

NOTICE: This is the author's version of a work accepted for publication by Springer. The final publication is available at www.springerlink.com:

<http://www.springerlink.com/content/7074566841305267/>

Plan-based Control of Joint Human-Robot Activities

Alexandra Kirsch · Thibault Kruse · E. Akin Sisbot · Rachid Alami · Martin Lawitzky · Dražen Bršćić · Sandra Hirche · Patrizia Basili · Stefan Glasauer

Received: date / Accepted: date

Abstract Cognition in technical systems is especially relevant for the interaction with humans. We present a newly emerging application for autonomous robots: companion robots that are not merely machines performing tasks *for* humans, but assistants that achieve joint goals *with* humans. This collaborative aspect entails specific challenges for AI and robotics. In this article, we describe several planning and action-related problems for human-robot collaboration and point out the challenges to implement cognitive robot assistants.

1 Motivation

Research on intelligent autonomous robots has mainly focused on machines that perform certain services *for* humans, such as robot butlers or autonomous cars [24, 26, 16, 7]. The interaction in such systems is restricted to a user interface to give a task to the robot and when acting, such a robot has to ensure the safety of humans in its vicinity.

A new class of robot applications takes a different view: a robot assistant performs tasks *together with* a human in order to achieve a joint goal. This approach combines the cognitive capabilities and manual dexterity of humans with the strength, endurance and accuracy of robots. Rosenthal et al. [20] use the term “symbiotic relationship” to describe the reciprocal complementation of skills.

On the one hand, the combination of human and robot skills moves useful robot applications closer to the state of marketable products. Whereas a robot butler would be expected to prepare a cup of tea all by itself, a robot assistant might only be expected to bring the kettle and the tea for a person to prepare the tea herself. On the other hand, the close interaction and collaboration with humans requires a range of new cognitive capabilities for robots in order to coordinate the human’s and robot’s actions in a way that feels natural to people.

This paper is concerned with the action side of human-robot collaboration. This includes AI techniques for planning and plan execution as well as specific action primitives for navigation and joint manipulation tasks.

1.1 Application Scenarios

To illustrate the idea of assistant robots, we present two application scenarios where a collaboration of a human and a robot can be of substantial help, even if the robot has restricted capabilities.

The aging society is a pressing problem in all industrialized countries. Attending elderly people in nursing homes is not only costly, but is also perceived by most elderly people as diminishing the quality of life. The need to move into a nursing home is often caused by the inability to perform activities of daily living such as eating, washing and dressing, which can be caused even by minor bodily or cognitive impairments.

Collaborative robot assistants are a promising approach to enable elderly people with minor disabilities to live an independent life for a longer time. In such a scenario, the robot would not take over all of the person’s tasks, but rather stimulate the human to be ac-

Dr. Alexandra Kirsch
Technische Universität München
CoTeSys Central Robotics Laboratory CCRL
Barer Str. 21
D-80290 München
Tel.: +49 89 289 25726
Fax: +49 89 289 25724
E-mail: kirsch@in.tum.de

tive and participating and assist in tasks that the person cannot perform on his own. This behavior would lead to a longer independence of elderly people in two ways: 1) the robot would assist in tasks that the person cannot perform because of some minor disabilities and 2) the person would be encouraged to keep performing as much of his everyday activities on his own. Note that such assistive robots do not need to be capable to perform *general* household activities completely autonomously. In consequence, the time to market for such robots is reduced in comparison to general fully autonomous systems.

Another interesting application domain for collaborative robots are workplaces such as factories or office buildings. The demands on modern manufacturing are moving towards mass customization, which means that production faces high flexibility requirements. The machines used in current production environments are very accurate and reliable, but lack the necessary flexibility. On the other hand, humans have the capability to adapt quickly to new requirements and situations. The combination of the flexibility of human workers and the strength and accuracy of machines is a promising approach for future manufacturing environments.

Although production environments can be adapted to the needs of robots (as is currently the case), the cognitive workload for humans in such an environment must not be neglected. The use of technology sometimes has not led to the expected advantages, but has created new problems. For example in the domain of automated airplane control missing transparency of the system and limited coordination between humans and machines has caused serious accidents [22]. Collaborative robot technology would provide a solution, which takes into account human needs at the workplace and provides the additional advantages of automation.

