

TOWARDS INTUITIVE HUMAN-ROBOT COOPERATION

Andreas J. Schmid, Heinz Wörn

Oliver C. Schrempf, Uwe D. Hanebeck

Institute for Process Control and Robotics
Universität Karlsruhe (TH)
Karlsruhe, Germany

Intelligent Sensor-Actuator-Systems Laboratory
Institute of Computer Science and Engineering
Universität Karlsruhe (TH)
Karlsruhe, Germany

ABSTRACT

Human-robot cooperation calls for the treatment of human-machine communication channels, especially if humanoid robots are involved. In this paper, we consider implicit non-verbal channels given by recognizing the partner's intention and proactive execution of tasks. We propose a method that keeps the human in the loop and allows for the systematic reduction of uncertainty inherent in implicit cooperation. We present a benchmark scenario as well as preliminary implementation results.

1. INTRODUCTION

Human centered robotics is a hot topic in research today. This is indicated by the large number of humanoid robotics projects world-wide like the Collaborative Research Center SFB 588 on "Humanoid Robots" of the German Research Foundation [1].

A key challenge within this area of research is to build robots that offer an intuitive and human-like way of interaction. Since such a task involves the application of sensors and actuators in a real world scenario dealing with humans, the development of highly sophisticated approaches is required.

To provide cooperation schemes that are appealing for humans, one needs to understand where the strength in humans cooperation lies. It is known that humans are very good at mutual control of their interaction. This is achieved by reading and interpreting the affective and social cues of their cooperation partners [2]. Hence, a robot system that is able to read the user's (non-)verbal cues in order to infer the user's intention is able to interact more intuitively from a human's perspective.

While figuring out their interaction partner's goals or desires, humans try to trigger reactions. An example for this is the waiter on a cocktail party who wants to know if somebody wants a refill. He presents the bottle, causing the guest to present his glass or to withdraw it. We call this action of

the waiter *proactive*, since he acts without an explicit command from the guest, by provoking a clarifying reaction from the guest and thus removing any uncertainty about the user's intention. As this example illustrates, humans are accustomed to performing intuitive cooperation. Thus, providing service robots with such a skill opens a new dimension in human-robot cooperation.

The crucial point of intuitive cooperation is the robot's ability to recognize the user's intention. Since even humans cannot do this perfectly, a probabilistic approach has to be used in order to describe the uncertainty involved when recognizing intentions. Proactive behavior of the robot can then be used to minimize this uncertainty. Hence, the challenge for the planner is to select an appropriate robot action in order to urge the user to react in a way that unravels his intention. It is obvious, that the corresponding robot actions need to be executed with care, since the recognized intention is uncertain. The human user is meant to close the loop of intention recognition and proactive action planning.

We proposed a theoretical framework for a system architecture in [3] that allows for intuitive human-robot cooperation in the sense of avoiding explicit clarification dialogs and explicit commands. In this work, we will furthermore consider implementation issues.

Our approach involves intention recognition, a discipline that is closely related to classical plan recognition. As we want to infer hidden user intentions, we are especially interested in the so called keyhole plan recognition as mentioned in [4]. In this approach no explicit help of the user for the inference process is expected. A popular approach in this field is the application of graphical models like Bayesian Networks. They provide a mathematical theory for reasoning under uncertainty and causal modeling. An example for the application of Bayesian Networks is the *Lumière* project [5] that tries to figure out the user's goals in office computer applications from tracking their inputs. A similar approach was successfully applied to affective state detection [6]. We have already presented a Bayesian network model for intention recognition in [7]. In this paper we will address this problem in a more traditional system theoretic manner, in order to

This work was supported in part by the German Research Foundation (DFG) within the Collaborative Research Center SFB 588 on "Humanoid robots — learning and cooperating multimodal robots".

