Crowd-Sourcing Real-World Human-Robot Dialogue and Teamwork through Online Multiplayer Games

Sonia Chernova, Nick DePalma, Cynthia Breazeal

■ *We present an innovative approach for large-scale data collection in human-robot interaction research through the use of online multiplayer games. By casting a robotic task as a collaborative game, we gather thousands of examples of human-human interactions online, and then leverage this corpus of action and dialogue data to create contextually relevant social and task-oriented behaviors for human-robot interaction in the real world. We demonstrate our work in a collaborative search and retrieval task requiring dialogue, action synchronization, and action sequencing between the human and robot partners. A user study performed at the Boston Museum of Science shows that the autonomous robot exhibits many of the same patterns of behavior that were observed in the online data set and survey results rate the robot similarly to human partners in several critical measures.*

We envision the need for robots to be not only functional, but adaptable, robust to the diversity of human behaviors and speech patterns, and capable of acting in a both task and socially appropriate manner. Natural and diverse human-robot interaction (HRI) of this kind has been a long-standing goal for robotics research, and a broad range of approaches have been proposed for the development of robots that support diverse interactions. Among proposed techniques, variants that are dependent on hand-coded rule sets and probabilistic single-task policy learning methods have proven to be too brittle for interactive applications, failing to generalize over the diversity of possible inputs. Such systems typically force the user to adapt their method of interaction to fit the coded requirements of the robot.

A different approach to creating more humanlike robotic systems has focused on imitating human cognitive processes by developing large scale cognitive architectures that support many modalities and interaction styles. While such systems have been shown to successfully support a broad range of interactions, they rely heavily on precoded data. For example, dialogue responses are typically limited to only one or two dozen phrases, which pales in comparison to the diversity of human speech.

We believe that in order for robotic systems to become a truly ubiquitous technology, robots must make sense of natural human behavior and engage with humans in a more humanlike way. Robots must become more like humans instead of forcing humans to be more like robots.

Much of human knowledge about the appropriateness of behavior, in terms of both speech and actions, comes from our personal experiences and our observations of others. Common

  ISSN 0738-4602

Item	Game Context	Generalization
Research Journal	top of stack of boxes	reachable by only one player
Captured Alien	on a raised platform	reachable by either player
Canister	on toxic barrels	reachable by either, but one is better suited
Memory Chip	appears when both players stand on a weight sensor	requires action synchronization
Sample Box	one of 100 identical boxes located on a high shelf	requires coupled actions and dialog

Table 1. Description of the Five Objects Players Must Obtain to Successfully Complete the Game.

behaviors that are repeated hundreds of times in different variations form a knowledge base from which we learn what to say and what actions to perform to achieve certain goals. Acquiring an equivalent data set for human-robot interaction has proven challenging, largely due to the costs of running user studies while recruiting hundreds of participants to perform the task of interest. As a result, while data-driven techniques and crowd sourcing have become the norm in other areas of robotics, they have not been widely adopted in HRI.

We propose an innovative approach for data collection through the use of virtual characters, and in particular online multiplayer games. Motivated by projects such as the Restaurant Game (Orkin and Roy 2007, 2009) and Games with a Purpose (von Ahn and Dabbish 2008), we show that by casting a robotic task as a collaborative multiplayer game we are able to gather thousands of examples of human-human interaction. We describe how we transfer this data into the physical world and use it to generate natural and robust human-robot interactive behavior in a similar real-world environment.

We demonstrate our work in a collaborative search and retrieval task requiring dialogue, action synchronization, and action sequencing between the human and robot partners. Using a custom game developed around the task, called Mars Escape, we randomly paired online players in the roles of a human and robot. We recorded data from more than 700 players, resulting in a diverse interaction corpus. In the following sections we describe the game and data set, and then discuss how the online data can be mapped into the physical world and leveraged to generate natural, task-centered interactive robot behaviors. While there are many possible techniques for utilizing this data, in this work we report on a memory-based approach based on case-based reasoning (CBR) (Micarelli, Panzieri, and Sansonetti 2007) as a means of studying to what degree crowd-sourced online data can be transfered directly to real-world domains.

We report results of a user study evaluating the resulting autonomous robot system at the Boston Museum of Science. We compare its performance to a teleoperated robot following a scripted task protocol and examine both the behavior of the robot and participant survey responses. We show that the robot successfully performs the collaborative task and exhibits many of the same patterns of behavior that were observed in the online data set. Finally, we discuss open research problems, such as robust transfer from online to physical worlds, integration of data-driven techniques with cognitive architectures, and generalization across domains.

Online Game

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. Figure 1 presents a screenshot of the game, showing the action menu and dialogue between players. We designed the game to model a collaborative search and retrieval task in which the players must locate and retrieve the following five items to successfully complete the mission: *research journal, captured alien, canister, memory chip*, and *sample box*. The object retrieval task is incorporated into the backstory of the game, in which players are told that the oxygen generator on their remote research station has failed, and that the pair must salvage the most critical items and return to the spaceship before oxygen supplies run out (10 minutes). Collaboration and communication between players are required to complete the entire task. Table 1 lists the location of each object and a description of the class of problems it represents.

During the game, 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. All player actions and dialogue are recorded by the game server. The game terminates when the players choose to exit, or when the game clock runs out. Players are then asked to complete a survey consisting of the following eight questions rated on a 5-point Likert scale (strongly agree to strongly disagree):