1.2 Challenges

The quality and usefulness of autonomous robots is usually measured in terms of time efficiency, quality of task achievement and reliability. When collaborating with a human, the focus shifts from these quantitative, measurable units to human-centered evaluation criteria. Most important for the interaction is that the human can predict and understand the robot's actions without additional need of communication. We use the term *legible* to describe behavior that is intuitively understood by humans as part of the joint plan [1].

Achieving *legible behavior* requires substantial knowledge of how humans achieve their tasks, especially when collaborating with others. Human behavior is usually quite efficient, but it is also largely guided by habits,

cultural conventions and social rules. For example, when constructing a product in a factory, a robot that brings the component parts must know in which order the human worker likes to combine those parts and when the robot should be ready to assist with a joint manipulation task. When a part is manipulated by both partners, the robot must also be aware of the specific preferences of the worker to apply the correct force for the joint manipulation task. Similarly, when navigating, the robot must make clear where it is heading and must make sure that humans near its path do not feel threatened.

Legibility is connected to the perceived *safety and comfort* of humans. Ensuring the safety of humans is a requirement for all kinds of robots. But while this criterion is relatively easy to fulfill in industrial environments, where robots are spatially separated from humans by fences or other safety measures, a robot acting in everyday household or work environments must possess enough cognitive capabilities to assess situations correctly and always act in a safe way. What is more, humans have a subjective feeling of safety and comfort. Even if a robot has been developed with the utmost safety standards, a person might feel uncomfortable if a robot passes closely behind his back or moves its arms very fast. Legibility is a prerequisite to establish human comfort — if the intention of the robot is understandable, its actions can be expected and are not felt as a threat — but it must be complemented by more specific considerations.

Another challenge is the human as a collaboration partner for the robot. Humans are very *flexible* in how they decide on and achieve their goals. A person might not even realize that she took different paths in a building to reach the same location or that she changed the order of constructing a product. But for a robot, such small changes raise the question if this is a normal behavior deviation or if the human has committed some kind of error or has changed her mind and is trying to achieve a new goal.

In the following sections, we explain in more detail our research on plan-based control for joint activities. We present the planning and plan execution mechanisms and the underlying actions for human-aware navigation and joint manipulation that are needed in joint human-robot plans. We will show the specific challenges for these tasks and explain our current research activities. We will then discuss the relation of these components and their integration into complete cognitive systems.

2 Collaborative Planning

Joint action plans are necessary to coordinate the activities of a robot and a human to achieve a joint goal. The planning problem in this collaborative setting comes with two additional requirements compared to single-robot planning: 1) there are special kinds of actions involving contributions from the human and the robot and 2) the primary objective measure to evaluate plans are legibility and human comfort.

These requirements imply the need for new plan representation mechanisms. A plan cannot only be composed of robot actions with pre- and postconditions, but has to include the actions of humans. This problem has received only little attention so far. Joint intention theory [4] provides a basis for modeling the commitment of the cooperation partners to a joint tasks. The Human-Aware Task Planner [2] combines this theory with a classical planning approach. It uses hierarchical task planning to plan joint courses of action and evaluates the generated plans with respect to social acceptance.

A problem with classical planning approaches is the high flexibility that is necessary in collaboration scenarios. One possibility is the use of reactive planning [6], which means that plans are represented in a richer language than partially ordered action sequences and allows to specify contingencies, parallel task execution and failure handling [15]. Reactive plans can take into account several courses of action of the human and thus are valid for a wider range of human behavior.

Reactive plans can be generated and modified by transformational planning. The transformational planner we use [17] has reduced the execution time of a single-robot plan for setting the table by up to 20%. In this scheme a robot searches through the space of possible plans by replacing parts of a plan by other sub-plans and then testing these plans by projecting them. Transformation rules are specified by patterns, which are matched against a plan and when applied produce the new plan. Projection is a kind of generalized simulation, which allows the robot to assess the quality of a plan.

To use transformational planning for joint human-robot plans [10], the transformation rules must be complemented with rules to change the executing agent of an action in a plan and to replace a sequence of single-agent actions by collaborative actions. Another challenge is the projection of a generated plan. For single-robot scenarios, plans can be projected by using a physical robot simulation and evaluating the output. In the collaborative case, we would have to simulate human behavior. As there is no accurate simulation of humans

available, the projection must be performed with models of human behavior, abilities, habits 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, 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.

