represent goals [PH99, MKB04]. A desire can become an intention when the agent tries to achieve a specific goal and corresponds to a plan to achieve this goal.

¹By legibility we describe the ability of humans to understand the intentions of the robot intuitively by observation of its behaviour, which makes it possible to anticipate the robot's next actions.

Under the term “intention reconsideration” Schut and Wooldridge [SW01] have described a general framework in the BDI model, in which the reconsideration of the current goal is treated like a physical action. In their framework, the cost of a reconsideration action is compared to the benefit that could be obtained by changing the course of action. In this way, an agent only reconsiders its intended course of action if it believes that it can significantly improve its performance.

Reconsideration strategies have also been explored by Pollack and Ringuette [PR90]. They introduced the notions of bold and cautious agents. They call an agent cautious when the agent is “very sensitive to its environment, willing to reconsider its plans in response to a wide range of events”. The same holds for opportunistic agents, the difference being that a cautious agent reconsiders to avoid making mistakes, an opportunistic agent reconsiders to be more acceptable and possibly more efficient.

Other examples for plan reconsideration in robotics are given by Helwig and Haddawy [HH96] for a truck–world simulation, or by Miura and Shirai [MS02] for an office robot. All these approaches work on operators defined by first–order predicates. But opportunities are often more interesting when considering more specific knowledge. This is taken into account by Parsons et al. [PPSW00] who used feedback from a fuzzy controller to recognise opportunities in the framework of Schut and Wooldridge [SW01].

Broz, Nourbakhsh, and Simmons [BNS08] consider the adequate timing of robot actions with respect to humans using POMDPs. Cirillo, Karlsson, and Saffiotti [CKS08] describe a planning framework in which a robot takes into account the activity of nearby humans and adapts its actions to not disturb or distress humans. This approach considers human comfort by avoiding the human, rather than contributing to a joint goal. In contrast, legibility in joint tasks is an explicit goal of the Human–Aware Task Planner by Alami et al. [RAC06]. They use an HTN planner to generate plans not only for the robot, but also for a human partner. Legibility is modelled as a cost function for plans. This planner is integrated into a classical three layer architecture [ACF⁺98] and thus cannot deal directly with the dynamics of the world. A new plan is only generated when the current one fails. This approach assumes that the robot can plan for itself and the human and that the human will accept this plan.

In other contexts, there are approaches that explicitly take into account the movement of humans [ZRG⁺09] or even the intentions of humans [CvSK⁺06, BTJB10]. Benson and Nilsson [BN95] use a more general model for opportunistic behaviour. Their architecture includes an inherent trade–off between the priority of an action and the estimated time to achieve it.

The models mentioned so far all assumed that the information about actions is available by definition [BN95, dST05], which is unrealistic in everyday environments, or is provided by a lower–level model that also achieves the given navigation goals [PPSW00, KK10], which might entail problems for execution efficiency.

To our knowledge, the acceptance of opportunistic behaviour has never been shown.

3 Method

We are specifically interested in using opportunistic methods in everyday real-world environment for a robot that collaborates with humans. Therefore a collaborative table-setting scenario was chosen for testing social acceptability.

To this end we have conducted a user study to evaluate the social acceptability of robot behaviour with a set of videos recorded with the Morse simulator, in which a human and a robot are setting the table together. The Morse ² simulator [LGK⁺12] is a project of LAAS-CNRS to supply researchers with a realistic, physical simulator for a wide variety of robot tasks. This simulator is based on the free 3D modelling software Blender and allows us to include an animated human that can move in the world and manipulate objects. We use the PR2 robot for demonstrating opportunistic planning. The basic functionality is provided by open-source ROS modules. Both the robot and the human can be controlled with the keyboard as in a computer game. For the recording the videos for our study the robot was controlled by a control script, while human avatar was moved manually.

The kitchen environment where the robot is bringing the plate and the human is holding the spoon is shown in Figure 2.



Figure 2: Robot and human performing the tasks in our kitchen environment.

