Exploring human-robot interaction through telepresence board games

— **Source link** ↗

Min Xin, Ehud Sharlin

Institutions: University of Calgary

Topics: Video game, Human–robot interaction, Game theory, User interface and Autonomous robot

Related papers:

- Linking physical and virtual interaction spaces

- Towards measuring the quality of interaction: communication through telepresence robots

- Artificial companion: Building a impacting relation

- The Role of the Vocal Stream in Telepresence Communication

- Intelligent virtual agents in collaborative scenarios

Exploring Human-Robot Interaction Through Telepresence Board Games

Min Xin and Ehud Sharlin

Interactions Laboratory, Department of Computer Science, University of Calgary
2500 University Drive NW, Calgary, Alberta, Canada, T2N 1N4
{xinm, ehud}@cpsc.ucalgary.ca

Abstract. This paper presents an experimental test bed for exploring and evaluating human-robot interaction (HRI). Our system is designed around the concept of playing board games involving collaboration between humans and robots in a shared physical environment. Unlike the classic human-versus-machine situation often established in computer-based board games, our test bed takes advantage of the rich interaction opportunities that arise when humans and robots play collaboratively as a team. To facilitate interaction within a shared physical environment, our game is played on a large checkerboard where human and robotic players can be situated and play as game pieces. With meaningful interaction occurring within this controlled setup, various aspects of human-robot interaction can be easily explored and evaluated such as interaction methods and robot behaviour. In this paper we present our test bed which uses a telepresence interface for playing the game and the results of a user study demonstrating the sensitivity of our system in assessing the effect of different robot behaviours on users.

1 Introduction

How will humans, intelligent computers, and robots coexist and interact? This question has motivated thinkers and writers for a long time, with ideas ranging from Licklider's Man-Computer Symbiosis Partnership [1] and Moravec's evolution of new intelligent superior species [2] to Philip Dick's masters-slaves society led by mistrust and fear [3]. Current researchers and designers working with human-robot interaction no longer see robots as fully-controlled subordinates but rather as peers and colleagues with a variety of social and emotional abilities (see for example [4] and [5]). It is logical that humans will find future intelligent robots more helpful if the robots act according to behavioural patterns that humans can understand and relate to.

With the growing demand and availability of interactive robots for varying applications, there is a need for effective and rapid prototyping and evaluation of human-robot interfaces. A considerable amount of research in human-robot interaction focuses on investigating interaction for particular applications such as entertainment robots, but the results gathered may not generalize well to other applications, and there are clearly inherent aspects of interaction shared by most applications that are not explored. Our goal is to construct a controlled environment and a manageable set of tasks that can serve as an interaction metaphor which encompasses the common

interactive qualities of major human-robot interaction applications. Our idea and solution is a collaborative board game involving humans and robots played in a shared physical environment. This setup serves as a test bed where various interfaces and robot behaviours can be developed to facilitate game play, and in turn, user evaluations can be performed to test their effectiveness.

Fig. 1. *Sheep and Wolves*: the AIBO robotic wolves closing in on the mixed reality sheep

Following this concept, we developed our test bed, *Sheep and Wolves* (Fig. 1), based on a classic checkerboard game but adapted to allow humans and robots to play together as a team of four wolves hunting a single sheep in an attempt to surround it. The game is played on a large checkerboard where humans, robots, and virtual entities act as game pieces, enabling a large variety of scenarios. Virtual entities were included in the game, using mixed reality, in order to highlight one of the robots' main advantages over humans: their ability to function in both the physical and virtual realms. Humans must rely on the robots' "senses" when it comes to the virtual entities, but for the robots, the virtual entities are as real as the physical components of the task.

Along with designing the test bed, we have also implemented a telepresence and mixed reality interface for playing the game and modeled two robot behaviours that are evaluated using *Sheep and Wolves*. We performed user evaluations asking human participants to play two games using the prototype interface. In one game the robots are always supportive and obedient, and in the other game the robots behave negatively and always ignore input from their human team-mate. The responses from the human participants were found to be sensitive to the contrasting robotic behaviours.