Examples for the first class of models are the most likely next action of a person, the space that will be occupied by a person in the near future or the inappropriateness of certain actions of the robot (to touch food, for instance). Some general knowledge about humans can be used implicitly by the choice of the action preferences of the robot. For example, we assume that in joint tasks it is appropriate to avoid unnecessary spatial conflicts. When choosing an action, our algorithm predicts possible conflicts by using the human-aware navigation planner described in the next section and selects an action with minimal spatial interference.

To support the second kind of models, a robot needs to monitor its environment constantly to detect useful experience for learning and updating models of the user. As a part of the transformational planning framework, we have developed the Robot Learning Language [9]. Using these language constructs, a learning problem and the necessary experience is defined by a declarative specification, which is then used to generate code that automatically recognizes relevant experience data and that can be called to execute a learning problem. This scheme allows a robot to repeat the learning process whenever new experience is available. It also makes it possible to adapt to changed requirements of a person. For example, the capabilities of an elderly person might change over time and a robot assistant must be able to adapt its plans accordingly.

3 Human-Aware Navigation

For carrying out joint action plans, the underlying capabilities of the robot also have to take into account the presence of the human in the task. In this section, we consider navigation tasks, in which a robot not only treats the human as a moving obstacle, but takes into account human comfort and social rules to maximize legibility of the overall task.

Although robot navigation has been subject to extensive research, there has been done only little work considering human comfort. Tadokoro et al. examine robot motion in the presence of humans [25] using a probabilistic prediction model of human motion. The




Fig. 1 (a) Path found by classical motion planning systems. (b) Human-aware plan found by HANP.

goals of this work was collision avoidance and efficiency, but social rules and comfort are not taken into account. In contrast, the social aspect of navigation was the focus of the work of Hansen et al. [8], who provide a solution for positioning a robot in relation to a human according to the situation classified by motion pattern analysis. Robots moving efficiently and acceptably in crowds of people have been studied by Müller et al. [18].

We are interested in robot navigation tasks to reach a certain goal position taking into account the safety and comfort of humans. For instance, Figure 1 (a) illustrates a path generated by a standard motion planner. It is uncomfortable, because the robot passes the person’s back at a very close distance.

A collaborative robot should also be able to determine where a given task should be achieved, how to place itself relative to a human, how to approach him, how to hand an object and how to move in a relatively constrained environment in the presence of humans. In this article, we focus on those decisions concerned with the navigation path in different situations.

The Human Aware Navigation Planner (HANP) [23] models the safety and comfort of humans with different kinds of cost functions, which can be adapted to specific situations and user needs. The most basic cost types to be considered in all cases are *safety*, which is modeled by a Gaussian cost function around the human and *hidden zones*, which assigns higher costs to areas behind large objects, which are not visible to a person. The hidden zones’ costs take into account the current distance between a human and the obstacle to raise costs when humans are near the other side of the obstacle.

Using the combination of all cost functions in a discretized representation of the world, HANP uses A* search to compute an optimal path in terms of human comfort.

This approach has been validated in different simulated domains with a varying number of humans and different pieces of furniture and obstacles. HANP is fast enough to replan and adapt its path along the execution. If a change occurs in the environment, like a change in human state, position, orientation or appearance of an obstacle, fast computation times allow online replanning and a smooth switch to the new path [23]. The quality of the found navigation paths is illustrated in Figure 1 (b), which shows an alternative path considering the cost functions for approaching humans.

3.1 Approaching Humans

When a robot and a human need to perform joint manipulation tasks, one challenge is how the robot should approach the human. We assume that the human to be approached is relatively static, either standing or sitting. The cost function for the safety grid is adapted to the human posture, assuming that a standing person will feel less threatened by an approaching robot than a sitting one. In addition to safety and hidden zones’ costs, we add *visibility* costs, which model the effort of the person to trace the robot. This means in particular, that a robot avoids to pass a person closely at the back, thus allowing the person to see the robot when it is approaching.

The definition of cost functions does not guarantee a human-friendly behavior unless the costs are calibrated with the actual comfort felt by humans. The data collected in two user studies was used to find appropriate cost parameters [5]. In these user studies, a robot approached a sitting person from different directions with different speeds with the purpose of handing over an object to the person.

3.2 Avoiding Humans

The method for approaching humans generally works well when humans in the environment are relatively static and there is enough room for the robot to find paths with low costs. Because HANP calculates paths very efficiently, it allows to recalculate paths constantly while the people are moving in the world. However, this constant reconsideration of the path may lead to robot behavior, where the robot constantly changes its intended course of action, which diminishes legibility. Moreover, the space in everyday environments like offices or households is very restricted, which means that an overcautious robot will end up waiting for a path to become free and never reaches its goal.