The user study was not held as Wizard of OZ experiment, because the task of controlling human avatar might both distract user and increase the number of possible reasons of certain robot behaviour observed. Moreover, in this case each participant would have different experience. A real robot would be even more complicated object to judge, therefore, keeping experiment fixed should allow us to figure out which exactly kind of robot behaviour would be more acceptable.

The videos for the user study were generated considering different combinations of opportunity models ³. We examined the following models: 1) spatial comfort for the human

²<http://morse.openrobots.org>

³These models were taken into account for constructing the scenes. The robot did not autonomously decide, but was scripted.

when the robot moves to an object; 2) the time to grasp an object, taking into account if cupboard doors are currently open or closed. In addition, we were interested in the degree of collaboration and mutual considerateness.

3.1 Participants

33 participants of the age 21–40 years with the average age of 28 years took part in our experiment — thereof 12 women and 21 men. 30 of them did not have experience with robots. The number of people taking the test simultaneously varied from one to four.

3.2 Procedure

After personal data collection and short introduction with the general task description (both printed and read aloud by the experimenter), users were asked to watch the videos in random order and to answer a set of questions on paper after each video. The experimenter's role was mainly to observe the users watching the videos.

Since the movements in Morse are rather slow, all videos were sped up to 218%. In a pilot study with three participants the normal speed of the movements was used and those results were comparable to those of the main study. This shows that the video speed does not seem to affect the users' opinions.

The videos can be seen at <http://vimeo.com/hcai/videos>. The first video, lasting for 2:33 minutes, was used as an introduction to make observers used to the robot in the simulated environment.

All videos include scripted behaviour of the robot except the last video with the shifting objects scenario, where the robot was controlled manually. Every video represents a certain situation from the real world.

3.3 Videos description

Each of the offered videos is based on one of scenarios, summarized in Table 1 with parameters. Opportunity is an action the robot can perform alternatively or following the current task considering the dynamics of the environment, such as picking a different object, which will be useful later rather than waiting for the free space to reach the object according to the current plan; or using a chance to take the item from an opened drawer, as the reachability of this item might become worse later.

The videos with opportunities used include situations when the man is blocking the way, and the robot has changed the original plan in favour of an alternative action, such as picking the object from the other side of the table; using a chance to take the object from the opened drawer or pick another item to avoid waiting.

Table 1: Video scenario characteristics

	Opportunity used	Collaboration	Robot waits	Human waits
Video 1 (Introduction)	yes	together	no	no
Video 2 (Human, then robot)	no	one by one	yes	no
Video 3 (Human waits)	no	together	no	yes
Video 4 (Two cups)	yes	together	no	no
Video 5 (Robot waits)	no	together	yes	no
Video 6 (Opened door)	yes	together	no	no
Video 7 (Shifting)	no	one by one	yes	-

Collaboration type can be mutually dependent (“together”), as the robot and human are moving and performing the tasks simultaneously, or almost independent (“one by one”), when the robot only acts during the time the human is far away or inactive.

We are using “robot waits” and “human waits” as parameters to indicate when pure waiting time of robot or human respectively was obvious.

Both qualities safe and polite are combined into one column in Table 1, as according to our imagination, in presented situations the robot is either both safe and polite, or unsafe and impolite.

3.4 Scenario

The main scenario of the videos is a man who brings the knife, the fork and the spoon to the table, while the robot is in charge of bringing the plate, milk and cup to the same table in the kitchen. We tried different variations of this scenario, so that situations when robot and man are completing the task one by one, with and without collision, with the opportunity used, with the opportunity used due to the human blocking the original way, with robot ignoring the human.

Let us describe the offered videos related to the opportunities provided:

Video 1:

The introduction video shows human and robot setting the table together one by one without having any conflicts till some point, when the human is blocking the way for the robot to reach the table, and the robot starts slightly shifting human to the side. As soon as the human leaves the way free, the robot continues completing the task of taking the blue cup, but chooses a new opportunity of taking it from the other side of the table. The results of this video are separated and are not included in evaluation.

