

Crowdsourcing Human-Robot Interaction: Application from Virtual to Physical Worlds

Sonia Chernova, Nick DePalma, Elisabeth Morant and Cynthia Breazeal

Abstract—The ability for robots to engage in interactive behavior with a broad range of people is critical for future development of social robotic applications. In this paper, we propose the use of online games as a means of generating large-scale data corpora for human-robot interaction research in order to create robust and diverse interaction models. We describe a data collection approach based on a multiplayer game that was used to collect movement, action and dialog data from hundreds of online users. We then study how these records of human-human interaction collected in a virtual world can be used to generate contextually correct social and task-oriented behaviors for a robot collaborating with a human in a similar real-world environment. We evaluate the resulting behavior model using a physical robot in the Boston Museum of Science, and show that the robot successfully performs the collaborative task and that its behavior is strongly influenced by patterns in the crowdsourced dataset.

I. INTRODUCTION

Robots require a broad range of interaction skills in order to work effectively alongside humans. They must have the ability to detect and recognize the actions and intentions of a person, to produce functionally valid and situationally appropriate actions, and to engage in social interactions through physical cues and dialog. A number of robotic platforms capable of these types of interactions have been developed for different applications, including museum guidance [1], reception desk assistance [2] and elder care [3].

Research on action and dialog generation has also been conducted in the gaming community in the context of character development for role-playing games [4], [5]. All of the above approaches present successful solutions for their respective applications based on carefully hand-crafted models for action and dialog generation. The typical result for this type of development process is a system that is capable of natural and engaging interaction for some range of topics, but only for those that were predetermined by the programmers.

Data-driven techniques present an alternate solution to hand-crafted models. These approaches utilize datasets of hundreds of example behaviors, often from a multitude of different users, to generate appropriate responses to input. Successful data-driven techniques have been demonstrated in a number of interactive applications, such as dialog management [6], [7].

The question we explore in this paper is whether data-driven approaches can be developed for interactive robotic systems. Can robot behaviors be *crowdsourced* to produce natural, engaging and functionally appropriate actions and dialog based on data from hundreds of people? Ideally, such an approach would benefit from the “power of the masses”, producing a more general result by incorporating examples from many users.

The challenge for crowdsourcing human-robot interaction is to develop a method for gathering interaction data on a large scale. One solution is to utilize the Wizard-of-Oz technique in which data is recorded as a robot is puppeteered through the task by a human subject. The significant drawbacks of this approach are that it is limited by the time required to recruit subjects, the number of physical robots that can be used to collect data at the same time, and the cost of maintaining these robots. The outcome of these factors is that Wizard-of-Oz techniques are typically limited to only several dozen participants.

In this paper, we propose the use of online games as a means of generating large-scale data corpora for human-robot interaction research. We present a data collection approach based on an online multiplayer game that was used to collect data from hundreds of users. We then present a data-driven policy generation technique utilizing Case-Based Reasoning (CBR) [8], [9] that emphasizes high-level memory learning and incorporates a large corpus of online, virtual interactions into a general task-oriented behavior model. We apply this model to a physical robot interacting with visitors at the Boston Museum of Science, and show that the robot successfully performs the collaborative task and that its behavior is strongly influenced by patterns in the crowdsourced dataset.

II. RELATED WORK

Our work is inspired by the Restaurant Game project [10], [11], in which data collected from thousands of players in an online game is used to create an automated data-driven behavior and dialog authoring system. The Restaurant Game is a minimal investment multiplayer online (MIMO) game that enables users to log in to a virtual environment and take the role of one of two characters, a customer or a waiter, at a restaurant. Players are randomly paired with another online player and can interact freely with each other and objects in the environment. In addition to standard game controls, the users can maintain dialog with each other, and other simulated characters, by typing freeform text. Logs of over 5,000 games were used by the authors to analyze

S. Chernova is with the Department of Computer Science, Worcester Polytechnic Institute, USA soniac@cs.wpi.edu

N. DePalma, E. Morant and C. Breazeal are with the Media Laboratory at the Massachusetts Institute of Technology, USA {ndepalma, cynthiab}@mit.edu

this interactive human behavior and acquire contextualized models of language and behavior for collaborative activities.