For the navigation tasks with moving humans, we lowered the safety costs to allow a robot to plan a potentially unsafe path along a position, which is currently occupied by a human. This measure assumes that the people are moving and might have cleared the way before the robot has reached their position. Besides, even if a person in the robot’s way did not have the intention to move, the approaching robot might initiate the human to step aside.

Another difference in the cost functions is that we replaced the visibility cost of the approaching case with a function that assigns a higher cost to the area in front of the human. At first, this seems a contradiction to the rationale for the visibility cost, because the robot will prefer paths behind the back of the person. But this cost function assumes a linear movement of the person and thus leads the robot to plan a path which makes way for the human to pass first.

The plans HANP generates with these modified costs are potentially unsafe. Therefore we use a controller that recognizes imminent conflict situations and stops the robot early enough to avoid collisions.

The navigation method for dynamic navigation tasks has been tested in a simulated kitchen domain [13]. The robot behavior was more legible in that the robot seemingly did not switch between paths (although the plan was recomputed constantly) and the intention of the robot to reach a certain goal position was better recognizable.

4 Joint Manipulation

Many assistive tasks require special actions in which both the human and the robot have to take part, for example jointly carrying an object, holding an object steady for the other agent to manipulate on it, lifting an elderly person from the bed or simply handing over an object from one agent to another.

In the following we first describe the link between human-aware navigation to approach a human and the manipulation part of indicating an intended handover. After that, we present the cognitive challenges of jointly carrying an object.

4.1 Initiation of Joint Manipulation Tasks

In Section 3.1 we have described the considerations for a robot to navigate towards a human with the goal of a joint action like handing over an object. In most robots, navigation and manipulation are distinct modules, which operate sequentially. This leads robots to approach a human, stand in front of her and then lift

the arm. In a realistic environment with uninstructed people, a person would probably not even wait for the robot to lift its arm, but consider it as an obstacle and proceed on her way. A robot must take care to make its intention legible for initiating a joint manipulation task.

A first step to implement joint manipulation and navigation is to understand how humans coordinate such movements. There have been several approaches to study the navigation behavior of humans and to use these findings for robot behavior [11,21], the combination of navigation and arm movement has not been studied so far.

We have investigated human behavior for preparing a handing over task with the goal of finding appropriate parameters for when and how a robot should start its manipulation actions [3]. Not surprisingly, the study found that the phases of approaching and the arm movement blend smoothly into each other. The exact parameters of when to lift the arm and at which speed varied among subjects, but the overall profile of the task execution was comparable.

The study also showed that the parameters of the handing over action seem to be independent of the receiving subject. In a second experiment, the receiving human was replaced by a table. Even though the goal position for “handing over” was different in the table case, the parameters for preparing this action were similar.

This behavior for handover tasks helps humans to predict the intentions of other people. In ongoing work, we implement the observed strategy of early task preparation on a robot. We will then investigate if this imitation of human behavior enhances legibility for the receiving human.

4.2 Physical Cooperation

Whereas in a handover task the agents jointly hold the object for a very short time, there are other forms of interaction, where an object is manipulated jointly over a longer time period. Figure 2 shows a scenario with a human and a robot jointly carrying a bulky object in a constrained environment.

In this kind of interaction, haptic feedback, i.e. the measurements from force sensors, can be used for controlling the robot’s movements. The direct coupling of the actions of humans and robots requires special attention regarding the stability of the coupled dynamical systems and human comfort and safety, and raises the demand for immediate compliant reaction under real-time constraints.

The classical approach for joint manipulation tasks is — similar to other aspects of human-robot interaction — that the robot is merely a passively following tool for the human. In this paradigm, the human leads the joint manipulation operation, while the robot serves as an additional load carrier, which reactively adapts its behavior based on force and position input from the human to share the overall task load [12, 28].

Similar to task planning for joint activities, this use of a robot as a pure tool limits the capabilities of the human-robot system. Combining abilities of both agents can make it possible to achieve goals that are out of reach with the human/master — robot/slave concept. For one thing, it has been shown mathematically that for two agents to navigate an object jointly in narrow space, two initiative taking partners are necessary to achieve the desired trajectory [14]. Also the different perception and manipulation capabilities of both agents and limited visibility are arguments to combine their cognitive capabilities.