2   Related Work

We first present an overview of research in human-robot interaction (HRI) which serves as a general motivation for the field. Then we provide rationale for the design of *Sheep and Wolves*. Finally we briefly outline two interaction techniques, namely telepresence and mixed reality, which are implemented in our prototype interface for playing the game.

Although the fields of human-computer interaction (HCI) and robotics have generated significant interest and made substantial progress, in the past there has been less attention paid to human-robot interaction (HRI). Certainly, arguments can be made that robots such as mechanical arms used in manufacturing can be operated with similar techniques which have already been explored in human-computer interaction. However, as increasingly intelligent and capable autonomous robots come into existence, traditional human-computer interaction approaches become less applicable, and special attention need to be given to the unique requirements and advantages embodied in intelligent autonomous robots [6]. People intuitively perceive robots to be more anthropomorphic than other computing systems [6]. This hints at the potential for more intimate and natural forms of interaction between humans and robots. Cynthia Breazeal, the head of MIT's Robotic Life Group, designed Kismet [5], an expressive anthropomorphic robot able to convey emotions and provide feedback, allowing humans to engage in natural social interaction. Kismet utilizes facial expression, body posture, gesture, gaze direction, and voice to intuitively communicate with humans. Another unique characteristic of autonomous robots is their ability to learn and make decisions based on information gathered from the physical environment. Many robots designed for entertainment such as Sony's AIBO^{TM} robotic dogs support a cognitive learning model which enables the robot to acknowledge various forms of human and environmental input and mold its behaviour accordingly. The Robotic Life Group's Leonardo [7], a life-like robot designed for social interaction, can interpret gestures and facial expressions from humans as a communication method for learning how to play games. User interaction with such autonomous robots tends to be richer and more intuitive than traditional human-computer interaction paradigms of clicking on icons or opening windows. Furthermore, mobile autonomous robots often bring interaction closer to the physical context of humans, allowing information and subtle social interaction cues to be readily exchanged. NASA's Robotnaut [8], a mobile autonomous humanoid robot, is being developed in an attempt to create future robot astronauts that are able to work alongside human astronauts in space. Efforts such as Kismet, Leonardo, and Robonaut are the prelude to a fascinating future for the field of human-robot interaction. Attitudes for designing robots are already shifting from a "robots as tools" approach to a "robots as partners" outlook.

Simulated computer agents playing games such as chess or checkers with or against humans is a familiar concept (see for example [9]). However, interaction and collaboration between humans and robots within a physical game environment is rare. One example is Carnegie Mellon University's Cognitive Robotics [10], which suggests means of implementing more involved physical interaction between robots and games, presenting a robot-based tic-tac-toe game in which the robot can move game pieces on a physical board. However the potential for human-robot interaction is still limited. Meaningful interaction between humans and robots within a game application can be enhanced by requiring humans and robots to play on the same team within a shared physical environment instead of against each other on a computer. The idea originates from using robots for search and rescue where performing collaborative tasks can be critical. Since human ability, artificial intelligence, and computational ability can be fairly balanced within a controlled game environment, it is conceivable to construct realistic scenarios where humans and robots collaborate as equals.

NASA developed the Web Interface for Telescience (WITS) [11] Software which linked a vehicle for Martian travel to Internet users. This allowed a group of high school students to actively participate in assisting researchers operate the vehicle during a field test. University of Southern California's Telegarden [12] enabled Internet users to operate a remote robotic arm centered in a garden in order to water and care for the plants inside. These projects demonstrate the power of telepresence in encouraging remote collaboration, active and assisted learning, and developing a sense of virtual partnership. With autonomous robots, the benefits of telepresence can be extended further. Arguments can be made that most current telepresence interaction techniques follow the "robots as tools" approach with users having to operate and control many mechanical aspects of the remote robot. Although the direct physical context is missing, previously mentioned interaction techniques based on the "robots as partners" perspective can still be applied by delivering video, sound, and other sensory and communication elements. The experience can be similar to existing interaction between humans online such as chatting using instant messaging programs, collaborating by voice in online games, and participating in video conferences. By exploring these interaction paradigms for telepresence, remote users can collaborate with a team of remote robots as participating members rather than superior operators having to control the entire team.