The use of games as a data mining technique has been shown to be highly successful in a number of applications. The term “Games With A Purpose”, or GWAP, was coined by Luis von Ahn to describe games that address computational problems [12]. The idea behind this approach is to make work fun by turning a scientific task into a game, thereby harnessing the computational power of internet users.

The idea of data-driven dialog systems has also been around for a long time, and there are a number of successful applications outside of gaming and robotics. For example, Reiser and Lemon present an automated system for information seeking dialogs, such a speech-driven navigation and track selection system for an MP3 player [13]. Another example of human-driven computation is the Open Mind Initiative, a worldwide research endeavor for developing intelligent software by leveraging human skills to train computers [14]. This approach is based on a network of volunteers that participate by providing answers to questions that computers can not yet answer, for example describing the content of an image.

Our work differs from all of the above approaches in that data collection is performed in a fundamentally different domain than the one in which it is applied. Furthermore, the goal of this work is to interleave spoken dialog and physical actions observed in humans into collaborative task-oriented robot behaviors.

III. APPROACH

The goal of our research is to study how records of human-human interaction collected in a virtual world can be used to generate natural and robust human-robot interactive behavior in similar real-world environments. In this section we first present a description of the data collection process and the resulting data set. We then present our data-driven approach for autonomous behavior generation utilizing Case-Based Reasoning, followed by a description of the real-world domain and user study. We designed the virtual and real-world domains as a collaborative search and retrieval task that can be generalized to many applications. The task has no strictly assigned social roles, however, the domain is designed to encourage collaborative behaviors such as action synchronization, sequencing and dialog.

A. Crowdsourcing Interaction using an Online Game

The purpose of the virtual world is to model a real-world environment in which a physical robot would perform a task with a human. *Mars Escape* is a two-player online game in which two randomly-paired players take on the roles of a human astronaut and a robot on Mars. The object retrieval task is incorporated into the back-story of the game, in which players are told that the oxygen generator on their remote research station has failed. In order to successfully complete the mission, the pair must locate and salvage five items and return to the spaceship before oxygen supplies run out. The




Fig. 1. A screenshot of the *Mars Escape* game showing the action menu and dialog between the players.

list of required items is presented in a side-bar inventory screen. The players have 10 minutes to complete the mission.

During the game, the players are able to navigate in the environment, manipulate objects using six predetermined actions (*pick up, put down, look at, go to, use, analyze*) and communicate with each other through in-game text-based chat (Fig. 1). All player actions and dialog are recorded by the game server. The game terminates when the players choose to exit, or when the game clock runs out.

We designed the domain such that the retrieval of each of the five items required a different type of collaboration in order to study a diverse set of behaviors. Table I presents a list of the five inventory items. For each item, the table includes a summary of how the object can be retrieved and a description of the class of problems the object represents. The object locations and avatar capabilities were designed to balance the role of both players, and to ensure that teamwork and communication were necessary to perform the complete mission successfully.

Following the completion of the game, players were asked to complete a survey evaluating their gaming experience and the performance of their partner. The survey contained the following eight questions and users were asked to rate their responses on a 5-point Likert scale from “strongly disagree” (1) to “strongly agree” (5):

- 1) My overall game experience was enjoyable.
- 2) The other player’s performance was an important contribution to the success of the team.
- 3) The human-robot team did well on the task.
- 4) The actions of the other player were rational.
- 5) The other player communicated in a clear manner.
- 6) The other player performed well as part of the team.
- 7) The other player’s behavior was predictable.
- 8) The other player was controlled by a human.