Video 2:

First the human brings the fork, the knife and the spoon, and only after the human

stops the robot brings the plate, milk and cup. This scenario does not contain opportunities, because robot's behaviour can be repeated with a usual plan if there is no human on the way.

Video 3:

Whenever there is a collision on the way, the human is waiting for the robot to finish its task. The robot does not use any opportunity, simply following the plan.

Video 4:

When the human is occupying the place near the table, the robot uses an opportunity to take the cup from another place with help of both hands to take two cups.

Video 5:

Whenever a collision occurs, the robot waits until the human frees the space and never uses any opportunity.

Video 6:

The human forgets to close the drawer, and the robot uses this opportunity to take the bowl from it. The human closes the drawer afterwards.

Video 7:

The human is bringing items to the table, and the robot is shifting them organizing more space while the human is far from the table. The robot uses opportunities to work around the table while the space is free from human.

3.5 Questionnaire

In order to evaluate social acceptability of the robot, the questions were created partially on the basis of the Godspeed test [BKCZ09]. The users were asked to choose the number from 1 (the best) to 5 (the worst) on a semantic differential scale to rate the selected qualities of the robot shown in each video: Safe–Unsafe, Pleasant–Unpleasant, Friendly–Unfriendly, Polite–Impolite, Competent–Incompetent, Knowledgeable–Ignorant, Responsible–Irresponsible, Intelligent–Unintelligent, Efficient–Inefficient, Helpful–Unhelpful. These items were chosen in order to keep the questionnaire clear, and to allow the user to describe the robot behaviour feedback also in other interesting aspects for better evaluation.

Despite some of these items seem to be closely related, we selected them for user rating to be sure in the final results [BM10]. These measurements were supplemented with a set of open questions, namely:

1. Please describe the strategy: did you notice any pattern of how the robot and human interacted?
2. Did you notice particularly bad or annoying robot behaviour?
3. Did you notice particularly good robot behaviour?

4. Where you surprised at some point? (What were you expecting instead?)
5. Does the robot predict human movement? (Yes/No)

The question about surprise is an additional subjective criterion to evaluate other aspects of the robot behaviour, helping to show if the robot's action and goal can be inferred correctly.

The last question was asking to put the overall rating for the video in the points from 1 to 5 expressing the willingness to buy the robot behaving as in the shown video. After completing the survey for all videos, the participants were asked to respond whether the robot or a human was faster, and to provide any other comments.

4 Results and evaluation

The results of the user ratings are presented in Figures 3 and 4 showing the chosen values from 1 to 5 for the certain quality in the group of videos with similar parameter according to Table 1. The smaller number corresponds to more positive rank, e.g. "1" is safe and "5" is unsafe. The biggest difference in the results is seen for the groups of videos related to collaboration: together or one by one, and related to the opportunity usage: with or without opportunities. Both figures show, that there is only slight change in the ratings for "polite" quality, in the one-by-one collaboration scenario the robot is safer, more competent, helpful and the will to buy it is also higher. The user ratings of the same robot qualities for the opportunity usage video group barely change.

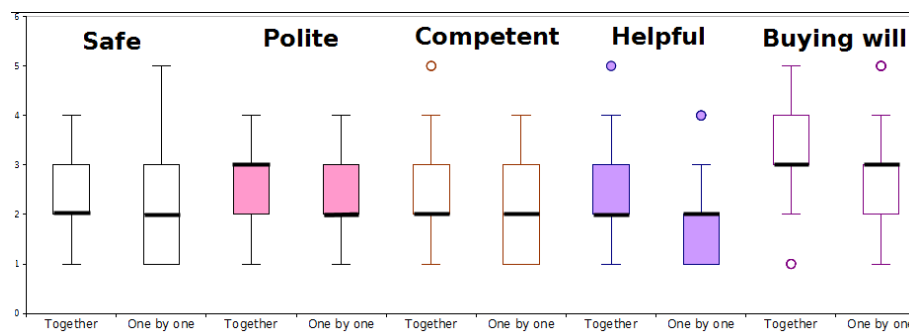


Figure 3: Collaboration dependent qualities rates

The robot only acting while the human inactive has the highest average ranks (2,6 for "I would buy it"), as well as the robot which has to wait for the human to free the space needed to complete the task (2,9 for "I would buy this robot").