Mixed reality gives humans the ability to access valuable information processed and stored within digital entities in the immediate context of a physical environment. This is ideal as a human-robot interaction paradigm because many robots operate within the physical world and can also obtain and process digital information. Applications such as the MagicBook [13] or the Human Pacman [14] allow humans to visualize and interact with virtual digital entities by superimposing computer-generated graphics onto physical scenes. In the MagicBook for example users can view a computer-generated animation on how to build a chair as they browse through a physical instruction book using a hand-held display [13]. Naturally the concept of visualizing contextual supplementary information can also be applied to robots, allowing them to directly express implicit information such as thoughts or synthetic emotions.

3   Test Bed Design

Sheep and Wolves (Fig. 1) is a human-robot interaction test bed following the goal of constructing a controlled environment and tasks that will serve as an interaction metaphor for major human-robot interaction applications. The environment and tasks we have devised are based on a classic turn-based game played on a checkerboard. The game involves five game pieces, four of which are the wolves, and one is the sheep. The wolves start on one end of the checkerboard, and the sheep starts on the other. The team of wolves are only allowed to move one wolf forward diagonally by one square during their turn. The wolves' objective is to surround the sheep so it cannot make any legal moves. Meanwhile, the sheep is allowed to move forward or backward diagonally by one square during its turn. Its objective is to move from one end of the checkerboard to the other. Obviously, while the sheep is more flexible in its moves, the wolves' strength is in their numbers and ability to move as a pack.

We chose this game because it is simple yet able to support collaborative game play. The metaphor of the game can be extended to various applications where humans and robots are required to share information, opinions, and resources in order to effectively complete a task. By performing a collaborative task in a controlled physical game environment instead of complex real-world environments, we are able to focus on interaction. Also, since implementing artificial intelligence for checkerboard games is relatively simple, we are able to easily adjust the intelligence of the robots in order to develop varying believable robot behaviours.

In our game we have elected to use Sony's AIBO dogs as our robot participants. These fairly capable commercial robots allow us to rapidly build prototype interfaces for evaluation. For the physical environment of the game, we use a 264 cm (104 in) by 264 cm RolaBoard^{TM} with the standard black and white checkerboard pattern. Each square measures 33 cm (13 in) by 33 cm, providing sufficient room for an AIBO wolf to sit on or humans to stand on. This controlled shared space is ideal for robots to navigate in. The lines and corners of the checkerboard serve as readily available navigation markers for movement on the checkerboard, and camera calibration can also be achieved using corner points to allow for mixed reality interfaces and localization of humans on the checkerboard.

4   Game Implementation

In the setup for our telepresence and mixed reality game, all four wolves are represented by the AIBOs and the sheep is a virtual entity. The AIBOs physically rise, move, and sit down on the checkerboard to indicate position and movement of the wolves in the game. A human player controls a single AIBO wolf at a remote computer using a telepresence interface, personifying the robotic entity within the game. Other uncontrolled AIBO wolves are autonomous robot team-mates which the human player must collaborate with. Live video of the physical game environment from the controlled AIBO's point of view is provided to the remote human player, and mixed reality is utilized for visualizing the virtual sheep's location and moves. Winning the game as wolves requires teamwork: the human player has to provide suggestions to the team and consider propositions made by other team-mates in order to help the team reach intelligent decisions on the moves the team should make.

4.1   Checkerboard Traversal and Augmentations

One of our goals is to introduce physical elements into the board game. By playing the game on a large checkerboard we define a simple environment in which the robotic game entities can easily operate. As a result of the rules of the game, AIBO wolves are only required to traverse the checkerboard moving forward diagonally one square at each turn. This can be achieved using a simple localized vision algorithm without having to map the physical environment of the checkerboard. When an AIBO wolf is about to move, it stands up on all four legs with its snout facing straight down. Since the camera is located in the AIBO's snout, this posture provides a bird's eye view of the board which is also very limited due to the camera's field of view and the relative closeness of the camera to the checkerboard. This limited bird's eye view