B. Interaction Dataset

During the first three months of the release of the game we captured data from 558 two-player games. Of these, approximately 700 player logs were retained for analysis after excluding logs in which a player exited the game prematurely by quitting the application and not filling out

Item	Game Context	Generalization
<i>Research Journal</i>	Located on top of a stack of boxes. Reachable only by the astronaut.	A task that can be performed by only one of the players.
<i>Captured Alien</i>	Located on a raised platform. Reachable by either player after lowering a platform using wall mounted controls.	A task that can be performed by either player.
<i>Canister</i>	Located on top of toxic barrels. Reachable by either player, but the astronaut loses 10 points for coming in contact with chemicals.	A task for which one player is better suited than the other.
<i>Memory Chip</i>	Appears when both players stand on a weight sensor at the same time. Reachable by either player.	A task that requires action synchronization.
<i>Sample Box</i>	One of 100 identical boxes located on a high shelf. The astronaut can pick up and look at each box individually to locate the correct one. The robot can identify the exact box, but can not reach the box due to its height. Optimal solution is for the robot to scan the boxes and tell the astronaut the sample's exact location.	A task that requires coupled actions and dialog.

TABLE I
DESCRIPTION OF THE FIVE OBJECTS PLAYERS MUST OBTAIN TO SUCCESSFULLY COMPLETE THE GAME.

the survey. The following is an example transcript showing an interaction in which the astronaut (A) and the robot (R) retrieve the alien.

```
A: "hi"
R: "hey"
R: "i'll get the yellow can"
A: "ok, i'll get the book"
<astronaut picks up book>
<robot picks up canister>
<astronaut places book in container>
A: "lets do the weight sensor next"
<astronaut enters weight sensor>
<robot places book in container>
R: "ok"
<robot enters weight sensor>
```

The retrieval of different items provided different degrees of challenge to the players, leading to some objects being commonly retrieved before others. The majority of players first picked up those items that were in clearly visible locations and could be retrieved individually (i.e. the canister and the journal), delaying the retrieval of collaborative items. In later sections we discuss how these temporal patterns impact the behavior of the robot in real-world experiments.

Somewhat unexpectedly, we found that only 57% of player pairs successfully collected all five items in the duration of the game (86% collected three or more). Of the five items, in 75% of games the last item to be retrieved was the sample box, the item that required the greatest degree of collaboration and communication between players. Furthermore, of the games in which players collected 4 items and missed only one, 89% were missing the sample box. As a result of these behavioral patterns, the distribution of crowdsourced data is not uniform across the entire range of player behaviors. In our analysis of real-world interactions, we show how this bias affects the robot's behavior.

C. Data Processing

Our goal is to leverage the corpus of interaction data describing the movements, actions and spoken dialog of players in the virtual world, to generate contextually correct social and task-oriented robot behaviors in the real world. We are particularly interested in exploring the degree to which data gathered in the virtual world can be directly leveraged in the real world through data-driven techniques. To study

this question we chose a memory-based approach to behavior generation utilizing Case-Based Reasoning [8], [9].

Case-Based Reasoning utilizes a library of past experiences (cases) to solve new problems by finding a similar past cases and reusing them in the new situation. CBR has been successfully applied to autonomous robot control in many applications, including indoor navigation [15] and autonomous robot soccer [16]. In this work, we use the interaction corpus collected in the online game to create a case library and apply CBR retrieval to generate autonomous robot behavior in the real world. Our case library contains only examples recorded in the virtual world. Techniques for augmenting this dataset with new examples from the physical world are an interesting topic for future research.

We represent each case in our dataset using the following 13 features: the previous robot action, the previous astronaut action, the previous robot spoken phrase, the previous astronaut spoken phrase, object held by astronaut, object held by robot, and the area location (e.g., center, near shelf, near toxins, etc.) of the astronaut, robot, journal, alien, chip, canister and sample box. This information enables us to track the most recent events, in terms of agent actions and dialog, while also maintaining a coarse history of past events based on current object locations (each state vector represents a snapshot in time that describes which items have already been retrieved). Note that continuous navigation of characters across the domain is represented using the discrete *go to* commands between pre-defined discrete regions of the space in order to reduce number of cases and to generalize across similar actions.

Based on the representation described above, our dataset resulted in 82,479 cases. For case storage and retrieval we utilized the open source FreeCBR software package [17]. During case retrieval, the current state of the robot is encoded using the feature vector and compared to the library of recorded cases. Similarity between the query and cases in the library is calculated based on a weighted sum of differences between features. We selected the weight for each feature based on the accuracy of the measure of that feature. For example, the weighting for all object locations was high because we were able to track this information with high accuracy, whereas a low weight value was used to compare speech data due to noise in the speech recognition system.