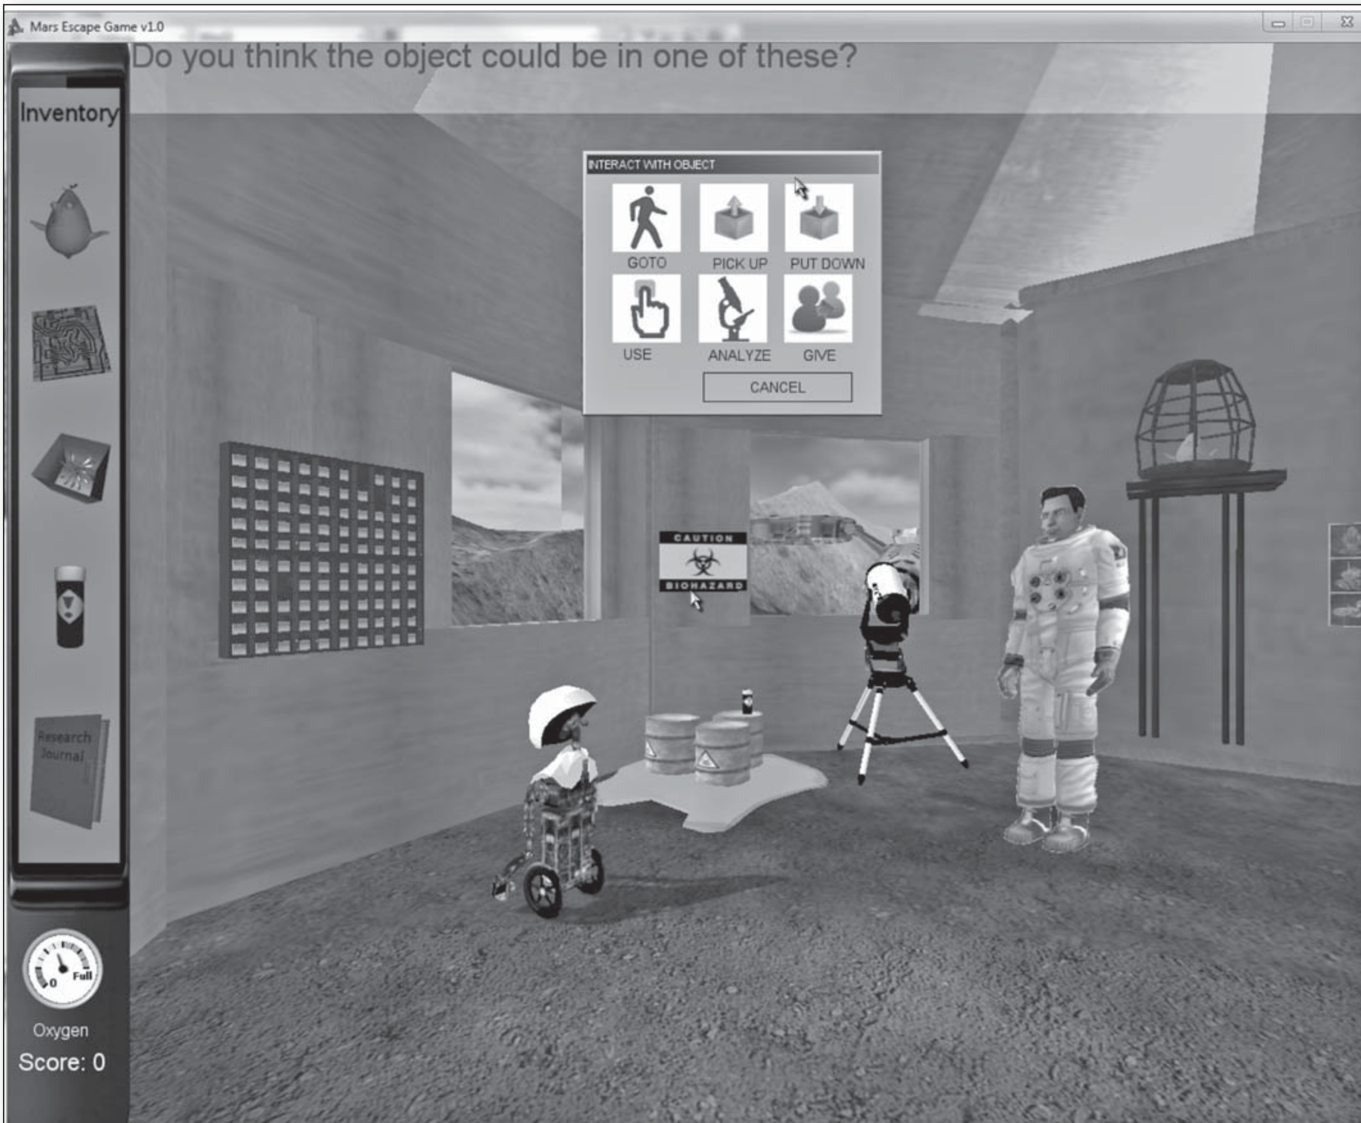

Figure 1. The Mars Escape Game.

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 actions of the other player were rational.
4. The human-robot team did well on the task.
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.

Online Interaction Data Set

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

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]