of the checkerboard is actually ideal for a simple traversal algorithm since there is very little perspective distortion, and for each frame of video obtained by the AIBO in the stand-up posture, we have only several distinct cases to consider for localizing and orientating the AIBO. For our walking algorithm, we decided to use lines and corners as means of localization and determining orientation. Working only with low resolution greyscale image data, we extract lines from the images by first applying a low-pass filter and then performing a binary threshold to generate resulting images similar to the ones shown in Fig. 2 (left). Next, we search for line end points around the perimeters of the images by simply performing exclusive-or operations of the tested pixel with each of its right and bottom neighbours. From the extracted line end points, we derive the line segments present in the image. The case with two line end points is trivial. To correctly match three or four line end points, we simply consider all possible pairings and calculate the resulting angles between the two line segments. Since the bird's eye view of the checkerboard does not suffer from perspective distortion, line segments within the limited view must be orthogonal to each other. Therefore, we can exclude pairings of line segments which are not orthogonal. In frames where two line segments can be extracted, we can also determine the position of a corner point by simply calculating the intersection between the two line segments. Corner points which can be inside or outside of an image are used to localize the AIBO on the checkerboard. The angles between extracted lines and the vertical axis are used to align the AIBO in a proper position.

Fig. 2. Images of bird's eye views of the board (left) and extracted corner points (right)

In order to visualize the virtual sheep and demonstrate the application of mixed reality, we enhance the live video provided by the AIBO's camera by superimposing a computer generated 3D sheep onto the scene (Fig. 1). To achieve this, we set up an OpenGL viewing frustum based on the camera's field of view and focal length. In the scene, a rectangle is placed at a distant location from the camera looking down the z-axis. The size and aspect ratio of the rectangle is calculated using the field of view and focal length of the camera to ensure it covers the entire viewing volume when displayed. Frames of video received from the AIBO's camera are then texture mapped onto the distant rectangle to provide a video background for the virtual 3D sheep in the scene. As AIBOs move on the board, the exact positions of the AIBOs' cameras are unknown after each move. To place the virtual sheep within the correct viewing context of the video background, continuous camera calibration is required. We designate the center of the checkerboard as the origin of our world coordinate system. Then, by keeping track of the game entities on the physical board, we know approximately the position of the camera. Next, we fine-tune the calibration using high resolution image data from the camera. First, we extract the checkerboard corner points from the image (Fig. 2 right). This is accomplished using a corner detection

algorithm. In most cases, we can extract at least three accurate points close to the camera which we use to perform a simplified camera calibration. After obtaining the corner points we inverse project these 2D points into our 3D world coordinate system. This is possible because we know the *y*-axis values of these potential 3D points are all supposed to be 0. We then pair the inverse projected 3D corner points together in attempt to find either a potential horizontal or vertical edge of a square. After calculating the angle between the vector resulting from such an edge and the corresponding horizontal or vertical vector, we rotate the virtual scene around the y-axis by the calculated angle to correct misalignments caused by the AIBO not always facing exactly forward. We assume that the AIBO's camera does not require roll adjustment and that its height remains constant. Following this calibration procedure we are able to correctly superimpose the sheep on the live video most of the time. Challenging cases such as the loss of corner points due to occlusions and the introduction of false corner points created by a black AIBO sitting on a white square can result in imprecision.

4.2 Game Play

The game algorithm for both the sheep and the autonomous wolves are implemented based on the concept of searching for paths from the sheep to the other end of the checkerboard. If multiple paths are available, the sheep will move following the shortest path. Otherwise it will make a random move with a preference for moving forward instead of backward and moving toward the center instead of to the side. The non-human controlled members of the wolves will suggest the move which results in the longest available path for the sheep, or which will lead the sheep to a dead-end.

For each turn, the sheep or the wolf pack has sixty seconds to arrive at a decision for the move. As in the classic board game, the wolves win when the sheep can no longer move, and the sheep wins if it gets past the last wolf on its way to the other end of the board. At the end of the sheep's turn, autonomous AIBO wolves assess the game and start to make suggestions to the rest of the team. The human player can also communicate with the team using a text messaging interface, and other autonomous AIBOs provide either positive or negative feedback depending on whether the others' suggestions match their opinion. One advantage of forcing the human player to play the game from a single robot's perspective is the limitation of field of view. Without an overview of the checkerboard, the human player has to deal with vagueness, uncertainties and lack of information. Although we provide the player with the option to pan the head of the AIBO to further explore the checkerboard, situations where the human player is completely clueless as to what the next move should be can occur. This forces the human player to utilize suggestions from the autonomous AIBO wolves, allowing us to explore issues of trust between humans and robots.