A challenging aspect of joint physical actions is that the task of negotiating a feasible path and the action execution are performed concurrently. Consider as an example two humans jointly carrying a bulky object through a narrow passage with the person at the object’s front walking backwards and the partner at rear facing forwards. Without explicit prior negotiation of a joint plan or any other than haptic interaction during task execution, from our experience humans will be able to fulfill the task. Both partners combine their perceptual capabilities to produce a collision-free path, indicating proximity and divergent individual plans through interaction forces. The agents’ force contributions to the physical trajectory following task are dynamically shared, based on capabilities and user-preference policies. Task sharing strategies and their influential factors are currently under investigation.

The generation of joint plans for motion and dynamic interaction behavior faces similar challenges as the formulation of task plans: the human part of the action must be taken into account at planning time, but only the robot’s force input and its compliance is controllable. Thus, such plans must be flexible enough to account for unexpected behavior of the human partner.

The generation of motion plans for motion and dynamic interaction behavior faces similar challenges as the formulation of task plans: the human part of the action must be taken into account at planning time, but only the robot’s force input and its compliance is controllable. Thus, such plans must be flexible enough to account for unexpected behavior of the human partner.




Fig. 2 Joint manipulation task of a human and a robot using haptic interaction.

Our planning problem is defined by a set of motion primitives, which are inspired by observations in nature [19]. The parallel and sequential composition of such motion primitives is a motion plan. We use an efficient sampling-based hypothesize-and-test method to generate those plans. Because of its efficiency, the algorithm can be repeated constantly to ensure the necessary flexibility to react to human motions. An extension to planning over a set of sensori-motor primitives, i.e. including both partners’ reactive dynamic behaviors is envisaged.

As in navigation and task planning, a planner for joint motion needs an objective function to assess the comfort, legibility and respect for social rules in combination with standard criteria such as minimizing the applied force or reducing the duration of the task. Finding appropriate objective functions for human-robot joint physical actions is subject to ongoing research.

5 Discussion

Controlling a robot to collaborate closely with humans includes all levels of activities from task planning and plan execution to specific capabilities such as navigation and manipulation. The different aspects we have presented show that the challenges listed in Section 1.2 apply to all the levels of robot control. However, the solutions have to depend on the specific problem: flexibility on a task planning level can be achieved by reactive planning methods and needs sophisticated skills for assessing unexpected human behavior. On less abstract levels, the control loops are much tighter and the algorithms are efficient enough to allow continuous recalculation of values as long as this does not interfere with the legibility of the behavior.

Another observation from this article is that the different aspects of joint activities should not be regarded as completely separate problems: On the level of task planning and plan execution, the spatial considerations

of the environment play an important role (Section 2). This means that already on a task planning level the later navigation and manipulation behavior of the robot is important to allow legible actions that respect human comfort. Also navigation and manipulation should not be seen as completely separate. The decision of the navigation goal position is highly influenced by subsequent manipulation actions. And Section 4.1 has made the point that the intention to initiate a joint manipulation task can be communicated implicitly by a smooth transition between the navigation and manipulation action.

In this work we have focused on the action side of robot control for assistive technology. Except for the physical cooperation, where the readings of the force sensors can be used directly, the necessary perception to perform the actions is a complex issue of its own. For the work presented here, a robot must know the positions, orientations and to some extent also the hand movements of humans. For the reactive task planning, more abstract interpretations are needed to recognize the current action (like reaching, gripping or searching) [27] and predict subsequent actions.

These perception problems are subject to ongoing research, but are not yet at a reliable stage to use them as “plug-and-play” input components for the control methods presented in this paper. To make the development of these methods efficient and scientifically valid¹, we have designed specific evaluation frameworks or provide the perception by task-specific perception methods [23] or marker-based tracking systems.

For instance, to work on reactive task planning, we have extended an existing physical simulation of a kitchen robot with a user interface to control a robot in simulation. This interface allows the user to move the robot in all directions and to grip and put down objects (which are preprogrammed methods). This computer-game-like interaction makes it possible to place two agents in the world: one is controlled by a human operator, the other works as an autonomous robot (Figure 3). Using a physical simulation allows to reproduce real-world inaccuracies in plan execution for the robot, but the perception can be regulated from complete observability to noisy images. Evaluating plan-based control in simulation has the additional advantages of guaranteeing safety for human subjects and making the experiments repeatable (which is next to impossible on current completely integrated robots). The physical simulation in combination with a geometric simulation has also been used to develop and evaluate the navigation capabilities described in Section 3.2.