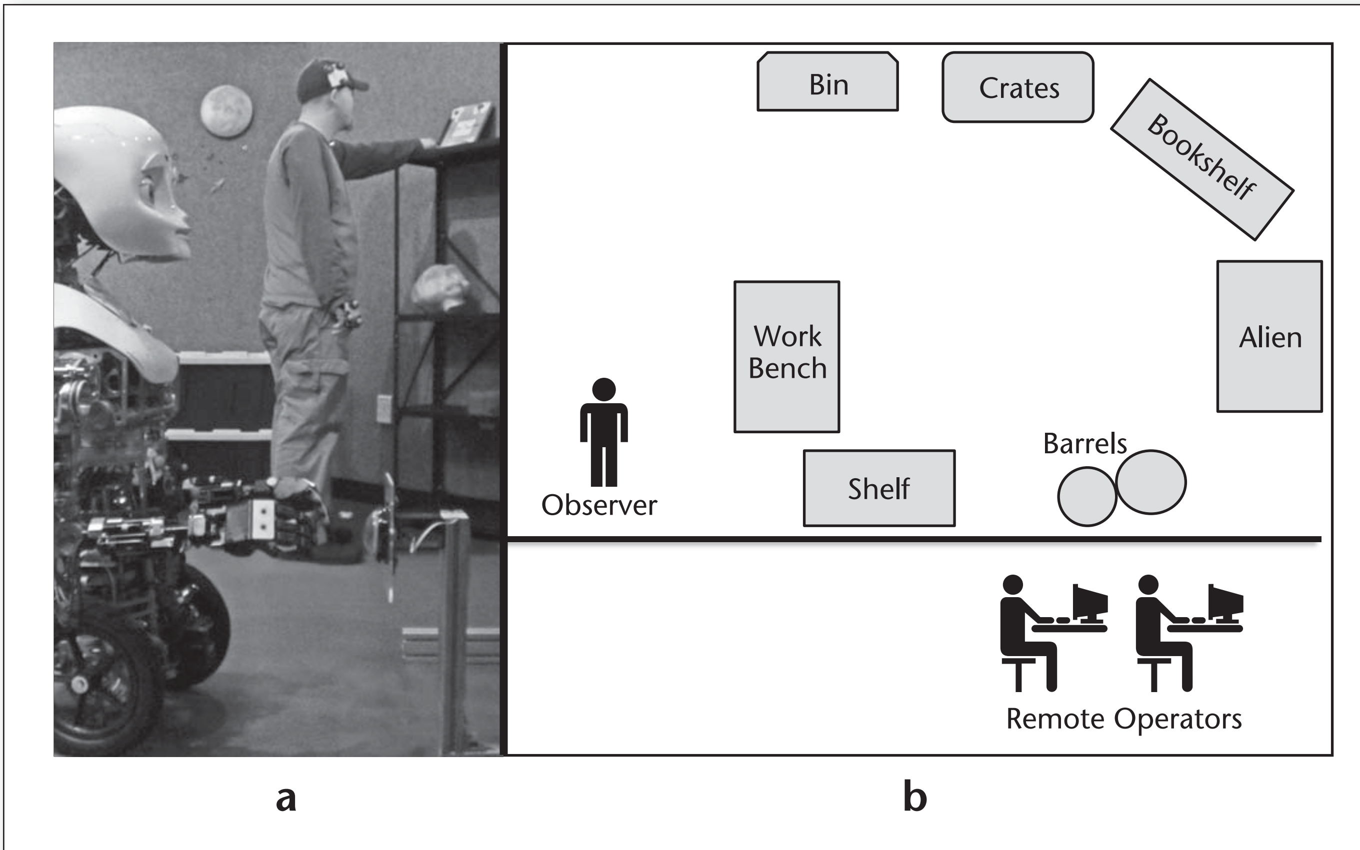

Figure 2. The MDS Robot Platform (a) and Study Setup at the Museum of Science (b).

As a result of the dialogue exchange, both players become physically colocated on the weight sensor. This example highlights the interleaved nature of dialogue and actions within the data set. On average, players exchanged 19.8 utterances per game, 10.6 in the robot role and 9.2 in the astronaut role. We found that first-time players engaged in dialogue more frequently, while players who were already experienced at the task focused on task-specific actions. A total of 708 unique phrases were recorded, most of which occurred only a single time in the data set.

The dialogue corpus highlights some of the benefits and challenges of crowd-sourcing interactive behaviors through online games. The dialogue data set is extremely diverse, far more so than any hand-coded system of comparable scope. This diversity can be leveraged to enable the robot to understand and contextually relate a broad range of phrases, as well as to produce varied, more humanlike dialogue. The challenge introduced by this data-gathering technique is one that's shared by most crowd-sourcing applications — the issue of data quality. Player dialogue frequently included topics that are not relevant to the task or are inappropriate in a real-world context. Examples included discussion of computer interfaces ("the font is really big," "click on the box"), personal comments ("hi dad!"), Internet slang ("lmao"), profanity, and dialogue in other languages, including Spanish, Portuguese, and Chinese. While our current solution to this problem is to filter the data manually, we anticipate that automated solutions can be developed in the future, possibly through additional crowd sourcing.

The retrieval of different items provided different degrees of challenge to the players, leading to several commonly observed patterns of behavior. For example, the majority of players first picked up those items that were in clearly visible locations and could be retrieved individually (that is, the canister and the journal), delaying the retrieval of collaborative items. Somewhat unexpectedly, we also found that only 57 percent of player pairs successfully collected all five items in the duration of the game (86 percent collected three or more). Of the five items, in 75 percent of games the last item to be retrieved was the sample box, the item that was most difficult to find and required the greatest degree of collaboration and communication between players. Furthermore, of the games in which players collected four items and missed only one, 89 percent were missing the sample box. All of these behavioral patterns affect the ordering, density, and distribution of data across the range of recorded player behaviors. In later sections we discuss how these patterns are reproduced by the robot in real-world experiments.

Transfer to Physical World

The robot used in this research is the MDS platform, which combines a mobile base with a socially expressive face and two dexterous hands that provide the capability to grasp and lift objects (figure 2a). 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 through a wearable microphone. Speech recognition utilized Sphinx 4,[1] with a grammar file generated from the online corpus. Speech synthesis and synchronized jaw movements were generated through the Cereproc software.[2] Navigation was implemented using a standard A* algorithm. Manipulation of objects was teleoperated by a hidden operator using a standard off-the-shelf 6DOF Connexion Mouse.

The Environment

In order to evaluate the behavior of the physical MDS robot in the collaborative task, we recreated the Mars Escape environment at the Boston Museum of Science. The physical setup of the space, shown on the right side of figure 2, 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 on 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. The left side of figure 2 shows the robot reaching for the button that activates the moving platform while the human participant retrieves the journal from the shelf.

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 and provide a degree of environmental awareness comparable to that of a human. The Vicon system was used to track the position of the robot, the objects in the environment, and the participant's head and right hand in real time using lightweight reflective markers attached to object surfaces.

Since the collaborative task presented many challenges in terms of robot sensing and mobility, a remote operator monitored the robot's progress from a hidden location and was able to intervene if necessary. Possible interventions included correcting the sensor state (for example, when a person picked up an object, but long-term occlusion prevented the system from recording the event), and overriding the autonomous behavior of the robot by selecting a different action. Examples of overrides are discussed in the evaluation section.

Data Transfer

In the transfer of a virtual character to an embodied agent, a one-to-one mapping must be created between the referent objects, the action set, and the world state of both worlds. The action set of the robot contained the same six actions as the online game — *pickup, put down, go to, use, look at,* and *analyze* — plus the additional action, *speak*, for dialogue. Physically impossible actions, such as *pick up astronaut*, were ignored.