As expected, the situation when the human has to wait for the robot is the least pleasant (average rank 3) and safe (2,45). Also, the least safe and pleasant the robot seems, the less will to buy the robot exists, despite the other qualities were rated pretty high.

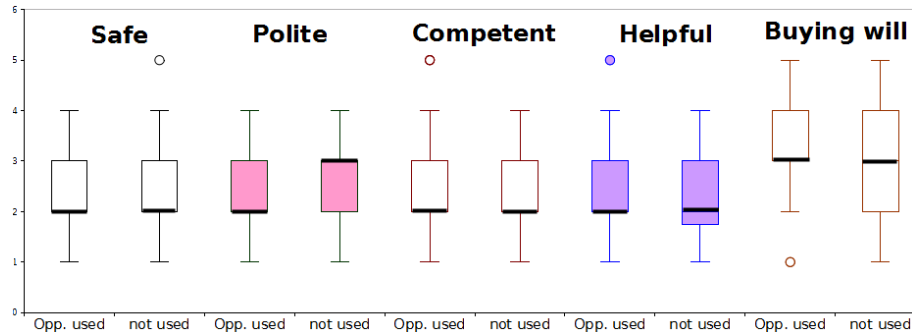


Figure 4: Opportunities usage dependent qualities rates

The robot is evaluated to be less helpful when it makes the human wait. The scenario with shifting items on the table (video 7) shows the robot behaviour becomes efficient, but does not make people want to buy the robot.

During the user study participants had small doubts such as wonder why did not robot close the drawer in the video 6, or why would one human need two cups in video 4. Participants thought the robot remembers how many people live in the apartment or is perfectly aware of the human habits and therefore predicts his intentions. A couple of users were sure the robot repeats after the human trying to help, learning from the environment right away. The scores for the opportunistic behaviour are slightly better in safety dimension, but less in helpfulness in the scenarios with the opportunities used.

Interesting comments from the study came from three users who were sure there was verbal communication between the human and the robot, despite the videos being silent. Two users thought there is a remote control for the robot. Some people proposed to avoid collision situations with robot by separating the workspace or giving a certain task to the robot based on how heavy the object is.

Some participants mentioned not only the human safety, but also the safety of objects placed: in their opinion, video 7 demonstrates how careful the robot is to keep the items far from the edge of the table, or to find a proper place for them.

We used only some variables in the previous evaluation, because the others are highly correlated, which can be seen from Table 2. It also shows how independent are the qualities intelligent or efficient from friendly, pleasant, polite or safe. The coefficient which is 0.5 or higher is considered as strong correlation [PS08], and is highlighted.

All participants claim the speed of the robot was slow, but this fact brings contradicting thoughts: slow is good for safety, but sometimes not efficient enough.

The answers on the question if the robot predict human stay on the same shape of the lines presented in Table 3. Also our study shows the more robot predicts human behaviour, the more interest in buying the robot is present.

Table 2: Pearson correlation coefficient for all videos

	Safe	Pleasant	Friendly	Polite	Competent	Knowledgeable	Responsible	Intelligent	Efficient	Helpful
Pleasant	0.53									
Friendly	0.46	0.69								
Polite	0.51	0.59	0.76							
Competent	0.43	0.49	0.4	0.37						
Knowledgeable	0.4	0.46	0.4	0.46	0.73					
Responsible	0.59	0.49	0.53	0.6	0.59	0.68				
Intelligent	0.26	0.34	0.26	0.2	0.55	0.6	0.55			
Efficient	0.28	0.36	0.32	0.34	0.53	0.5	0.51	0.47		
Helpful	0.39	0.47	0.46	0.45	0.63	0.6	0.55	0.52	0.68	
I would buy it	0.38	0.42	0.4	0.39	0.49	0.46	0.52	0.48	0.46	0.49

Table 3: User opinion on whether robot predicts human in each video.