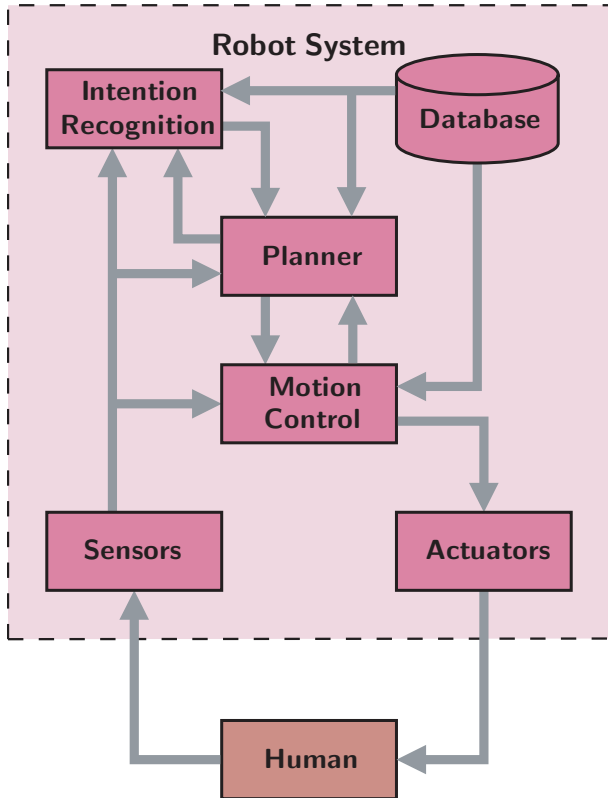


Fig. 1. Robot control architecture.

show the genericity of our approach.

The other vital concept in our approach is proactive execution. Although many applications of proactive behavior are located within the realm of business and finances, there have been attempts to apply it to robotics. Proactive planning is mentioned in [8] in the case of probabilistic determination of the results of an action of a mobile robot. An architecture for autonomous agents that include a proactive behavior component is outlined in [9]. Achieving proactive behavior of agents through goal reprioritization is suggested in [10]. The unified planning and execution framework IDEA [11] allows agents to use the concept of proactive planner invocation in case the agents anticipate any problems.

Fig. 1 depicts our robot control architecture that includes an *Intention Recognition* module. This module fuses the information that is available from the *Sensors* and the *Database* using probabilistic methods. The result is fed to the *Planner*. In case the information about the human intention is too uncertain, the *Planner* is forced to execute tasks proactively.

The remainder of the paper is structured as follows. The probabilistic approach to intention recognition is explained in section 2.1, and section 2.2 illustrates proactive execution. We propose a benchmark scenario in section 3.1. Section 3.2 presents details about the implementation and our experiments, and we conclude the paper in section 4.

2. PROPOSED METHOD

2.1. Intention Recognition

Assisting a user based on implicit communication requires the knowledge of the user’s aims, goals, or wishes. We summarize these as the user’s *intention*. Since intention is a state of mind, it cannot be measured directly. Nevertheless, humans are able to recognize intentions of their communication partners. This skill is extremely important, especially in non-verbal communications. Even though the estimation of a partner’s intention is usually uncertain, the gained information is still of great value. Hence, we need a model that allows for estimating the user’s intention from external clues while maintaining information concerning the uncertainty of the estimate.

The key to the hidden state of the user’s intention are the actions performed by the user. It can be assumed, that actions are directly caused by the intention, as long as the user is not trying to cheat. Hidden intentions drive the observable actions, thus, the model must describe how the actions depend on the intention. We call this a forward model, since it captures the causal dependencies — actions depend on intentions, not vice versa.

To estimate the user’s intention, we propose a dynamic stochastic model which is shown in Fig. 2.

User-intentions are often influenced by external circumstances. In other words, the intention is affected by the environment the user acts in. We cover these environmental influences by a random variable d containing “domain knowledge”.

The user’s intention is a hidden state that cannot be observed directly. Hence, we model the intention by means of a random variable i . For our application we assume this variable to be discrete since there are distinct intentions that we want to distinguish. Nevertheless, it is possible to define continuous intentions.