All mission-critical objects had equivalent counterparts between the virtual and physical worlds. The robot's state was modeled using the following 13 features: the last robot action, the last astronaut action, the last robot spoken phrase, the last astronaut spoken phrase, object held by astronaut, object held by robot, and the area location (for example, center, near shelf, near toxins, and so on) of the astronaut, robot, journal, alien, chip, canister, and sample box. This information enabled tracking of recent events in terms of agent actions and dialogue, while also maintaining a coarse long-term history based on current object locations (for example, which objects have already been retrieved). The continuous movement of characters in the online game was discretized into regions in order to allow for generalization across similar actions.

Memory-Based Test Condition

Our goal is to leverage the corpus of interaction data describing the movements, actions, and spoken dialogue of players in the virtual world, to generate contextually correct social and task-oriented robot behaviors in the physical 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 (Kolodner 1993, Aamodt and Plaza 1994).

Case-based reasoning utilizes a library of past experiences (cases) to solve new problems by finding 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 (Micarelli, Panzieri, and Sansonetti 2007) and autonomous robot soccer (Ros et al. 2009). 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 physical world. Our case library contains only examples recorded in the virtual world. Methods for augmenting this data set with new examples from the physical world will be explored in future work.

Using the state representation described in the previous section, our data set resulted in 82,479 unique cases. For case storage and retrieval we utilized the open source FreeCBR software package.[3] 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. The case library includes all actions (both physical and speech) equally, enabling the same action selection mechanism to perform both behaviors.

Scripted Test Condition

Our second test condition, based on scripted Wizard-of-Oz control, represents hand-coded rule sets that are frequently found in robotic applications. In the scripted condition, a hidden operator teleoperates the robot through a preset sequence of dialogue and actions following a set script. The following example shows the first few actions of the script:

R: "Lets collect these items and get out of here."

R: "The toxic waste barrels look dangerous. I will go get the canister."

R: "Can you get the journal?" [go to toxic waste barrel] [pickup canister] [go to container] [drop canister]

R: "I'll go get the alien"

The scripted condition represents a precoded, open-loop behavior typical of many automated interactive systems. In the scripted condition the robot clearly states its intentions as it moves between different elements of the task. However, the robot is unable to adapt or respond to the speech and actions of the human participant. In our evaluation, we compare participant responses to the scripted and memory-based conditions, and show that although the scripted condition was rated more positively with respect to communication, it was rated significantly lower than the memory-based approach in performing as part of a team.

User Study Setup

We recruited 44 museum visitors as participants, none of whom had previous experience with humanoid robots. Data from 13 trials was thrown out due to robot error (dead battery, motor slippage, or critical perception problem), leaving 15 participants for the scripted condition (6 male, 9 female, 12–36 years of age), and 16 for the CBR condition (10 male, 6 female, 14–35 years of age). Before the start of the experiment participants were informed that the robot could understand speech, move around, pick up objects, and use its sensors to locate organic life-forms. They were also instructed that the toxins were hazardous to humans and should be avoided. The study continued until the team retrieved all five objects. Following the study, participants were asked to fill out the same questionnaire as in the online game.

Evaluation and Comparison

The behavior of museum 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 share of the task. A similar number of participants took a passive role, expected the robot to take charge and to provide all the answers. Below we present and discuss example interactions and then present a quantitative comparison of the experimental conditions.

The following example shows a hesitant participant (A), who checks with the robot (R) before performing any action:

A: "Do you want me to pick up the alien?"

R: "Climb up there, I can't go" [participant picks up alien]

A: "Do you want to put it in the bucket?"

R: "Please put the object in the bucket" [participant puts object in bucket]

Similar behaviors were observed in several other participants. We note that although the alien is successfully retrieved, the robot's response to the first question requires some interpretation. In fact, the "Climb up there …" statement was originally crowd-sourced from an online player in reference to the journal, not the alien. Reused by CBR in this context, the robot's phrase communicates the necessary information, but in an unnatural manner. The response given to the second question is more appropriate.

Another dialogue exchange, this time initiated by the robot, shows how successful dialogue paired