Video	"Yes" for robot predicts human, %
Video 1 (Introduction)	50.00
Video 2 (Human then robot)	44.44
Video 3 (Human waits)	55.56
Video 4 (Two cups)	50.00
Video 5 (Robot waits)	75.86
Video 6 (Opened door)	53.57
Video 7 (Shifting)	53.33

5 Conclusion

The question should the robot actually change the behaviour due to dynamics of the environment, or to do the first action intended was raised. We examined different scenarios with and without opportunity usage and asked users to rate them with respect to efficiency, friendliness, usefulness and legibility of the resulting behaviour. The spatial comfort and waiting effort of the human were taken into consideration to analyse the applicability to real-world environment.

To conclude, with our results in the social acceptability of opportunistic behaviour of an assistant robot in human-robot everyday collaboration, we could not support the finding of Norman saying that a robot does not need to be opportunistic, but it should rather be predictable [Nor07]. There is also no reliable indication that opportunism is better accepted, it is possible due to the small number of videos used, and due to existence of many other differences in scenarios, than just using or not using opportunities, which is also not easy to show clearly. The other distracting factors, such as robot movement, the task performed and various little details and changes influence user opinions more than opportunity-taking.

Furthermore, our results show how hard evaluation of autonomous behaviour is, as the real robot is even more uncontrolled than the robot with scripted behaviour we used basing on proposed model. We should be careful in setting up a truly realistic task, keeping in mind how important all the details are for user opinion, and to show as many interactions as possible. It would be also good to use larger semantic differential scale to be able to receive more precise answers for evaluation. For future studies objective measurements such as waiting times, time to reach the goal will be added.

We should mention that every person would behave differently in proposed situation. This research was using a passive study where users could only watch the videos. In future, the interactive model where users can control the human avatar in the simulated environment will be used for evaluation.

References

- [ACF⁺98] R. Alami, R. Chatila, S. Fleury, M. Ghallab, and F. Ingrand. An Architecture for Autonomy. *International Journal of Robotics Research*, 17(4), 1998.
- [BKCZ09] Christoph Bartneck, Dana Kulić, Elizabeth Croft, and Susana Zoghbi. Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots. *International Journal of Social Robotics*, 1(1):71–81, 2009.
- [BM10] Cindy L. Bethel and Robin R. Murphy. Review of Human Studies Methods in HRI and Recommendations. *International Journal of Social Robotics*, 2:347–359, 2010.
- [BN95] S. Benson and N. Nilsson. Reacting, Planning and Learning in an Autonomous Agent. In K. Furukawa, D. Michie, and S. Muggleton, editors, *Machine Intelligence 14*. Oxford: The Clarendon Press, 1995.
- [BNS08] Frank Broz, Illah Nourbakhsh, and Reid Simmons. Planning for Human-Robot Interaction using Time-State Aggregated POMDPs. In *Proceedings of the NCAI AAAI*, 2008.
- [Bot08] Matthew M. Botvinick. Hierarchical models of behavior and prefrontal function. *Trends in Cognitive Sciences*, 12(5), 2008.
- [Bra87] M. Bratman. *Intention, Plans, and Practical Reason*. Harvard University Press, Cambridge, Massachusetts, 1987.
- [BTJB10] Michael Beetz, Moritz Tenorth, Dominik Jain, and Jan Bandouch. Towards Automated Models of Activities of Daily Life. *Technology and Disability*, 22(1-2):27–40, 2010.
- [CKS08] M. Cirillo, L. Karlsson, and A. Saffiotti. A Framework For Human-Aware Robot Planning. In *Proc. of the Scandinavian Conf on Artificial Intelligence (SCAI)*, Stockholm, SE, 2008.
- [CvSK⁺06] Raymond H. Cuijpers, Hein T. van Schie, Mathieu Koppen, Wolfram Erlhagen, and Harold Bekkering. Goals and means in action observation: A computational approach. *Neural Networks*, 19(3):311–322, 2006.
- [dST05] E. de Sevin and D. Thalmann. A motivational model of action selection for virtual humans. In *CGI '05: Proceedings of the Computer Graphics International 2005*, pages 213–220, Washington, DC, USA, 2005. IEEE Computer Society.