To allow a human player to effectively control an AIBO wolf and interact with the rest of the team, we have devised an intuitive graphical user interface (Fig. 3). In the following paragraphs, the various parts of the interface will be outlined, and the motivations behind the design choices will be explained.

In the main area of the screen, live video of the game along with the virtual sheep is displayed. This allows the remote human player to see the physical board from the point of view of the controlled AIBO. The virtual sheep is visible to the player if it is occupying a square in the field of the view of the camera. At the bottom of the main

display, game information is provided, indicating what the game entities are doing (thinking or moving), whose turn it is, and the time remaining for making a decision.

On the top right of the interface, a game radar (Fig. 3) indicates the positions of the wolves relative to the edges of the checkerboard. Since our goal is to simulate search and rescue operations, we chose not to provide the human player with the position of the sheep and the grid of the checkerboard. This encourages the human player to actively interact with the physical environment of the checkerboard rather than utilizing the abstract radar to play the game. Each AIBO wolf is represented by a red dot. The AIBO wolf controlled by the human player is indicated with a blue ring around its dot. When an AIBO wolf moves, its dot will flash to indicate the movement. Displayed next to their corresponding dots are the AIBO wolves' nicknames. The nicknames along with their full names, Leonardo, Michelangelo, Donatello, and Raphael, can be used by the player to refer to a particular wolf in the game. For simplicity we designate the direction the wolves are initially facing as north, and therefore the green arrow in the radar always points north.
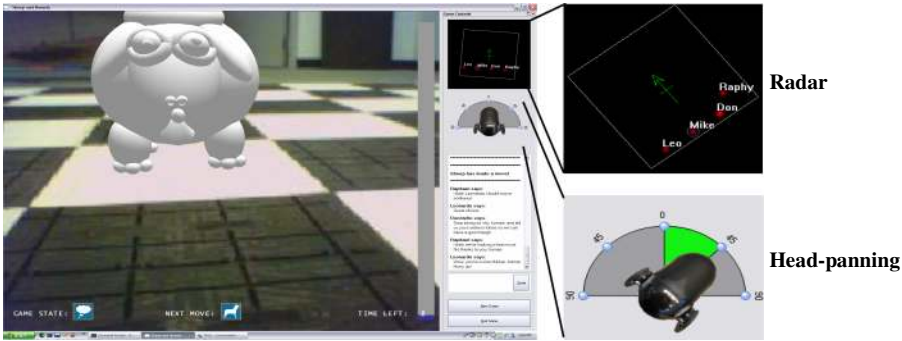

Fig. 3. Telepresence interface with mixed reality and text messaging

Underneath the game radar is the head-panning control (Fig. 3). Since the initial forward-facing view is limited we allow the human player to pan the head of the controlled AIBO 45° or 90° left or right (east or west). This feature can be used to explore the checkerboard, locate the sheep, observe other AIBO team-mates, or watch them move. When the player pans the head of the controlled AIBO, the game radar also rotates to match the orientation of this AIBO's current view and further assists in spatial awareness.

A crucial element of our design is the text messaging interface (Fig. 4). It allows the human player to communicate with the rest of the AIBO wolves following a familiar interaction paradigm. Although currently the richness of the conversation is limited, we believe that this interaction technique has potential in effectively engaging human users in active collaboration with robotic entities especially in telepresence applications because most human users are already familiar with instant messaging programs. In our game, conversation occurs amongst four team-mates. Due to the loss of context or the intended recipient of messages, effective communication can be difficult when the discussion is commencing at a rapid rate. To solve this problem,