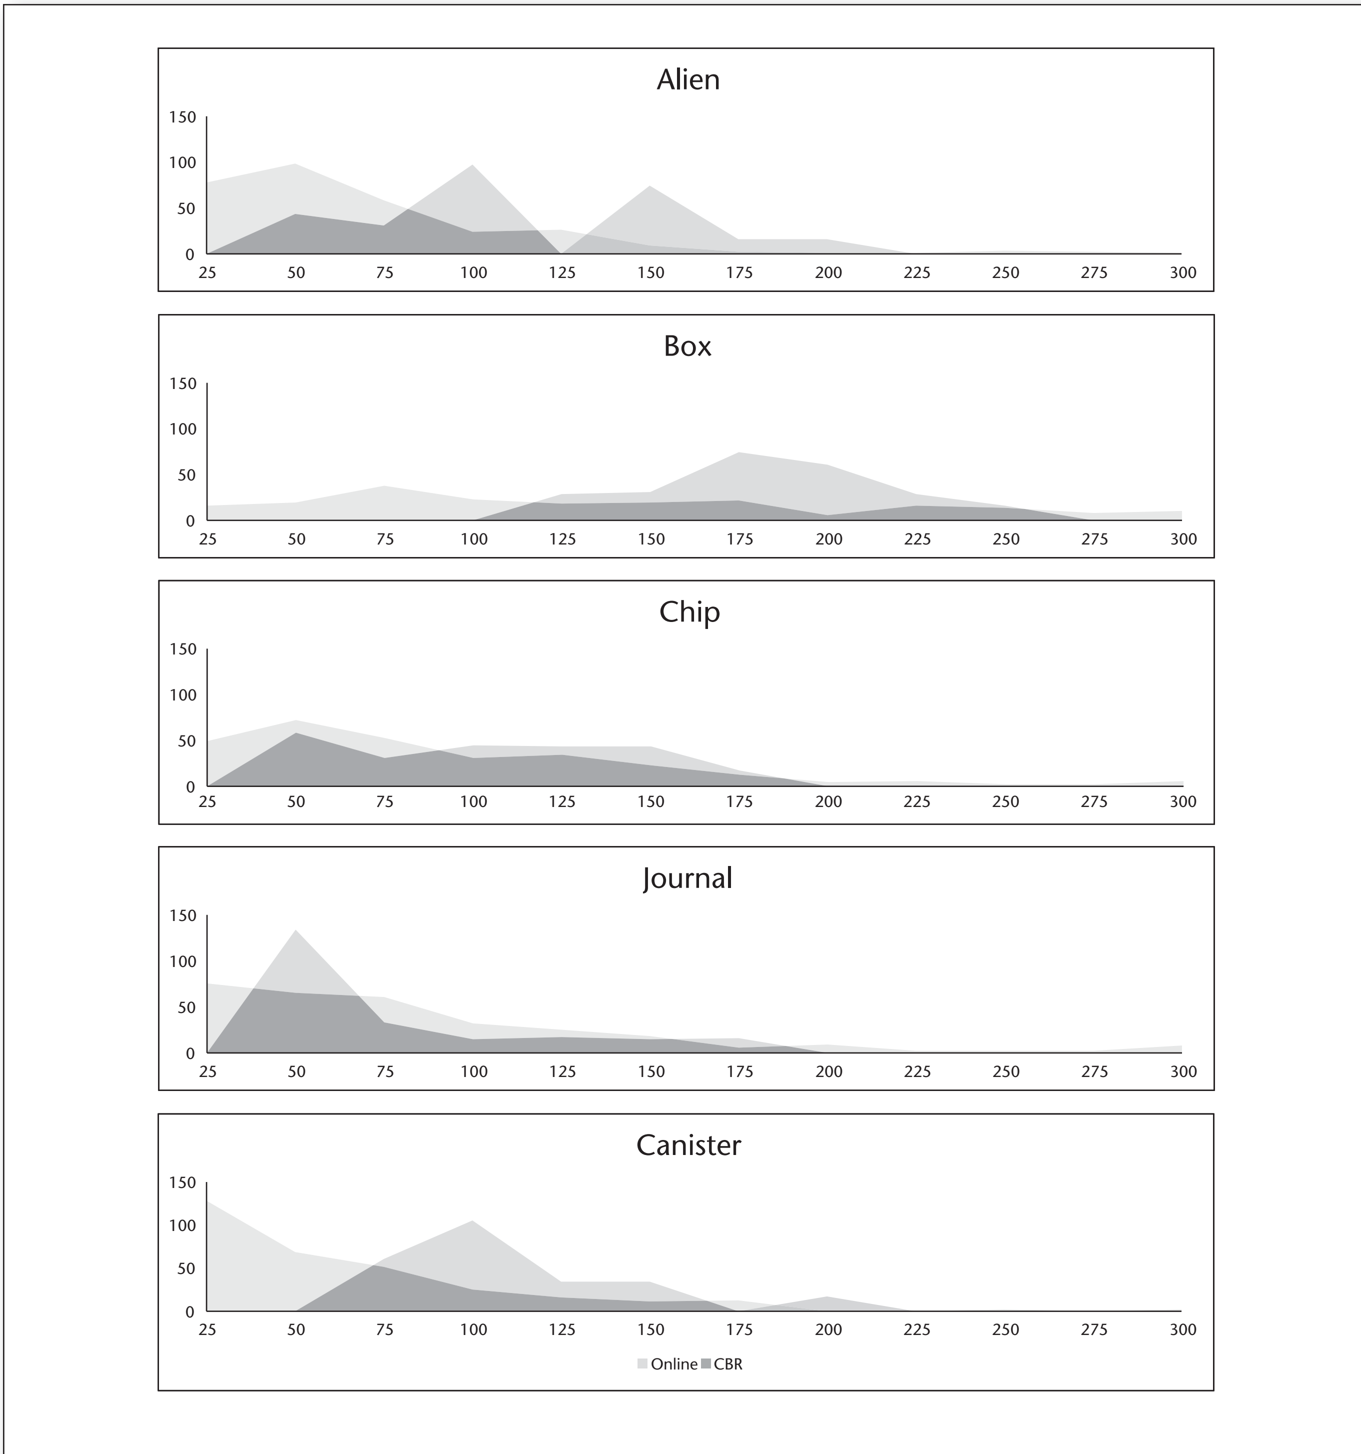

Figure 3. Comparison of Object Collection Times.

with social cues fosters human-robot collaboration and leads to successful retrieval of the sample box:

R: "Check the little box"

A: "What little box?"

R: "The object is in bin 56" [robot directs gaze at the shelf of boxes] [participant follows gaze and realizes the object location]

A: "Oh okay" [participant retrieves sample box]

Finally, we present an example in which several robot errors result in a confusing situation:

[robot picks up canister, moves toward participant while looking at him]

[canister becomes occluded from the perception system]

R: "I gave you the canister"

A: "The bucket is over there." [points]

R: "I gave you the canister" A: "Oh you want to give it to me?"

R: "Could you pick up the biohazard waste?" [Participant takes canister out of robots hand]

R: "Could you pick up the biohazard waste?"

A: "You want me to go over there and pick it up?" [participant picks up barrel] [robot moves to scale] [participant follows robot, places barrel on scale]

A number of factors played a role in this situation. First, due to perceptual error, the canister disappears from the robot's grasp. The phrase "I gave you the canister," which was originally used online when handing an object to the astronaut, is then matched by CBR as the closest relevant event in the interaction corpus. The participant, not knowing about the perception error, attempts to interpret the robot's intentions and take the canister. The interaction concerning biohazard waste is then retrieved by CBR from a different log in which the players consider placing a barrel on the scale. In the original online scenario the human player answers "it's too toxic" and does not retrieve the barrel. Interestingly, in the real-world scenario most people treated the robot as an expert and followed its requests even if they went against the guidelines given in the study.

Action Order

Figure 3 plots the distribution of times at which items were collected by participants in the online game and the memory-based study condition. The x-axis shows the elapsed time in the trial in seconds. We note several interesting patterns. In the online data, the sample box is typically retrieved much later than the other four items, with significantly fewer total successful retrievals. The chip, the other item that required collaboration between players, also shows a distribution skewed further along the time line than the easily accessible canister, journal, and alien.

Comparing the online data to the memory-based condition in the museum, 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 due to the fact that in the physical 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, driven by the crowd-sourced data set, also play a significant 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.