A user performs actions, modeled by the random variables a_i . These actions are depending on the intention. As already pointed out, the actions depend causally on the intention and *not* vice versa. We cover this fact by the application of a probabilistic forward model $f(a_{it}|i_t)$ for every known a_i . Due to the power of probabilistic reasoning we are able to infer the intention from observations on actions performed by the human. The vector representation of actions ($\underline{a} = [a_1, \dots, a_n]^T$) in Fig. 2 was chosen just for convenience. Since the actions a_i are independent they could be modeled in multiple separate blocks.

The robot has to recognize the performed actions from sensor measurements. Hence, we need an additional layer (measurement variables) in our model. Here we can apply standard measurement models known from dynamic systems theory.

To represent temporal dynamic behavior of a user, we introduce connections from time-step t of the intention variable




Fig. 2. Block diagram of intention forward model.

to time-step $t + 1$, in order to build a dynamic model. This enables us to cope with a user "changing his mind".

Actions may depend on the actions performed in the preceding time-step. Hence, a connection from every action to its corresponding variable in the next step is introduced. This allows for reasoning how likely it is, that the same action is performed twice, given a certain intention.

Since sensor measurements depend only on their originating action in the current time step and not on previous measurements, there are no connections from a measurement in time step t to the corresponding measurement in time step $t + 1$.

We now describe the estimator for the intention i_t given the measurement vector \hat{m}_t and the domain knowledge \hat{d}_t as depicted in Fig.3. As a result it computes a probability density over the intention i_t . The BF- and BB-blocks depict a

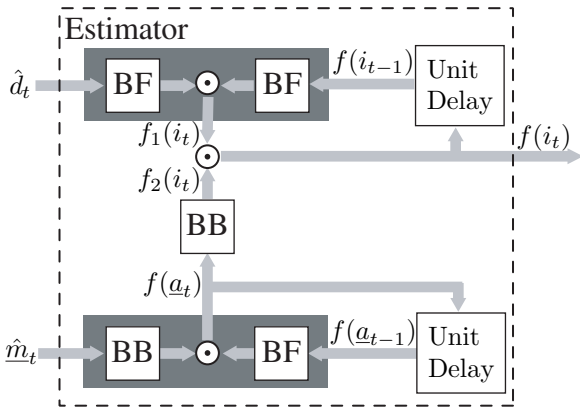


Fig. 3. The estimator computes a probability density over the intention i_t based on the current domain knowledge \hat{d}_t and the measurements \hat{m}_t via intermediate densities $f(a_t)$, $f_1(i_t)$, and $f_2(i_t)$. It consists of Bayesian forward (BF) and Bayesian backward (BB) inference blocks.

Bayesian forward and Bayesian backward inference respectively. In this way the density $f(i_t)$ is calculated via intermediate densities $f(a_t)$, $f_1(i_t)$, and $f_2(i_t)$.

The intermediate densities are multiplied, which is indicated by the dot in the circle. The dark blocks indicate the fusion of information from time-step t with information from time-step $t - 1$. This is to emphasize the fact that prediction- and filter-step are processed simultaneously.

A more in depth introduction to our approach to intention recognition can be found in [7].

2.2. Proactive Execution of Tasks

Our concept of intuitive interaction between a robot and a human involves the tight interaction of the intention recognition and the proactive execution module. The intention recognition provides the proactive execution with an estimated probability density over the possible intentions. In the other direction the proactive execution returns the task that is currently being executed, which serves as an input for the calculation of the conditional intention probabilities in the intention recognition module.

When the robot is supposed to act in response to the intention of a human user, the planner takes all known and available tasks into account, any explicit action requests through a user interface and the input from the intention recognition. With respect to the intention recognition we have to distinguish several cases:

- The first case is that no intention can be inferred from the available sensor data. As a consequence, the probabilities of all intentions are equal, and no preferred intention can be determined.
- Another case with similar symptoms arises when there are many intentions that seem to be equally likely according to the observations. Again, the probabilities of those intentions will all have similar values, and it is again not possible to choose a clear winner unless there is another intention that has a higher probability.
- Assuming that there are two or three candidates as likely estimates for the human intention we have the chance to make a guess about the "true" intention. In this third case of ambiguous results from the intention recognition, we can choose an appropriate action and monitor the development of the probability density over all intentions.
- The last case happens when there is indeed one single intention that obviously dominates the rest. This is the ideal case, as it gives the planner a clear idea of what task to execute.