we assign four time slots 15 seconds apart within the 60 second decision-making du-
ration. Only one randomly selected autonomous AIBO wolf is to make a suggestion
at each time slot, and a response to a suggestion made by any member of the team is
generated by another randomly selected autonomous AIBO wolf 2.5 seconds after the
suggestion was made. This helps to reduce the number of messages displayed and the
rate at which they must be processed by the human player. Using this interface
(Fig. 4), the human player is able to make a suggestion using the following syntax:
*{AIBO's name or nickname} move {the direction of the target square, either northwest,
northeast, or nw, ne}*. Currently messages not following this syntax cannot be inter-
preted by the autonomous AIBOs and result in a response indicating incomprehen-
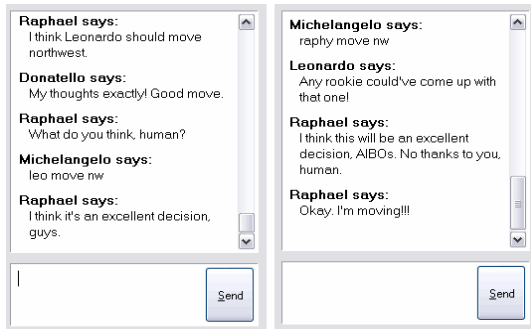sion, but we found the simple syntax to be sufficient for our present test conditions.

Fig. 4. Conversation samples between the team of wolves (Michelangelo is the human)

5   Two Robot Behaviours

Along with the game-playing interface, we have also designed two simple contrasting
robot behaviours for the autonomous AIBO wolves to test their effect on human-robot
interaction within the game, namely the human-centric condition and the robot-centric
condition. The robot behaviour which humans are most accustomed to is obedience.
The game's human-centric behaviour is designed with that human perception in mind.
When playing the game with human-centric control, the autonomous AIBO wolves
always follow suggestions given by human players. To further invoke a feeling of
superiority, we direct the autonomous AIBO wolves to praise human players for their
input, and all comments provided are communicated in a supportive manner. The op-
posite of obedience is defiance, and this is reflected in our robot-centric behaviour,
which attempts to agitate human players by placing them in a position of inferiority.
With this behaviour, the game becomes completely controlled by the three autono-
mous AIBO wolves, thinking alike and neglecting any advice from their human team-
mate. To make the situation worse, we direct the autonomous AIBOs to mock the
human player for any mistakes and suggestions that do not match their own. Even
when human players suggest a move that corresponds with the opinion of the rest of
the team, they are greeted with contempt.

6 User Evaluation

We performed a user evaluation of the two robot behaviours using our telepresence and mixed reality interface to demonstrate the utility of *Sheep and Wolves*. Our goal is to explore interaction issues between humans and robots and see if our test bed is sensitive in discovering possible interaction methods and obstacles. For this evaluation, we want to measure the human response to robot team-mates with different behaviours when immersed in a collaborative task.

The participants played games at a remote computer where the physical checkerboard was not visible. The evaluation was conducted following a written protocol to make sure each participant received the same information and followed the same experimental procedure. We introduced participants to the purpose of our study, showed them the rules and concepts of the game, and familiarized them with the remote user interface. They were told that the game supported a democratic decision-making process for the team of wolves with the decision receiving the majority of votes being selected by the team. Participants were encouraged to actively collaborate with their robot counterparts, either trying to convince the robots to support a decision or trusting the robot's decision when they are unsure about the next move. The robot-centric/human-centric deception underlying the robotic behaviour patterns was not revealed to the participants until the end of the experiment.

To explain occasional misalignment of the sheep due to camera calibration errors, we told participants that the sheep can be tricky at times and may jump from square to square on the checkerboard. Participants were told they may have to trust the advice of their robot team-mates if they are not sure where the virtual sheep is and cannot derive an intelligent move.

Each participant played one game in the human-centric condition and another in the robot-centric condition. Afterward, the participants were given a short questionnaire. Questions such as, "How well do you think you collaborated with your robot team-mates?", "How much trust did you have for your robot team-mates' suggestions?", and "How responsible are the following team-mates for the outcome of the game?" are asked. Participants were instructed to answer these questions by drawing a mark on a line segment to indicate their position between two extremes. Later, the distances denoted by the marks were measured and normalized between 0 and 1.

We started the evaluation with a pilot study involving five computer science graduate students that were not involved with our research. Several interface issues were discovered and corrected during the pilot study, and our questionnaire was refined accordingly. After the revisions we recruited fourteen participants for the main study. The demographic includes University students, professors, staff, as well as members of the general public. They each played two games, one in each condition. We counterbalanced the ordering effect by asking half to play the human-centric condition first and asking the other half to play the robot-centric condition first.