Level of Autonomy

In the memory-based condition, the robot took an average of 24.0 actions (4.4 utterances) to complete the task, compared to 25 (8 utterances) in the scripted condition. An average of 64.1 ± 4.4 percent of all robot actions in the memory-based condition were selected autonomously, as defined by the number of autonomous actions divided by the total number of actions. The remaining actions were manually triggered by an operator as an override. Unsurprisingly, locating the sample box proved to be a significant challenge for the human-robot teams, resulting in the greatest 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 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). Approximately half of the participants appropriately responded by picking up the sample box item, while others 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 redirect the robot back to the shelf to repeat its instructions. The second most common override instruction was in relation to repeat a failed attempt to pick up the canister.

Survey Results

Figure 4 presents a comparison of the survey results for the online, scripted, and memory-based conditions. Note that for clarity of presentation, the 5-point Likert scale has been collapsed to 3 categories by combining the "strongly disagree"/"dis-

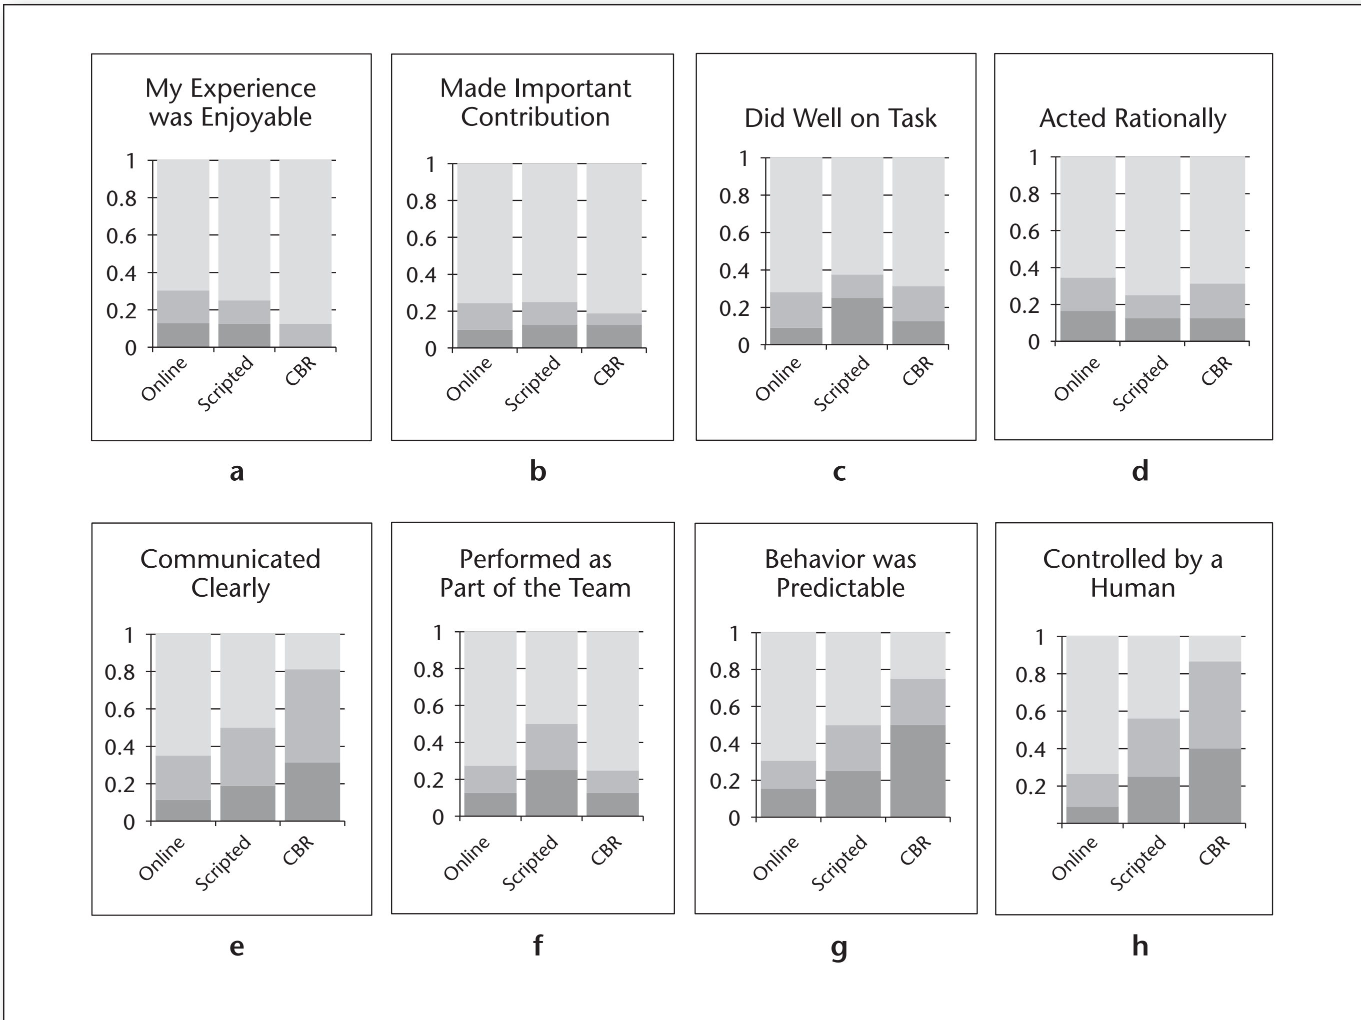

Figure 4. Comparison of the Results of the Eight Survey Questions between the Online, Scripted, and Case-Based Reasoning Conditions.

Questions 1–8 correspond to figures (a)–(h), respectively.

agree" categories (1–2), and the "strongly agree" / "agree" categories (4–5).

The survey shows that the vast majority of participants enjoyed taking part in the study, especially in the real-world scenarios. In all three study conditions, participants agree that the robot made an important contribution to the task and acted rationally (b–c). In fact, although 33 percent of participants reported uncertainty or disagreement with the statement that the robot's actions were rational, the same numbers are reported for the rationality of human players in the online game! As a result, we view the autonomous memory-based approach as highly successful in scoring on par with both human and scripted behavior in these critical measures.