D. User Study using a Physical Robot

One of the central motivations for this research is that crowdsourcing human-robot interaction will enable the development of interaction models that are diverse and robust, capable of natural interaction with people of different ages and technical backgrounds. To test this theory we performed a user study by recreating the virtual *Mars Escape* environment in the real world at the Boston Museum of Science. Museum visitors were recruited one at a time to perform the task in collaboration with our autonomous MDS robot *Nexi* (Fig. 2(a)). The MDS robot platform combines a mobile base with a socially expressive face and two dexterous hands that provide the capability to grasp and lift objects. The robot is equipped with a biologically-inspired vision system that supports animate vision for shared attention to visually communicate the robot's intentions to human observers. Auditory inputs support a microphone array for sound localization, as well as a dedicated channel for speech recognition via a wearable microphone.

Due to the complexity of the search and retrieval task, a high precision offboard Vicon MX camera system was used to supplement the robot's onboard sensors. The Vicon system was used to track the position of the robot, human, and objects in the environment in real time using lightweight reflective markers attached to object surfaces. This tracking system enabled the robot to have the degree of environmental awareness that is comparable to that of a human.

The physical setup of the space closely modeled the game environment. It contained five mission objects, in similar placements to their in-game counterparts, including a tall shelf to keep the journal out of the robot's reach; a raised platform that could be lowered to access the alien; toxic barrels near which the human user was warned to step away; a box that would automatically open to reveal a chip when both players stepped onto the scale; a shelf unit containing several dozen small numbered boxes, one of which contained a sample; and a number of other props, such as empty crates and tools. Figure 2(b) shows the robot reaching for the button that activates the moving platform while the human




Fig. 2. (a) The MDS robot platform. (b) The robot and a participant performing the study at the Boston Museum of Science.

participant retrieves the journal from the shelf.

A total of 18 participants were recruited to perform the study, 11 male and 7 female, with ages ranging from 13 to 35. None of the participants had previously played the *Mars Escape* game online. The study environment was curtained off from the rest of the museum, preventing participants from observing others performing the task. As a result, participants had neither familiarity with the robot nor knowledge of the space layout and object locations before beginning the study.

Before the start of the experiment participants were given basic instructions and were informed that the robot could understand speech, move around, pick up objects and use its sensors to locate organic lifeforms. Upon entering the study area, participants were greeted by the robot using the following standard greeting: "Hello. Lets find the five items on our list". After the greeting, the robot began autonomous action selection using the CBR algorithm. We chose to use a standard greeting after observing several preliminary trials in which the robot immediately began autonomous behaviors related to the task without a greeting. This led to confusion with some study participants who were expecting a response to their appearance in the room and were uncertain about whether the study had started. Note that while our dataset contains some greetings and introductions between players, many players did not engage in this behavior in the virtual world, and as a result the robot may not do so either. The use of a standard greeting eliminated this uncertainty for the study participants.

For each participant, the study continued until either the team retrieved all five objects, the robot failed (e.g., low battery), or the study participant requested to stop the session (none of the participants requested to stop). Data from incomplete trials was discarded. Following the study, participants were asked to fill out the same questionnaire as in Section III-A.

Of the 18 participants, 16 successfully completed the study¹. The behavior of participants toward the robot and the task varied greatly. Some users acted independently, freely explored the space, quickly found all objects and tried to help the robot with its parts of the task. A similar number of users took a passive role, expected the robot to take charge and to provide all the answers. As a result of not exploring the space, passive users often had a hard time finding some of the items. Note that unlike the online game, the study only ended once all five objects were retrieved. If a user had trouble finding an item but did not request to end the study, a behind the scenes robot operator would briefly override the autonomous behavior selection mechanism of the robot to provide extra guidance to the person.