Fig. 3 Simulated collaboration of a human and a robot in a kitchen environment.

Although there are currently ways to develop action and perception for collaborative robots more or less independently, an important next step is to integrate both sides more closely. It cannot be expected that automatic perception will reach the accuracy of ground-truth data in simulation and the robot’s actions can support perception, for example by actively searching for objects. In this respect, completely integrated cognitive systems are an important basis for collaborating assistive robots.

6 Conclusion

Robots that are able to engage in joint activities with humans are a promising scheme for assistive technology, for example in the context of work places or elder-care. On the one hand, the combination of capabilities reduces the demand on dexterity and perceptual capabilities of the robot, because these are typical human skills. On the other hand, the need for cognitive interaction with humans is especially strong.

The challenges for implementing assistive robots are the same at all levels of planning and action execution: 1) How can a robot show legible behavior? It has to make its intentions clear to the human, but this process is mostly implicit and happens during the joint task execution. 2) How can not only the objective safety of humans be maximized, but also the subjective feeling of comfort? A crucial problem here is the modeling of human preferences, which comprises general models of human behavior and specific models for cultural conventions and personal habits. In Section 4.1 we have given an example of how neuroscience and psychology can provide valuable input for AI and we are convinced that such collaborations are an important step towards legible, safe robot behavior. 3) How to make robot behavior flexible enough to deal with human deviations from presumed fixed plans or action patterns? In particular, it is difficult to decide if a person’s behavior is compliant with the joint goal or if there is some kind of failure or the human has decided to abandon the joint goal.

¹ When executing imperfect action with imperfect perception, it is not possible to assign the success or failure to one of the subsystems.

In this article we have presented three specific aspects of joint activities from our current research activities: collaborative task planning, human-aware navigation and joint manipulation. We have briefly sketched the necessity to integrate all the levels of plan-based control as well as perception skills in a complete cognitive framework.