The last case is the easiest case to handle. The planner chooses the appropriate task and the robot thus acts according to the recognized intention. The other cases are a lot harder to deal with. In the cases where no intention was recognized with a sufficient certainty, the planner selects either an idle task or a task that tries to capture the human user’s attention and communicate that the robot is idling and waiting for a command.

For the third case of two or three plausible intentions to choose from, we developed the concept of the *proactive execution* of a task. This means that instead of idling, we pick an intention according to a measure of optimality and pretend that this is the wanted intention, and select an appropriate task. Subsequently we start executing this task, closely monitoring how the values from the intention recognition develop. In case the similar probabilities tip in favor of our chosen intention we keep executing the task as usual. On the other hand, if it becomes clear that this task does *not* match the human’s intention we stop execution, maybe roll back some movements, and start all over. Should there be no significant change of the confidence in these intentions we just keep executing the task.

The challenge here is the optimal selection of an intention from the two or three candidates. A practical strategy is to select the intention that triggers the execution of a task that lends itself to a segmentation into several parts naturally. This is true for most tasks that are specified by a finite state machine consisting of more than two states. Another strategy takes the issue of human safety into account and therefore the intention that triggers a robot action that is deemed the safest of all possible activities.

The strategy we propose here, however, is to pick the intention whose corresponding robot action will maximally decrease the uncertainty we have about the correct intention. If we denote the random variable for the intentions with I , we can specify this uncertainty as the entropy

$$H(I) = - \sum_j p(i_j) \lg p(i_j) .$$

Let the random variable for the actions be A , then after picking an action the uncertainty of our system is reduced to the conditional entropy $H(I|A)$. We calculate $H(I|A)$ as

$$H(I|A) = - \sum_i p(a_i) \sum_j p(i_j|a_i) \lg p(i_j|a_i) .$$

Using Bayes’ rule we can express the unknown $p(i_j|a_i)$ with the known $p(a_i|i_j)$ and thus obtain

$$\begin{aligned} H(I|A) &= - \sum_i p(a_i) \sum_j \frac{p(a_i|i_j) p(i_j)}{p(a_i)} \lg \frac{p(a_i|i_j) p(i_j)}{p(a_i)} \\ &= - \sum_i \sum_j p(a_i|i_j) p(i_j) \lg \frac{p(a_i|i_j) p(i_j)}{p(a_i)} . \end{aligned}$$

By computing this value for all possible actions and comparing the results, we are able to determine the action \tilde{a} that has the lowest conditional entropy value and thus leaves us with the least uncertainty, that is

$$\tilde{a} = \arg_A \min H(I|A) .$$

Example: Consider the following probability values for 3 intentions i_j : $p(i_j) = \{0.4, 0.3, 0.3\}$ and 2 possible actions a_i . The selection of the action is done according to table 1.

Table 1. Action selection depending on intentions $\equiv p(a_i|i_j)$

	i_1	i_2	i_3
a_1	1	0	0
a_2	0	1	1

The entropy of the intentions I is $H(I) = 1.571$. Plugging in our values of i_j and table 1 and using $p(a_i|i_j) = 0$ when $p(a_i) = 0$, we obtain $H(I|A) = 0.529$ when choosing action a_1 (i.e., $p(a_i) = \{1, 0\}$), and $H(I|A) = 1.042$ when choosing action a_2 (i.e., $p(a_i) = \{0, 1\}$). Hence we would pick action $\tilde{a} = a_1$ in this situation because it leaves us with the least uncertainty.

3. EXPERIMENTAL RESULTS AND DISCUSSION

3.1. Benchmark Scenario

As a benchmark that can be used to effectively demonstrate and evaluate the proactive execution of tasks, we propose a scenario that involves two competing intentions and corresponding actions. Fig. 4 shows the state machine that describes this scenario.