Unsurprisingly, locating the sample box proved the biggest challenge for the human-robot teams, resulting in a number of operator interventions. A common interaction would be for the robot to approach the shelf of boxes while the human was paying attention elsewhere. The robot would scan the shelf, announce the location of the box containing the sample, then

¹A sample video is available at <http://www.vimeo.com/24546560>

continue on to other tasks with the assumption that the user would pick up the box (the robot is unable to retrieve the box in both the online and real-world versions of the task). While approximately half of the participants appropriately responded by picking up the sample box item, others either did not pay attention to the message or simply observed the robot and did not take initiative to move toward the shelf. In these cases, if the sample box was not retrieved after some time, the remote operator would manually re-direct the robot back to the shelf to repeat its instructions. Other overrides included error-recovery, such as asking the person to help pick up an object that the robot had tried to pick up but dropped. On average, 7 overrides were performed during the course of a trial, accounting for approximately 36% of the total number of robot actions.

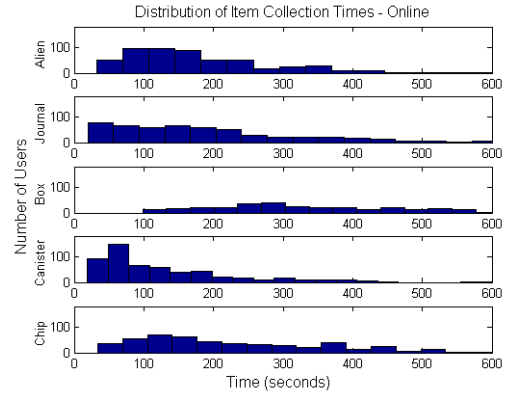
IV. COMPARISON AND DISCUSSION

In this section we compare the behavior of real-world human-robot teams to human-human interactions in the virtual world, and examine the survey responses of participants following both types of studies. We are particularly interested in examining how the broad patterns of behavior observed in online games, with respect to the order in which users retrieved objects and the relative difficulty of locating different objects, translates into the real world where human users will have different perceptual abilities than their avatar.

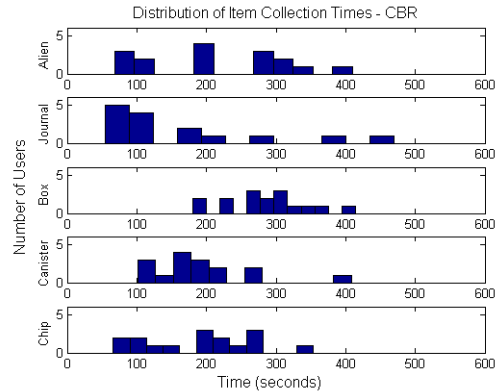
Figure 3 presents histograms showing the distribution of times at which items were collected by participants in the online and real-world tasks. The x-axis shows the elapsed time in the trial in seconds. We note several interesting patterns in the data. The online data shows a clear distinction between the sample box item (labeled “Box”) and the rest of the items. The sample box distribution comes temporally much later, and we clearly observe that fewer players successfully retrieved that item. The chip, the other item that required collaboration between players, also shows a distribution skewed further along the timeline than the easily accessible canister, journal and alien.

Comparing data from the online study (Figure 3(a)) and the museum (Figure 3(b)) we observe similar distributions in the data for all five items, with several noteworthy differences. The time between the beginning of the study and the retrieval of the first object is longer in the real-world scenario. This can be attributed to the participants taking time to evaluate the surroundings and observe the robot. Additionally, the time at which the canister is picked up tends to occur later in the experiment. This is a result of the fact that in the real world, the robot’s pickup action requires approximately 60 seconds to complete, whereas in the online game this action was instantaneous. The overall run time for the online and real-world conditions was similar.

Critically, this data highlights the fact that many of the same high level behavior patterns are present in both the virtual and real-world applications of this task. While some of this effect is likely due to the structure of the domain itself, such as the placement of visible objects, we hypothesize that the robot’s action selection choices also play a significant



(a) Online Game



(b) Museum of Science

Fig. 3. A comparison of the distribution of times at which items were collected by the players.

role. For example, had the robot chosen always to locate the sample box before retrieving the canister, the distribution of the data would be very different. Further studies are needed to verify this hypothesis, and to test to what degree the action selection of the robot influences the behavior of the human.

