Proceedings of the 17th IEEE International Symposium on Robot and Human Interactive Communication, Technische Universität München, Munich, Germany, August 1-3, 2008

“Good Robot”, “Bad Robot” –Analyzing Users’ Feedback in a Human-Robot Teaching Task

Anja Austermann, Seiji Yamada

***Abstract*—This paper describes an experimental study in which we analyze how users give multimodal positive and negative feedback by speech, gesture and touch when teaching easy game-tasks to a pet robot. The tasks are designed to allow the robot to freely explore and provoke human reward behavior. By choosing game-based tasks, we ensure that the training can be carried out without stressing or boring the user. This way, we can observe natural, situated reward behavior.**

I. Introduction

When teaching a robot by natural interaction one important task for the human teacher is to evaluate the robot’s performance and give positive and negative reward. This is especially true in case of reinforcement learning, where positive and negative reward is the only means for the user to guide the behavior of the robot, but feedback is also used frequently in general teaching tasks.

In this paper we describe a recently conducted experimental study which aimed at finding out, how users employ different modalities to give positive and negative feedback to a robot. We analyzed how speech, touch and gestures were useed for giving reward. We also assessed the variability of positive and negative feedback between users and between different tasks that the robot performs. By analyzing and comparing users’ feedback in different training tasks, we investigated which features of a task encourage or hinder the user to give feedback in a natural way.

The long-term-goal of our research, which motivated this study, is to develop a method for learning commands as well as positive and negative rewards that can be used for controlling the service and entertainment functions of service-robots and pet-robots through the interaction with a user. We chose the learning of positive and negative rewards as the starting point of our work, because they are the smallest useful set of commands, that can be used to teach a robot, for example by reinforcement learning. We are using a combination of Hidden Markov Models [11] and classical conditioning [3] for enabling the robot to learn its user’s preferred ways of giving reward and instruction. Details of the learning algorithm are given in [1] and [2].

Manuscript received May, 20. 2008. A. Austermann is with the Graduate University for Advanced Studies (SOKENDAI), Tokyo, 101-8430 Japan (e-mail:anja@nii.ac.jp)

S. Yamada is with the National Institute of Informatics and the Graduate University for Advanced Studies (SOKENDAI) Tokyo, 101-8430 Japan (e-mail: seiji@nii.ac.jp)

The work presented in this paper tries to answer three questions which are a basis for our research:

- *Is there a benefit in having a robot learn multimodal feedback from its user?* This is only true, if different people give reward in different ways which are hard to handle by using hard-coded feedback patterns.
- *Is it possible to learn multimodal user feedback in a training phase in a reasonable amount of time?* Only if reward behavior used by a single person does not vary excessively and is similar between different tasks, it can be learned effectively in a training phase
- *Which features of the training tasks are important for learning natural user feedback?*

In our study four different tasks are used to provoke positive and negative feedback from the user. The tasks are modeled to resemble simple games, suitable for young children, such as “Pairs” and “Connect Four”.

As the tasks are used for automatically recording and learning users’ feedback, they must be designed in a way that they allow the robot to provoke natural feedback from the user while not boring or stressing him or her. Moreover, the robot needs to be able to explore its user’s reward behavior autonomously without remote control. We decided to use an approach for provoking user’s feedback that is inspired by the Wizard-of-Oz principle, which is widely used in research on

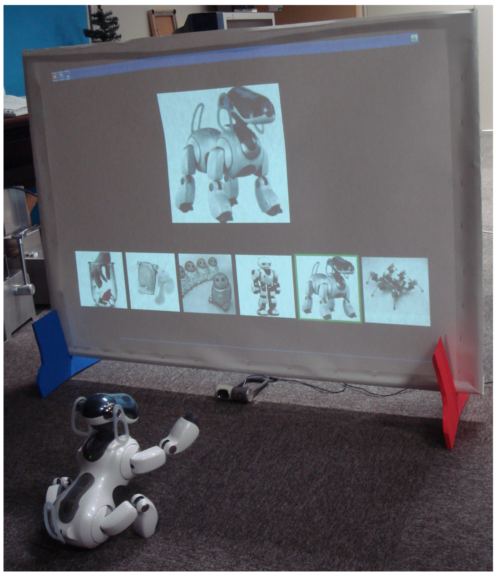

Figure 1: AIBO playing. The goal of the task, is to find the picture that corresponds to the sample above

Human-Robot-Interaction [4]. As in a Wizard-of-Oz scenario, we aim at giving the user the impression that the robot adequately reacts to his or her commands at a stage, where it actually does not understand them. However, different from a regular Wizard-of-Oz scenario, the training can be performed without remote controlling the robot.