It starts out with a dialog between robot and human where the human asks the robot to fetch a pot. The robot then navigates to the pot, grasps it, and comes back to the human. Now the intention recognition comes into play. The human is holding a tray in one hand and a cup in another. By presenting the cup to the robot, the latter should interpret this implicit communication as the human’s intention of having himself a cup poured. As a consequence, the robot should fill the cup. If the human moves the tray forward, the robot should recognize that it is asked to place the pot on the tray and release its grip.

When the user indicates neither desire, the intention recognition should realize this and present similar probability values for both intentions. The planner then switches to proactive execution, and the following three steps are performed in a loop: First, the planner selects a task to execute tentatively. Then the robot starts or continues to execute the given task. Lastly, after some short interval, the planner revisits the inputs it receives from the intention recognition and checks if the currently selected intention is still supported by the sensory evidence. After that the next loop iteration begins.

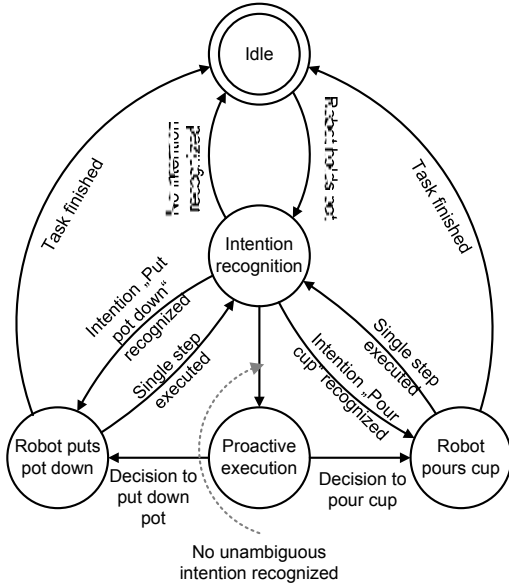


Fig. 4. Finite state machine for the demonstration of the proactive approach.

Upon successful completion of one of the tasks the robot should go back to the idle state.

3.2. Implementation and Experiments

As a software framework for the entire robot control architecture we use the Modular Controller Architecture (MCA) that runs both under Linux and Windows [12]. Its predominant paradigm is that of a controller with sensory and control I/O. Each control *module* receives sensory input from lower-level modules and returns control output to them. It passes on sensory output to higher-level modules (possibly after some processing) and receives in turn control input. Modules are connected via these I/O-edges, and multiple modules can be grouped together. One or more modules and groups are then collected in an executable *part*. Parts communicate through TCP/IP across platforms if necessary. An administration tool that visualizes the module hierarchy and a GUI tool for easy user control complete the framework.

You can see a picture of our current evaluation platform in Fig. 5. Eventually it will run on the demonstrator robot of our collaborative research center shown in Fig. 6.

We have implemented the intention recognition software as a C++ library. The library functions are called from within an MCA module that is connected with all available perception modules for input values, and the proactive execution module that receives the output as a probability density vector over the human intentions.

The proactive execution module first filters and then evaluates that output as described in section 2.2. Also, we enhanced the module’s decision capability by a variety of addi-




Fig. 5. Our current robotic evaluation platform.

tional methods that are beyond the scope of this paper. The result is a suggestion to the central planner of the robot control architecture as to what task to execute next.

The present implementation of our intuitive human-robot interaction software is tailored to perform the benchmark described in the previous section 3.1. We are planning to broaden the range of scenarios by increasing the number of recognized intentions and the complexity of processed sensor information.

4. CONCLUSIONS AND FUTURE WORK

We have presented an intuitive human-robot cooperation system. It allows for the recognition of the human’s intention and the planning of corresponding robot actions. The system’s probabilistic nature is designed to deal with the inherent uncertainty of implicit non-verbal communication and perception. As a result, we are able to close the loop involving human and robot by sensing the human’s intentions and feeding back the findings through the robot’s actions at any time and at any level of certainty.