Figure 4 presents a comparison of the survey results from both the online and museum study conditions. We observe that the majority of players in both conditions enjoyed taking part in the study and felt that the robot contributed to the success of the team, performed well at the task and behaved rationally.

The greatest differences in the survey responses are observed with respect to communication (question 5, (e)), predictability (question 7, (g)), and whether the robot was controlled by a human (question 8, (h)). Participants wanted better communication with the robot, with one participant writing in the comment section of the survey: “It should have asked me how it could help/tell me what it was capable of/knew”. The perceived predictability of the robot was significantly lower than that of its online counterparts. We hypothesize that this is due to the nature of the Case-Based Reasoning algorithm that retrieves actions one at a time and maintains little hysteresis about past actions, resulting in a behaviors that may switching between different activities at

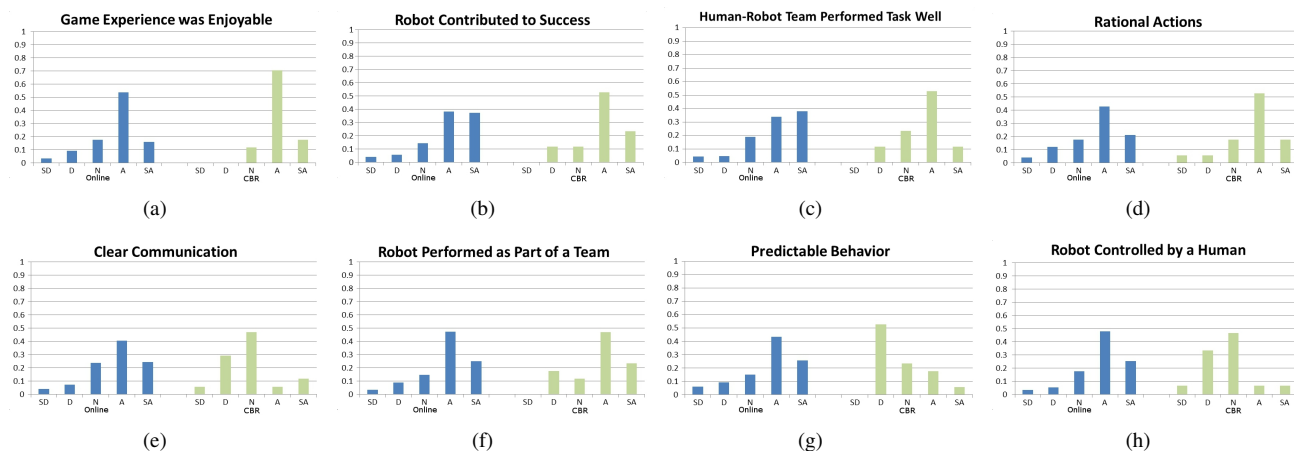


Fig. 4. Comparison of the results of the 8 survey questions between the online (blue, left) and museum (green, right) conditions. Histograms show percentage of responders identifying with each point on the Likert scale: strongly disagree (SD), disagree (D), neutral (N), agree (A), and strongly agree (SA). Questions 1-8 correspond to figures (a)-(h), respectively.

random. In future work, we will examine other algorithms that will generate behaviors based on longer sequences of actions. Finally, potentially as a result of the communication and predictability factors, most users correctly believed that the robot was indeed not controlled by a human. Interestingly, a significant portion of online players were also uncertain as to whether their teammate was human or not.

V. CONCLUSION

The ability for robots to engage in interactive behavior with a broad range of people is critical for future development of social robotic applications. In this paper, we presented a novel approach to generating task-specific social behaviors based on crowdsourcing human-robot interaction in virtual worlds. We described a data collection approach based on an online multiplayer game that was used to collect data from hundreds of users. We then leveraged the corpus of interaction data that described the movements, actions and spoken dialog of players in the virtual world, to generate contextually correct social and task-oriented behaviors for a robot operating in a real-world environment. We showed that very similar patterns of behavior are observed in both the online and real-world task.