Instead of that, the tasks and the behavior of the robot are designed with the objective to ensure that the user and the robot share the same understanding of the goal of the task as well as the way to reach it. This is done in two ways: The rules of the game task as well as an evaluation function to calculate the quality of all possible moves are hard-coded in the robot, so that it knows whether a move is good or bad in a certain situation. On the other hand, we chose tasks, which are easy enough to allow the user an instant and correct evaluation of the moves of the robot and give feedback accordingly. So the robot does not need to actually understand the user's commands to know how to correctly perform the task and can deliberately provoke feedback from the user by good or bad performance.

In our experiment, three out of four training tasks are "virtual" tasks. The robot plays on a computer-generated game board, which is projected from the back to a white screen, as can be seen in figure 1. All relevant information on the state of the training task can be accessed directly from the game server. This way, we can ensure that the task state is always interpreted correctly because the robot does not depend on potentially erroneous processing of sensor data. Further explanations on the training tasks are given in section III of this paper.

II. Related Work

Various studies have been conducted to investigate how users apply different modalities when interacting with a computer system. In [6] Oviatt gives an overview of some major results of recent research in multimodality and points out ten of the most common misconceptions about how people interact with multimodal voice- and gesture-based systems.

However, all quantitative analyses on multimodal user behavior, that we are aware of, have been conducted using GUI-based systems. Although several implementations of multimodal systems for Human-Robot-Interaction can be found in literature [5][7] little or no information exists on user preferences for natural, unrestricted multimodal interaction with robots. As embodiment needs to be considered as an important factor in Human-Robot-Interaction [10], it is unclear in how far results from the interaction with GUI-based systems can be transferred to Human-Robot-Interaction.

Several studies on how people teach robots or other artificial creatures, such as virtual characters through positive and negative reward have been conducted in recent years.

Ullerstam and Mizukawa [9] developed a method to teach complex behavior patterns to an AIBO pet robot based on reinforcement learning. In their system, reward is given through two predefined positive and negative utterances, as well as pressing the robots head or back sensor.

Thomaz et al. described an experimental setting for assessing human reward behavior and its contingency [8]. The participants of the study could give positive as well as negative reward to teach the virtual character Sophie to bake a cake in the "Sophie's World" scenario. Reward could be given by an interactive reward interface that allowed the user to assign any reward on a scale from -1 to +1 either to a certain object or to the world state. In their experiments they found a strong bias towards positive reward and discovered a phenomenon that they described as anticipatory rewards, positive rewards that were assigned to an object that the character has to use in a later step. This kind of reward can be interpreted as guidance for the character.

III. Training Tasks

In order to allow the robot to explore the way, its user gives reward and in order to record positive and negative feedback for learning and further analysis, we use specially designed training tasks. The main difficulty that has to be addressed during the training is that the robot needs to perform autonomously while not having any prior knowledge about its user and his or her way of interacting. This leads to a number of requirements that need to be met by the training task.

A. *Requirements for suitable training tasks*

Deliberately provoking positive and negative rewards from a user is only possible for the robot within a task where the human and the robot have the same understanding of which moves are desirable or undesirable. This allows the robot to anticipate the user's commands and explore his or her reward behavior by performing compliant or non-compliant actions.