In the memory-based condition, the robot is rated similarly to online human partners with respect to its performance at the overall task (d); both the online and memory-based conditions score slightly higher than the scripted condition. A similar, and even more pronounced pattern of responses is observed with respect to whether the robot performed as part of a team (f). These results highlight that the nonadaptable nature of a predefined policy adversely affects both task performance and collaboration, supporting the need for further research into adaptable systems.

Finally, human partners in the online game significantly outperformed both robot conditions in three measures: clear communication (e), predictable behavior (g), and being human (h). In all three measures, the scripted condition scored better than the memory-based condition. However, despite the ordered nature of the script, a significant number of participants found the robot's behavior difficult to predict, which highlights the potential need for additional social cues to communicate intentions more clearly. We believe that the measure of whether the robot was autonomous or controlled by a human is strongly dependent on the clarity of communication and predictability of actions. We anticipate that improvements in these areas will lead to perceptions of more humanlike behavior for the entire system. Interestingly, we note that more than 20 percent of online players were uncertain as to whether their teammate was human or not.

Discussion

We view the previously described approach as one that sets a baseline for what can be achieved through data-driven methods directly mapping from the virtual to the physical world. In this section we discuss lessons learned through the course of this study, important directions for future research, and related work.

Obtaining the Data Set

Arguably, the greatest challenge that any crowd-sourced system must address is to ensure that the sourced data accurately represents, and effectively maps to, the goal task. In the collection of an online data set that will be transferred to the physical world, game designers must ensure that the domain attributes that are most critical to the success of the real-world task are accurately represented in the online game. For example, due to its focus on collaboration, the Mars Escape game accurately modeled the movement speeds and reach abilities of the characters, which we believe aided in the transfer between domains.

Additionally, games must be designed with care to incentivize the desired behaviors, which can prove to be a challenge. The Mars Escape game included time pressure, in the form of a time limit, in order to encourage players to focus on performing task-relevant actions. While successful (most players stayed on task), an undesired side effect of this feature was that players attempted to minimize dialogue engagement in order to finish the task more quickly. This effect was not intended and could be counteracted in the future by removing the time pressure or allowing players to communicate through a hands-free head set.

Transfer to the Physical World

The biggest challenges we observed with respect to the transfer of data from the virtual to the physical world were in relation to dialogue. One challenge is filtering out phrases that are not only off topic, but highly unlikely to be encountered in the real world (for example, "this font is big"). While crowd sourcing is the source of this problem, it is also a likely source of a future automated solution. Anonymous reviewers contracted through services such as Amazon's Mechanical Turk[4] could be used to flag and remove inappropriate phrases in the future.

Another challenge is that of generalization — how can data crowd-sourced for one application be applied for a highly related similar task? For example, what if we were to perform the same search and retrieval task but in a different setting, replacing the book on top of the boxes with a vase on top of a table? Although both the dialogue and physical actions would need to be adapted to the new task, changes to dialogue pose a greater challenge due to its diversity. Direct search and replace methods substituting one word for another are unlikely to succeed due to both the diversity of human speech (for example, terms used to describe the alien included *bird*, *green thing*, and *toy*, among others) and the possible need for changes in the grammatical structure of the resulting phrase. More advanced techniques for generalization and transfer between domains remain an interesting topic for future work.

More generally, simulated worlds are far simpler and more predictable than their real counterparts, making the transfer of policies between these representations difficult. Potential methods for addressing this problem include building more accurate simulations of the physical world by modeling the stochasticity of real-world environments, developing base behavior models through crowd sourcing in simulation and then adaptively correcting them in the physical world, and learning the complete task in situ in the real world. Our own work so far has explored crowd-sourcing base policies through simulation, and we are in the process of exploring techniques for refining these policies through human instruction in the real world. In situ crowd-sourced learning has previously been explored by Crick et al. (2011), who demonstrated a technique for collecting a diverse set of demonstrations by enabling online users to directly teleoperate a robot through a maze.

Finally, the presented work focuses on a data-driven approach, exploring what can be achieved with a direct mapping from a virtual to a physical world. A natural future step is to integrate crowd-sourced data into existing systems, such as higher-order planning and cognitive architectures.

Related Work

Research on crowd-sourcing HRI derives ideas from many related works in machine learning and robotics. Projects leveraging Internet users to collect large-scale data corpora for different applications include Soylent (Bernstein et al. 2010), Open-Mind (Singh et al. 2010), OpenStreetMap (Haklay and Weber 2008), and the ESP Game (von Ahn and Dabbish 2008). These projects range from collecting common sense knowledge to providing copy editing help for writers. Our work is particularly motivated by two projects, Games With A Purpose, which aims to address computational problems through the creation of online games (von Ahn and Dabbish 2008), and the Restaurant Game, in which data collected from thousands of players in an online game is used to create an automated data-driven behavior and dialogue authoring system (Orkin and Roy 2007; 2009). Our work differs from that of Orkin and Roy most significantly in that data collection is performed in a fundamentally different domain than the one in which it is applied.

Within robotics, crowd sourcing has been applied to many subfields including vision (Sorokin and Forsyth 2008), grasping (Sorokin et al. 2010), and navigation (Crick et al. 2011). On a smaller scale, data collection from nonexpert users, both in the real world and in simulation, has been used in learning from demonstration research, such as interactive reinforcement learning (Thomaz and Breazeal 2007), the TAMER learning framework (Knox and Stone 2009) and confidence-based autonomy (Chernova and Veloso 2009). All of these approaches differ from our work in that they utilize humans as data sources for policy learning of narrowly defined tasks, whereas our goal is to model and create systems capable of natural humanlike interaction. A closely related work in the area of robotics is that of Kollar et al. (2010), in which a large corpus of task-constrained language is used to develop a robotic system capable of following natural language instruction. This work is highly relevant but does not extend to collaborative and social aspects of robot interaction and behavior.

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. Our work presents a novel approach to generating task-specific social behaviors based on crowdsourcing human-robot interaction in virtual worlds. We show that crowd-sourced interaction data describing the movements, actions and spoken dialogue of players in the virtual world can be used to generate contextually correct social and task-oriented behaviors for a robot operating in a real-world environment, allowing the robot to exhibit similar patterns of behavior to those observed in online players.