The two modules we use to realize our concept are the intention recognition module and the proactive execution module. The former facilitates communication between human and robot on an intuitive level, using affective and social cues rather than explicit commands. The latter selects the tasks to be executed according to the intentions the former has determined. This is straightforward for unambiguous results of the intention recognition. In the difficult case of high uncertainty we opt for the proactive execution of tasks rather than idling. Thus we display our information of the human’s intentions back to him and provoke his reactions that we use in turn to confirm or disconfirm our choice for the correct intention. This intention is chosen such that we maximize the information we can obtain from the user’s reaction and at the same time minimize our system’s uncertainty.



Fig. 6. The demonstrator robot of our collaborative research center.

We have already implemented first versions of both the intention recognition module and the proactive execution module within our robot control framework software MCA2, and they are working very well. Progress in the implementation and experiments with humans are beyond the scope of this paper and therefore subject to future publications.

5. REFERENCES

- [1] Regine Becher, Peter Steinhaus, and Ruediger Dillmann, "The Collaborative Research Center 588: Humanoid Robots - Learning and Cooperating Multimodal Robots," in *Proceedings of Humanoids 2003, Karlsruhe, Germany*, 2003.
- [2] Cynthia Breazeal, "Robots in Society: Friend or Appliance?," in *Agents99 Workshop on Emotion-based Agent Architecture*, Seattle, WA, 1999, pp. 18–26.
- [3] Andreas J. Schmid, Oliver C. Schrempf, Uwe D. Hanebeck, and Heinz Woern, "A Novel Approach to Proactive Human-Robot Cooperation," in *Proceedings of the 14th IEEE International Workshop on Robot and Human Interactive Communication (RO-MAN 2005)*, Nashville, USA, August 2005.
- [4] David W. Albrecht, Ingrid Zukerman, and Anne E. Nicholson, "Bayesian Models for Keyhole Plan Recognition in an Adventure Game," *User Modeling and User-Adapted Interaction*, vol. 8, no. 1-2, pp. 5–47, 1998.
- [5] Eric Horvitz, Jack Breese, David Heckerman, David Hovel, and Koos Rommelse, "The Lumière Project: Bayesian User Modelling for Inferring the Goals and Needs of Software Users," in *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*. AUAI, 1998, pp. 256–265, Morgan Kaufman.
- [6] Xiangyang Li and Qiang Ji, "Active Affective State Detection and User Assistance With Dynamic Bayesian Networks," *Transactions on Systems, Man, and Cybernetics – Part A: Systems and Humans*, vol. 35, no. 1, pp. 93–105, January 2005.
- [7] Oliver C. Schrempf and Uwe D. Hanebeck, "A Generic Model for Estimating User-Intentions in Human-Robot Cooperation," in *Proceedings of the 2nd International Conference on Informatics in Control, Automation and Robotics, ICINCO 05*, 2005.
- [8] Jun Miura and Yoshiaki Shirai, "Parallel Scheduling of Planning and Action of a Mobile Robot Based on Planning-Action Consistency," in *Proceedings of the IJCAI-99 Workshop on Scheduling meet Real-time Monitoring in a Dynamic and Uncertain World; Stockholm, Sweden*, 1999.
- [9] Giuliano Armano, Giancarlo Cherchi, and Eloisa Vargiu, "An Agent Architecture for Planning in a Dynamic Environment," in *Proceedings of the 7th Congress of the Italian Association for Artificial Intelligence on Advances in Artificial Intelligence (AI*IA 2001)*, 2001.
- [10] James P. Gunderson, "Adaptive Goal Prioritization by Agents in Dynamic Environments," in *Proceedings of the 2000 IEEE International Conference on Systems, Man, and Cybernetics*, 2000.
- [11] Nicola Muscettola, Gregory A. Dorais, Chuck Fry, Richard Levinson, and Christian Plaunt, "IDEA: Planning at the Core of Autonomous Reactive Agents," in *Proceedings of the Workshops at the AIPS-2002 Conference, Toulouse*, 2002.
- [12] Kay-Ulrich Scholl, Jan Albiez, and Bernd Gassmann, "MCA - An Expandable Modular Controller Architecture," *3rd Real-Time Linux Workshop*, 2001.