In our experiments, two of the tasks allowed the user to give reward as well as instructions, while in the remaining two tasks, the users were asked to teach the robot by only giving positive and negative reward. Tasks that only allow positive and negative feedback from the user for training the robot can be designed relatively freely as long as their goal is known to both, the user and the robot. However, tasks that allow the user to give commands to the robot have to be designed carefully to ensure, that the user and the robot maintain the same understanding of whether an action of the robot is positive or negative. When users feel, their current level of instruction is not interpreted correctly by the robot they tend to give more fine-grained instructions. E.g. when the robot did not understand "Put the red stone to the yellow field", the participants of our preliminary experiments typically changed their instructions to "go forward… stop… pick up the red stone…". In tasks where the order of their subtasks can be changed, fine grained instruction can lead to a situation where the robot's actions do not comply with the user's instructions, although they lead to the correct goal state. Therefore, when allowing instructions from the user, tasks have to be designed in a way that knowing the goal state imposes an unambiguous order on their subtasks.

However, the user's judgment of the situation remains a potential source of errors, so the task must be designed in a way that the situation is easy to evaluate by the user. This is true for three out of our four tasks that we used in our experiments. In the "Pairs" game and the "Same Image" game the user only had to evaluate whether the two cards, chosen by the robot, show the same picture or not. In the "Dog Training" game, the users had to evaluate, whether the motion of the robot corresponded to the command given. In case of the "connect four" game, evaluation is more difficult for the user and there are possible moves that are not clearly good or bad. This allowed us to assess changes in the participants' behavior depending on whether they are confident in the feedback, they provide, or not.

B. *Using "virtual" tasks*

An issue the impact of which became obvious in the preparation and during the execution of preliminary experiments is the very limited ability of the AIBO robot to physically manipulate its environment and to move precisely.

In our preliminary study [1], this issue was addressed by equipping the robot with a shovel to facilitate moving objects around. However, it is still possible that the robot does not detect detecting errors during task-execution, such as failing to pick up the correct object. This poses a risk for misinterpreting the current status of the task. Therefore we decided to implement the training task in a way that the robot can complete it without having to directly manipulate its environment. When using a computer-based task, the current situation of the robot can be assessed instantly and correctly by the software at any time. During the experiments, the image of the playfield is generated by a computer and projected from the back to the physical playfield, as seen in figure 1. The robot visualizes its moves by motion and sounds and reacts to the moves of its computer opponent by looking at the appropriate positions on the playfield.

By reducing the risk that either the robot or the user misinterprets the situation, we ensure that the robot is able to anticipate the user's next reward or instruction correctly. This is necessary for associating the observed behavior correctly with a positive or negative reward.

IV. ASSUMPTIONS

Our further work on learning multimodal reward patterns relies on various assumptions which were investigated upon in the experiments, described in section V of this paper.

We assume that patterns of interaction between humans and robots range from rather universal ones, like pointing gestures, which are roughly the same between different individuals to highly individual patterns, like giving positive/negative reward.

Patterns that are universal can be pre-trained and adapted to a certain user during task execution. Only patterns that vary substantially between users need to be trained in a training phase that precedes the actual use of the robot. In our work we focus on the latter kind of interaction patterns.

We further assume that each user has a limited inventory of interaction patterns to express a certain command or reward. The interaction patterns, that are typically used, can change slowly over time. Moreover, we assume that interaction patterns used by one user for the same instruction do not vary excessively between different tasks. Otherwise, it would not be possible to learn these patterns within a training task in a reasonable amount of time.

V. EXPERIMENTS

We conducted two experimental studies to analyze human reward behavior during the interaction with a pet robot. Our aim was to get a better understanding of how people naturally employ different kinds of feedback and different modalities to give reward to a robot, which modalities are most important and how the reward given, varies for different persons and throughout different tasks.

In a preliminary study, which is described in more detail in [1] and [2] the participants were asked to instruct the robot to move colored objects to their target places and give reward to the robot for correct and incorrect actions. We imposed different restrictions on the reward behaviors that the participants were allowed to use in order to find out, how restrictions in allowed reward modalities affect the frequency of reward given by the users. We confirmed that the participants reacted quite sensitive to restrictions in the allowed way of giving feedback and that they tended to unintentionally give additional feedback that did not conform to the given restrictions. We also recognized some issues, such as the low ratio of actual interaction time to the time that the user was waiting for the robot to do something due to the robot's slow movements and the not very enjoyable nature of the service-task used in our preliminary experiments.