This is the first study that we are aware of examining large-scale online crowd sourcing for human-robot interaction. While many existing approaches have explored learning in a virtual world, in real-world environments, and through games, we are unaware of any work that examines the transfer of social and collaborative robot behaviors between virtual to physical worlds at this scale. 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. In comparison to a scripted interaction, participants reported the autonomous robot to be a better team member, but poor at communication. 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 crowd sourcing in virtual worlds has the potential to become a powerful tool in human-robot interaction research.

Acknowledgements

This work was supported by Microsoft Research and the Office of Naval Research Award Numbers N000140910112 and N000140710749. We would also like to thank Dan Noren and the staff of Cahners Computer Place at the Boston Museum of Science for their support and assistance.

Notes

1. See cmusphinx.sourceforge.net.
2. See www.cereproc.com.
3. Johanson, L. 2010. FreeCBR. (freecbr.sourceforge.net).
4. See www.mturk.com.

References

Aamodt, A., and Plaza, E. 1994. Case-Based Reasoning; Foundational Issues, Methodological Variations, and System Approaches. *AI Communications* 7(1): 39–59.

Bernstein, M.; Little, G.; Miller, R.; Hartmann, B.; Ackerman, M.; Karger, D.; Crowell, D.; and Panovich, K. 2010. Soylent: A Word Processor with a Crowd Inside. In *Proceedings of the 23nd Annual ACM Symposium on User Interface Software and Technology*, 313–322. New York: Association for Computing Machinery.

Chernova, S., and Veloso, M. 2009. Interactive Policy Learning through Confidence-Based Autonomy. *Journal of Artificial Inelligence Research* 34(1): 1–25.

Crick, C.; Osentoski, S.; Jay, G.; and Jenkins, O. C. 2011. Human and Robot Perception in Large-Scale Learning from Demonstration. In *Proceedings of the 6th International Conference on Human-Robot Interaction*, 339–346. New York: Association for Computing Machinery.

Haklay, M., and Weber, P. 2008. OpenStreetMap: User-Generated Street Maps. *IEEE Pervasive Computing* 7(4): 12–18.

Knox, W., and Stone, P. 2009. Interactively Shaping Agents via Human Reinforcement: The TAMER Framework. In *Proceedings of the Fifth International Conference on Knowledge Capture*, 9–16. New York: Association for Computing Machinery.

Kollar, T.; Tellex, S.; Roy, D.; and Roy, N. 2010. Grounding Verbs of Motion in Natural Language Commands to Robots. In *Proceedings of the 12th International Symposium on Experimental Robotics*. Berlin: Springer.

Kolodner, J. 1993. *Case-Based Reasoning*. San Francisco: Morgan Kaufmann Publishers Inc.

Micarelli, A.; Panzieri, S.; and Sansonetti, G. 2007. Case-Based Reasoning in Robot Indoor Navigation. In *Proceedings of the Seventh International Conference on Case-Based Reasoning: Case-Based Reasoning Research and Development*, 284–298. Berlin: Springer-Verlag.

Orkin, J., and Roy, D. 2007. The Restaurant Game: Learning Social Behavior and Language from Thousands of Players Online. *Journal of Game Development* 3(1): 39–60.

Orkin, J., and Roy, D. 2009. Automatic Learning and Generation of Social Behavior from Collective Human Gameplay. In *Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems*, 385–392. Richland, SC: International Foundation for Autonomous Agenets and Multiagent Systems.

Ros, R.; Arcos, J. L.; Lopez de Mantaras, R.; and Veloso, M. 2009. A Case-Based Approach for Coordinated Action Selection in Robot Soccer. *Artificial Intelligence* 173(9–10): 1014–1039.

Singh, P.; Lin, T.; Mueller, E.; Lim, G.; Perkins, T.; and Li Zhu, W. 2010. Open Mind Common Sense: Knowledge Acquisition from the General Public. *On the Move to Meaningful Internet Systems 2002: CoopIS, DOA, and ODBASE* 1223–1237. Berlin: Springer.

Sorokin, A., and Forsyth, D. 2008. Utility Data Annotation with Amazon's Mechanical Turk. In *Proceedings of the 2008 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, 1–8. Piscataway, NJ: Institute of Electrical and Electronics Engineers.

Sorokin, A.; Berenson, D.; Srinivasa, S.; and Hebert, M. 2010. People Helping Robots Helping People: Crowdsourcing for Grasping Novel Objects. In *Proceedings of the 19th IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2117–2122. Piscataway, NJ: Institute of Electrical and Electronics Engineers.

Thomaz, A., and Breazeal, C. 2007. Asymmetric Interpretations of Positive and Negative Human Feedback for a Social Learning Agent. In *Proceedings of the 16th IEEE International Symposium on Robot and Human-Interactive Communication*, 720–725. Institute of Electrical and Electronics Engineers.

von Ahn, L., and Dabbish, L. 2008. Designing Games with a Purpose. *Communications of the ACM* 51(8): 58–67.

Sonia Chernova is an assistant professor at Worcester Polytechnic Institute (WPI). Her research interests lie in interactive machine learning, adjustable autonomy, human computation, and human-robot interaction. Her work focuses on the development algorithms that enable robots to learn through social interaction with humans. Chernova received her Ph.D. from Carnegie Mellon University in 2009. Prior to joining WPI she worked as a postdoc at the MIT Media Lab with Cynthia Breazeal.

Nick DePalma is a student in the MIT Media Lab. He holds a BS and an MS in computer science from the Georgia Institute of Technology with focuses in graphics, user interface, artificial intelligence, and robotics. He has worked in the game and computer vision industry and intends to build helpful, entertaining, and enjoyable cooperative robots for those that need them most.

Cynthia Breazeal is an associate professor at the MIT Media Lab where she founded and directs the Personal Robots Group. She is a pioneer of social robotics and human robot interaction. She is author of *Designing Sociable Robots* and has published more than 100 peer-reviewed articles on these topics. Her research program develops personal robots with interpersonal skills that enable them to work and learn collaboratively with people. Recent work focuses on socially assistive robots targeting applications in education and health that require long-term interaction. She received her Sc.D. in electrical engineering and computer science from the Massachusetts Institute of Technology in 2000.