This is the first study that we are aware of examining large-scale online crowdsourcing for human-robot interaction. The comparison of questionnaire answers across both the online and real-world conditions shows that participants enjoyed taking part in the interaction and rated the robot similarly to human partners in several critical measures. This work sets a baseline for what can be achieved through direct data-driven methods, and suggests many interesting directions for future research in this area. Ultimately, we believe that crowdsourcing in virtual worlds has the potential to become a powerful tool in human-robot interaction research.

VI. ACKNOWLEDGEMENTS

This work was supported by the Office of Naval Research Award Numbers N000140910112 and N000140710749, and

the Microsoft Corporation. We would also like to thank Dan Noren and the staff of Cahners Computer Place at the Boston Museum of Science for helping organize the study exhibit.

REFERENCES

- [1] W. Burgard, A. B. Cremers, D. Fox, D. Hähnel, G. Lakemeyer, D. Schulz, W. Steiner, and S. Thrun, "The interactive museum tour-guide robot," in *AAAI/IAAI*, 1998, pp. 11–18.
- [2] M. K. Lee and M. Makatchev, "How do people talk with a robot?: an analysis of human-robot dialogues in the real world," in *CHI*. New York, NY, USA: ACM, 2009, pp. 3769–3774.
- [3] B. Graf, M. Hans, and R. D. Schraft, "Care-o-bot II—development of a next generation robotic home assistant," *Auton. Robots*, vol. 16, no. 2, pp. 193–205, 2004.
- [4] M. McNaughton, J. Schaeffer, D. Szafron, D. Parker, and J. Redford, "Code generation for ai scripting in computer role-playing games," in *Challenges in Game AI Workshop at AAAI-04*, 2004.
- [5] G. Kacmarcik, "Question-answering in role-playing games," in *Workshop on Question Answering in Restricted Domains, AAAI-05*, 2005.
- [6] A. L. Gorin, G. Riccardi, and J. H. Wright, "How may I help you?" *Speech Commun.*, vol. 23, no. 1-2, pp. 113–127, 1997.
- [7] S. Singh, D. Litman, M. Kearns, and M. Walker, "Optimizing dialogue management with reinforcement learning: experiments with the NJFun system," *J. Artif. Int. Res.*, vol. 16, no. 1, pp. 105–133, 2002.
- [8] J. Kolodner, *Case-based reasoning*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1993.
- [9] A. Aamodt and E. Plaza, "Case-based reasoning; foundational issues, methodological variations, and system approaches," *AI COMMUNICATIONS*, vol. 7, no. 1, pp. 39–59, 1994.
- [10] J. Orkin and D. Roy, "The restaurant game: Learning social behavior and language from thousands of players online," *Journal of Game Development*, 2007.
- [11] —, "Automatic learning and generation of social behavior from collective human gameplay," in *AAMAS*, 2009, pp. 385–392.
- [12] L. von Ahn and L. Dabbish, "Designing games with a purpose," *Communications of the ACM*, vol. 51, no. 8, pp. 58–67, 2008.
- [13] V. Rieser and O. Lemon, "Learning human multimodal dialogue strategies," *Nat. Lang. Eng.*, vol. 16, pp. 3–23, 2010.
- [14] D. G. Stork, "The open mind initiative," *IEEE Intelligent Systems and Their Applications*, vol. 14, no. 3, pp. 19–20, 1999.
- [15] A. Micarelli, S. Panziera, and G. Sansonetti, "Case-based reasoning in robot indoor navigation," in *Proceedings of the 7th international conference on Case-Based Reasoning: Case-Based Reasoning Research and Development*, ser. ICCBR. Berlin, Heidelberg: Springer-Verlag, 2007, pp. 284–298.
- [16] R. Ros, J. L. Arcos, R. Lopez de Mantaras, and M. Veloso, "A case-based approach for coordinated action selection in robot soccer," *Artif. Intell.*, vol. 173, pp. 1014–1039, June 2009.
- [17] L. Johanson, "FreeCBR," <http://freecbr.sourceforge.net>.