We used these findings from the preliminary study to improve the design of the training task for our second study. The participants interacted with an AIBO pet robot in four different game-like tasks. Three of which were "virtual" tasks as described in section III.B.

Audio and stereo video data as well as information perceived through the sensors of the robot were recorded. After the experiment and after each of the subtasks, we asked the participants to answer a questionnaire about their experience.

A. *Experimental setting and instruction*

Ten persons participated in our study. All of them were Japanese graduate students or employees at the National Institute of Informatics in Tokyo. Five of them were females, five males. Their ages ranged from 23 to 47. All participants have experience in using computers. Two of them have previous experience in interacting with entertainment robots. Interaction with the robot was done in Japanese. During the experiment, we recorded roughly 5.5 hours of audio and video data containing 2409 reward instances.

Based on this data we analyzed the usage of the different modalities speech, gesture and touch for giving feedback to a robot. Moreover, we investigated on how the feedback, given to a robot differs between different people and how it varies for an individual person during one task and when training the robot in different tasks.

1) Experimental setting

After analyzing the videos and the questionnaire responses of the preliminary study, the task for the next series of experiments was adapted according to our findings. The main reasons for the modifications were to gather more user feedback in a short time, to design the training task in a way that it does not resemble as much to a simple service task as the one used for the preliminary study, and to restrict the participants' possible movements and postures in a non-intrusive way. Moreover, the task should be feasible for implementation without the need to remote control the robot. We selected training tasks based on easy games.

In the experiments the participants were asked to teach the robot, how to correctly play these games by giving instructions and positive/negative reward for the robot's moves. They were instructed to interact with the robot freely by gesture, voice and by touching the robots touch sensors. They were told that the robot learns to play the different games through their instructions. In order to endorse the impression, that the system actually processes and learns from gestures and speech data a stereo camera was placed in 2.5 meters distance, facing the participant, and a microphone was attached to his/her clothes. The locations of the touch sensors on the back and the head of the robot were explained to the participants.

In order to find out, whether certain features of the task like the presence of an opponent in a game or the opportunity to give not only reward but also instructions have an influence on the user behavior or the user's evaluation of the task, we decided to use four different training tasks with different properties.

a) Find same images

In the *Find same images*-Task, the users taught the robot to compare six different pictures to a sample image in the middle of the screen and to choose the one that corresponds to the sample. This task is shown in figures 1 and 3. While playing, the picture, that the robot is currently looking at, is marked with a green frame to make it easier for the user to understand the robot's viewing direction. When AIBO chooses an image, the robot points at it and a red frame is shown around it. By waving its tail and moving its head, the robot indicates, it is waiting for feedback from its user. Feedback is detected using the touch sensors as well as voice activity detection. When feedback has been detected or a timeout has exceeded the robot proceeds with the next image. The users were asked to provide *instruction as well as reward* to the robot to make it learn to perform the task correctly. The system was implemented in a way that the rate of correct choices and the speed of finding the correct image increased over time.

b) Pairs

In the *Pairs* game, the robot played the children's game *Pairs*, which can be seen in the second picture of figure 2. At the beginning of a *Pairs* game, all cards are displayed upside down on the playfield. The robot chooses two cards to turn around by looking and pointing at them. In case, they show the same image, the cards remain open on the playfield. Otherwise, they are turned upside down again. The goal of the game is, to find all pairs of cards with same images in as little moves as possible. In this task, the participants were asked *not to give instruction* to the robot, which card to chose but *teach the robot to play the game by only giving positive and negative feedback for its moves.* Waiting for and reacting to feedback is realized in the same way as in the *Same Images* task.

c) Connect Four

In the *Connect Four* game, the robot plays the game *Connect Four* against a computer player. Both players take turns to insert one stone into one of the rows in the playfield, which then drops to the lowest free space in that row. The goal of the game is, to align four stones of one's own color either vertically, horizontally or diagonally. As in the *Pairs* task, the participants were asked to *not to give instructions to the robot but provide feedback for good and bad moves* in order to make the robot learn how to win against the computer player.