References

- Rachid Alami, Aurélie Clodic, Vincent Montreuil, Emrah Akin Sisbot, and Raja Chatila. Toward human-aware robot task planning. In *AAAI Spring Symposium*, 2006.
- Samir Alili, Rachid Alami, and Vincent Montreuil. A task planner for an autonomous social robot. In *Distributed Autonomous Robotic Systems 8*. Springer Verlag, 2009.
- P. Basili, M. Huber, T. Brandt, S. Hirche, and S. Glasauer. Investigating human-human approach and hand-over. *Human Centered Robot Systems: Cognition, Interaction, Technology*, pages 151–160, 2009.
- P. Cohen and H. Levesque. Teamwork. *Nous, Special Issue on Cognitive Science and AI*, 25(4):487–512, 1991.
- K. Dautenhahn, M. Walters, S. Woods, K. L. Koay, C. L. Nehaniv, E. A. Sisbot, R. Alami, and T. Siméon. How may i serve you?: a robot companion approaching a seated person in a helping context. In *ACM SIGCHI/SIGART International Conference on Human-Robot Interaction, HRI*, pages 172–179, Utah, USA, March 2006.
- J. Firby. An investigation into reactive planning in complex domains. In *Proceedings of the Sixth National Conference on Artificial Intelligence*, pages 202–206, Seattle, WA, 1987.
- Matthias Goebel, Matthias Althoff, Martin Buss, Georg Färber, Falk Hecker, Bernd Heissing, Sven Kraus, Robert Nagel, Fernando Puente Leon, Florian Rattei, Martin Russ, Michael Schweitzer, Michael Thuy, Cheng Wang, and Hans-Joachim Wünsche. Design and capabilities of the munich cognitive automobile. In *In Proc. IEEE Intelligent Vehicles Symposium*, 2008.
- Søren Tranberg Hansen, Mikael Svenstrup, Hans Jørgen Andersen, and Thomas Bak. Adaptive human aware navigation based on motion pattern analysis. In *Robot and Human Interactive Communication, 2009.*, Toyama, Japan, Sept.-Oct. 2009.
- Alexandra Kirsch. Robot learning language — integrating programming and learning for cognitive systems. *Robotics and Autonomous Systems Journal*, 57(9):943–954, 2009.
- Alexandra Kirsch, Thibault Kruse, and Lorenz Mösenlechner. An integrated planning and learning framework for human-robot interaction. In *4th Workshop on Planning and Plan Execution for Real-World Systems (held in conjunction with ICAPS 09)*, 2009.
- K.L. Koay, E.A. Sisbot, D.S. Syrdal, M.L. Walters, K. Dautenhahn, and R. Alami. 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*, 2007.
- K. Kosuge, H. Yoshida, and T. Fukuda. Dynamic control for robot-human collaboration. In *2nd IEEE International Workshop on Robot and Human Communication*, pages 398–401, 1993.
- Thibault Kruse, Alexandra Kirsch, E. Akin Sisbot, and Rachid 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. accepted for publication.
- M. Lawitzky, A. Mörtl, and S. Hirche. Human-robot cooperative task sharing. In *19th IEEE International Symposium in Robot and Human Interactive Communication (RO-MAN)*, 2010. (accepted for publication).
- Drew McDermott. A reactive plan language. Technical report, Yale University, Computer Science Dept., 1993.
- Michael Montemerlo, Jan Becker, Suhrid Bhat, Hendrik Dahlkamp, Dmitri Dolgov, Scott Ettinger, Dirk Haehnel, Tim Hilden, Gabe Hoffmann, Burkhard Huhnke, Doug Johnston, Stefan Klumpp, Dirk Langer, Anthony Levandowski, Jesse Levinson, Julien Marcil, David Orenstein, Johannes Paefgen, Isaac Penny, Anna Petrovskaya, Mike Pflueger, Ganymed Stanek, David Stavens, Antone Vogt, and Sebastian Thrun. Junior: The stanford entry in the urban challenge. *Journal of Field Robotics*, 25(9):569–597, 2008.
- Armin Müller, Alexandra Kirsch, and Michael Beetz. Transformational planning for everyday activity. In *Proceedings of the 17th International Conference on Automated Planning and Scheduling (ICAPS'07)*, pages 248–255, Providence, USA, September 2007.
- Jörg Müller, Cyrill Stachniss, Kai O. Arras, and Wolfram Burgard. Socially inspired motion planning for mobile robots in populated environments. In *International Conference on Cognitive Systems (CogSys'08)*, Karlsruhe, Germany, 2008.
- F. Nori and R. Frezza. Nonlinear control by a finite set of motion primitives. In *Proceedings of the 6th IFAC Symposium on Nonlinear Control Systems*, 2004.
- Stephanie Rosenthal, Joydeep Biswas, and Manuela Veloso. An effective personal mobile robot agent through a symbiotic human-robot interaction. In *Proceedings of the Ninth International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2010.
- Satake S., Kanda T., Glas D. F., Imai M., Ishiguro H, and Hagita N. How to approach humans? strategies for social robots to initiate interaction. In *ACM/IEEE International Conference on Human-Robot Interaction*, 2009.
- 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.
- 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.
- Siddhartha Srinivasa, David Ferguson, Casey Helfrich, Dmitry Berenson, Alvaro Collet Romea, Rosen Diankov, Garratt Gallagher, Geoffrey Hollinger, James Kuffner, and J Michael Vandeweghe. Herb: a home exploring robotic butler. *Autonomous Robots*, 28(1):5–20, 2010.
- S. Tadokoro, M. Hayashi, Y. Manabe, Y. Nakami, and T. Takamori. On motion planning of mobile robots which co-exist and cooperate with human. *IROS, IEEE/RSJ*, 2:2518, 1995.
- Yeow Kee Tan, Dilip Limbu Kumar, Ridong Jiang, Liyuan Li, Kah Eng Hoe, Xinguo Yu, Li Dong, Chern Yuen Wong, and Haizhou Li. An interactive robot butler. In J.A. Jacko, editor, *Human-Computer Interaction. Novel Interaction Methods and Techniques*, pages 385–394. 2009.
- Moritz Tenorth, Jan Bandouch, and Michael Beetz. The TUM Kitchen Data Set of Everyday Manipulation Activities for Motion Tracking and Action Recognition. In *IEEE Int. Workshop on Tracking Humans for the Evaluation of their Motion in Image Sequences (THEMIS). In conjunction with ICCV2009*, 2009.
- T. Wojtara, M. Uchihara, H. Murayama, S. Shimoda, S. Sakai, H. Fujimoto, and H. Kimura. Human-robot collaboration in precise positioning of a three-dimensional object. *Automatica*, 45(2):333–342, 2009.