- [GPP⁺98] Michael P. Georgeff, Barney Pell, Martha E. Pollack, Milind Tambe, and Michael Wooldridge. The Belief-Desire-Intention Model of Agency. In *ATAL*, pages 1–10, 1998.
- [HH96] James Helwig and Peter Haddawy. An Abstraction-Based Approach to Interleaving Planning and Execution in Partially-Observable Domains. In *AAAI Technical Report FS-96-01*, 1996.
- [KK10] Thibault Kruse and Alexandra Kirsch. Towards Opportunistic Action Selection in Human-Robot Cooperation. In *33rd Annual German Conference on Artificial Intelligence (KI 2010)*, 2010.
- [LGK⁺12] S. Lemaignan, Echeverria G., M. Karg, M. Mainprice, A. Kirsch, and R. Alami. Human-Robot Interaction in the MORSE Simulator. In *Proceedings of the 2012 Human-Robot Interaction Conference (late breaking report)*, 2012.
- [MKA10] Matthew Molineaux, Matthew Klenk, and David Aha. Goal-Driven Autonomy in a Navy Strategy Simulation. In *AAAI Conference on Artificial Intelligence*, 2010.
- [MKB04] Armin Müller, Alexandra Kirsch, and Michael Beetz. Object-oriented Model-based Extensions of Robot Control Languages. In *27th German Conference on Artificial Intelligence*, 2004.
- [MS02] Jun Miura and Yoshiaki Shirai. Parallel scheduling of planning and action for realizing an efficient and reactive robotic system. In *ICARCV*, pages 246–251, 2002.
- [Nor07] Don Norman. *The Design of Future Things, Chap. 3*. Basic Books, 2007.
- [PH99] M. Pollack and J. Horty. There’s More to Life than Making Plans: Plan Management in Dynamic, Multi-Agent Environments. *AI Magazine*, 20(4):71–84, 1999.
- [PPSW00] Simon Parsons, Ola Pettersson, Alessandro Saffiotti, and Michael Wooldridge. Intention reconsideration in theory and practice. In *Proceedings of Fourteenth European Conference on Arti Intelligence (ECAI-2000)*, 2000.
- [PR90] M. Pollack and M. Ringuette. Introducing the Tileworld: Experimentally evaluating agent architectures. In *Proceedings of the Eighth National Conference on Artificial Intelligence*, pages 183–189, Boston, MA, 1990.
- [PS08] Frank Renkewitz Peter Sedlmeier. *Forschungsmethoden und Statistik in der Psychologie*. Pearson, 2008.
- [RAC06] Vincent Montreuil Emrah Akin Sisbot Rachid Alami, Aurelie Clodic and Raja Chatila. Toward Human-Aware Robot Task Planning. In *AAAI Spring Symposium*, 2006.
- [SW01] Martijn Schut and Michael Wooldridge. Principles of intention reconsideration. In *AGENTS ’01: Proceedings of the fifth international conference on Autonomous agents*, 2001.
- [SWB97] N.B. Sarter, D. D. Woods, and C.E. Billings. Automation Surprises. In G. Salvendy, editor, *Handbook of Human Factors & Ergonomics, second edition*. Wiley, 1997.
- [ZRG⁺09] Brian Ziebart, Nathan Ratliff, Garratt Gallagher, Christoph Mertz, Kevin Peterson, J. Andrew (Drew) Bagnell, Martial Hebert, Anind Dey, and Siddhartha Srinivasa. Planning-based Prediction for Pedestrians. In *Proc. IROS 2009*, October 2009.