d) Dog training

In the *Dog Training* task, the participants were asked to teach the speech commands "forward" (mae), "back" (ushiro), "left" (hidari), "right" (migi), "sit down" (suwatte) and "stand up" (tatte) to the robot. The "dog training" task is the only task that does not use the "virtual playfield". In this task, we remote controlled the robot, so the participants could give instructions in the order they liked without the restrictions for

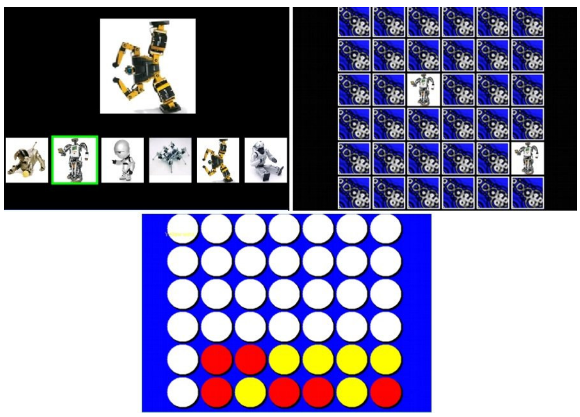

Figure 2: Playfields for the different game tasks: 1) Find same images, 2) Pairs, 3) Connect four

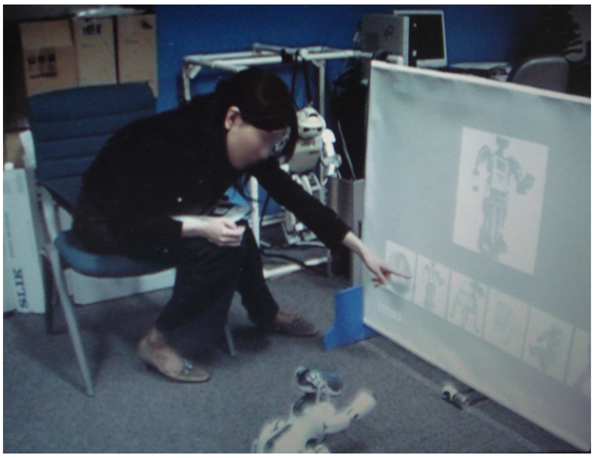

Figure 3: Scene from the video taken during the experiments.

suitable tasks defined in section III. It is the only task, where the robot actually reacts to its user's commands. It is used as a control task to determine, whether the virtual playfield and the game-like nature of the other tasks affect the participants' way of giving reward in any way.

2) Results

As for the modalities used for giving reward, we found a strong preference for speech-based reward. Among 2409 rewards given, 1888 (78.37%) were given by speech, 504 (20.92%) were given by touching the robot and 17 (0.71%) were given by gestures. For the different users, the percentage of speech-based rewards ranged from 52.25% to 97.75%. Gestures were frequently employed by the participants for giving instructions, but we almost did not observe gestures being used for giving positive or negative reward.

Typically, multiple rewards were given for a single positive or negative behavior of the robot. Counting only the rewards given during the time, when the robot signaled that it was waiting for feedback after an action, 3.43 rewards were given for one action on average, usually including one touch reward and one to four utterances. One utterance was counted as one reward. Repetitions of an utterance were counted as multiple rewards. In case of touch reward, one or multiple contacts with the robot's touch sensors were counted as one reward, as long as the participant kept his/her hand close to the sensor.

The favorite verbal feedback differed between the users especially in case of positive reward. None of the utterances, used for positive feedback, appeared within the first six most frequently used utterances for all ten participants. On average, each person shared his/her overall most frequently used positive feedback with one other person. In case of negative reward, the feedback, given by the participants was more homogenous. The most frequently used feedback - "wrong" (chigau) - was preferred by eight out of ten persons. For the two remaining persons, it was the second and third most frequently used feedback utterance.

As for the variability of the feedback, given to the robot by an individual user: On average, participants used 12.3 different verbal expressions to convey positive feedback and 13.4 different expressions to express negative feedback. However, this number varies strongly between individuals: One person always used the same utterance for giving positive feedback and a second utterance for giving negative feedback while the person with the most variable feedback used 30 different expressions for giving positive and 28 different expressions for giving negative feedback. 55.61% of all verbal feedback was given by the participants using their preferred feedback utterance. 88.73% of a user's verbal feedback was given using one of his/her six most frequently used positive/negative utterances, so understanding a relatively small number of different utterances suffices to cover most of a participant's feedback. .

For positive feedback, four out of ten participants had one preferred utterance which did not vary between the four training tasks. In case of negative reward, this was true for five people. For eight out of ten participants in case of positive reward and six participants in case of negative reward, their overall most frequently used feedback utterance was among the top three feedback utterances in each individual task.

In the cases, where the preferred feedback was not the same in all tasks, it typically differed for the "Connect Four" task, while in the three other tasks, including the "Dog Training" control task similar feedback was used as described above. As in the "Connect Four" task it was difficult for the users to judge, whether a move was good or bad in order to provide immediate reward, feedback tended to be very sparse and tentative like "not really good" (amari yokunai), "Is this good?" (ii kana?) or "good, isn't it" (ii deshou).

In a questionnaire, we asked the participants to evaluate their experience throughout the four different training tasks. The participants could rate their agreement with the statements, shown in table 1, on a scale from 1 to 5, where 1 was the best and 5 the lowest rating.

TABLE1: RESULTS OF THE QUESTIONNAIRE

	Same	Pairs	Four	Dog
Teaching the robot through the given task was enjoyable	1.81 1.04	1.90 0.83	1.81 0.89	1.63 0.81
The robot understood my feedback	1.27 0.4	1.81 0.74	2.90 0.85	1.81 0.30
The robot learned through my feedback	1.36 0.59	2.81 0.93	3.45 0.95	1.54 0.69
The robot adapted to my way of teaching	1.45 0.66	2.63 1.05	3.45 1.04	1.64 0.58
I was able to teach the robot in a natural way	2.18 0.96	2.09 0.86	2.54 1.12	1.64 0.69
I always knew, which instruction or reward to give to the robot	2 0.72	2.09 0.86	2.90 1.02	1.91 0.83

(First value: average, second value: standard deviation)

The four tasks were considered almost equally enjoyable by the participants. For the *Find same Images* task and the *Dog Training* task, the participants' impression that the robot actually learned through their feedback and adapted to their way of teaching was strongest. Those two tasks allowed the participants to not only give feedback to the robot but also provide instructions to it. Moreover, they were designed in a way that the robot's performance improved over time.

In the *Dog Training* task, the robot was remote controlled to react to the user's commands and feedback in a typical Wizard of OZ-Scenario. However, in the *Same Image* task, the user's instructions and feedback were not actually understood by the robot but anticipated from the state of the training task which was designed using the restrictions defined in section III. This did not have a negative impact on the participants' impression that the robot understood their feedback, learned through it and adapted to their way of teaching, compared to the Wizard of Oz scenario.

The lowest ratings were given for the "Connect Four" task. As the robot's moves could not be evaluated as easily, as in the other tasks, the participants were unsure, which rewards to give and therefore did not experience an effective teaching situation.

VI. Discussion

From our experiments, we found that for learning and understanding users' positive/negative feedback, emphasis should be placed on the modalities speech and touch. Gesture is mainly used for instruction and explanation.

Natural feedback given by different users can vary strongly, especially in case of positive rewards. Therefore, learning to understand the feedback that a certain user employs instead of using hard-coded and potentially unintuitive commands, which have to be learned by the user, helps to ensure natural interaction and a positive user experience.

Learning to understand feedback through a training task is only feasible and useful, if the feedback given by one user is similar within different tasks. The results from the experiments suggest that this is actually the case and that typically a limited number of utterances are used by an individual to convey positive and negative reward.

However, there are cases where the contents of the utterances alone may not be correctly understood as a positive or negative reward: For instance, some of the users occasionally just repeated their previous command in a stricter tone before or instead of giving other negative feedback to the robot. In these cases, analyzing and learning the prosody which determines the sentence melody of typical positive and negative feedback utterances can be expected to improve the recognition accuracy.

Problems arise, if the user is not exactly sure, how to judge the robot's behavior, as in the "Connect Four" task. Therefore, for automatically learning rewards, the task has to be designed in a way that it is easy to understand for the user. Otherwise the amount and reliability of the given reward, as well as the user's motivation to complete the training, decrease.

Comparing the results of our study with existing work on how people use multimodality is difficult because our study differs from most existing work in two major aspects: It focuses on positive and negative reward instead of instructions or commands and it uses a robot as interaction partner, not a GUI-based system.

Oviatt states in [6] that multimodal combining of speech and gesture is mainly used to describe spatial relations. In her experiments, less than 1% of the participants expressed themselves multimodally when performing actions without any spatial component. This could be one explanation, why we observed almost no gestures for giving positive and negative feedback, while gestures could be observed when the participants gave instructions to the robot like, for instance, pointing to an image to chose.

Our experiments have been conducted with a dog-shaped AIBO pet robot. This might have persuaded the users to interact with the robot in a similar way as with an actual pet. Especially the frequent use of touch to express approval or disapproval may be caused by the dog-like appearance of the robot. To ensure the transferability of our findings to general human-robot-interaction tasks, experiments with other types of robots would be required.

REFERENCES

[1] A. Austermann, S. Yamada, "Learning Reward Modalities for Human-Robot-Interaction in a Cooperative Training Task", Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication 2007

[2] A. Austermann, S. Yamada, " Learning to Understand Multimodal Rewards for Human-Robot-Interaction using Hidden Markov Models and Classical Conditioning" *To appear: Proceedings of the WCCI 2008*

[3] C. Balkenius and J. Morn. "Computational models of classical conditioning: a comparative study." Proceedings of the Fifth International Conference on Simulation of Adaptive Behavior, 1998

[4] K. Dautenhahn, "Methodology & Themes of Human-Robot-Interaction: A Growing Research Field", International Journal of Advanced Robotic Systems, Special Issue on Human-Robot-Interaction,Vol. 4, No. 1, pp. 103-108 , 2007

[5] A. Haasch , S. Hohenner, S. Hüwel, M. Kleinehagenbrock, S. Lang, I. Toptsis, G. A. Fink, J. Fritsch, B. Wrede, G. Sagerer, "BIRON - The Bielefeld Robot Companion" (In E. Prassler, G. Lawitzky, P. Fiorini, and M. Hägele, editors), Proc. Int. Workshop on Advances in Service Robotics, pages 27-32, Stuttgart, Germany, May 2004. Fraunhofer IRB Verlag.

[6] S. Oviatt, "Ten myths of multimodal interaction" Communications of the ACM, Vol. 42 , No. 11, pp. 74 – 81, 1999

[7] R. Stiefelhagen, C. Fugen, P. Gieselmann, H. Holzapfel, K. Nickel and A. Waibel "Natural Human-Robot Interaction using Speech, Head Pose and Gestures" Proceedings. 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems, 2004. (IROS 2004). 2422- 2427 vol.3, 2004

[8] A. L. Thomaz, G. Hoffman, and C. Breazeal. "Reinforcement Learning with Human Teachers: Understanding how people want to teach robots." Proceedings of the 15th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN), 2006.

[9] M. Ullerstam, M. Mizukawa, "Teaching robots behavior patterns by using reinforcement learning: how to raise pet robots with a remote control", SICE 2004 Annual Conference Volume 1, Issue , 4-6 Aug. 2004 Page(s): 143 - 146 vol. 1

[10] J. Wainer, D. J. Feil-Seifer, D. A. Shell, M. J. Mataric "Embodiment and Human-Robot Interaction: A Task-Based Perspective", , Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication 2007

[11] S. Young, G. Evermann, M. Gales, T. Hain, D. Kershaw, X. Liu, G. Moore, J. Odell, D. Ollason, D. Povey, V. Valtchev, P.Woodland "The HTK Book (for HTK Version 3.4)", 2006