In the pilot study four out of five participants indicated that they collaborated better with their robotic team-mates in the human-centric condition, had a stronger sense of control in the human-centric condition, but had a greater sense of trust for their robotic team-mate's suggestions in the robot-centric condition. For our main study we performed statistical analysis on the data gathered using ANOVA. We found that participants felt they collaborated better in the human-centric condition, and that they

had a stronger sense of control in the human-centric condition as well, which is consistent with the pilot study results. However, unlike the pilot study, the effects on trust were inconclusive in the main study. Participants were also asked to evaluate each robotic team-mate's performance during the games, and most of them gave different scores for each team-mate although the autonomous robots ran the same algorithm and always agreed with each other.

7 Discussion and Conclusion

The *Sheep and Wolves* study presented here is an early and limited experiment performed mainly for exploratory purposes and the evaluation of the test bed's capabilities. It is hard to derive solid conclusions from the current measures that, other than the game's final outcome, are mostly qualitative and subjective in nature.

Overall, the *Sheep and Wolves* test bed, hardware and software, performed quite well. Although we had the odd traversal error in games this was fixed quickly and did not affect the game experience. Participants managed to interact with the application and play the two games in full, usually enjoying the experience.

A confirmation from both the pilot study and the main study is that players felt they collaborated better and had more control in the human-centric condition relative to the robot-centric condition. We expected this outcome which indicates that the test bed is not generating arbitrary results. In the pilot study we were surprised to find out that four out of five participants indicated that they trusted suggestions made by the robots in the robot-centric condition more than the human-centric condition. This finding is unexpected since suggestions in the robot-centric condition were forceful, less polite, and even aggressive in tone, and we were expecting them to be generally annoying. The results suggest players translated assertiveness to credibility, and trusted their robotic team-mates more when their suggestions had an added quality of effrontery (or *robotic chutzpah)*. However, the follow up results from the main study were inconclusive. This may be due to the fact that participants who played the robot-centric condition in the pilot study won the game most of the time, and in the main study most of them lost the game. We were also surprised to find that most participants believed they were interacting with three autonomous robotic entities with different characteristics and abilities even though the robots ran the same algorithm and always agreed with each other. Most participants gave different scores when asked to evaluate their team-mates' performances, and some even indicated to one of the robots as being seemingly less dependable. This demonstrates the tendency for humans to anthropomorphize life-like computing platforms like the AIBO robotic dogs used in this test bed. Although our autonomous AIBO wolves do not qualify as truly intelligent agents with their own personalities, we are able to produce a convincing collaborative experience by simply using physical life-like robots programmed with limited communication capabilities. Human participants seem to naturally project individual characteristics onto them and perceive them as autonomous team-mates with believable personalities and behaviour patterns.

Is Sheep and Wolves a useful tool for assessment of human-robot interaction paradigms? We believe it is a promising tool. The hardware and software we used and developed are reliable, replicable, and relatively affordable, allowing studies of

elaborate human-robot interaction paradigms in laboratory conditions. We think that the use of a mixed reality interface between the robots and humans highlights the unique nature of human-robot interaction tasks and the role and advantages robots will have in future applications, merging the physical and virtual domains and performing actions and accessing information in both realms.

How can we improve Sheep and Wolves? Following the test bed concept, there are many directions in which future research can progress. One is developing novel interaction techniques that will facilitate collaborative game play. We are currently designing a set of more advance mixed reality techniques for interacting with robots, allowing humans to play physically on the game board with a group of autonomous robots (much like the Harry Potter Wizard's Chess). Another interesting area for further experimentation is robot behaviour. Although this can be tested on a PC with agents, we have demonstrated that it is easier to produce a perception of realistic interaction with real robots in a real setting.

What does Sheep and Wolves signify to the domain of human-robot interaction*?* We are by far not alone in advocating the need to search for effective new interaction paradigms between humans and robots. We believe *Sheep and Wolves* and similar systems will allow high-level human-robot interaction ideas and philosophies to be easily designed, tested, and improved in research laboratory settings.

We presented the idea of constructing an effective test bed for human-robot interaction. With our prototype test bed, *Sheep and Wolves*, we are able to explore and evaluate a telepresence interface using mixed reality and two contrasting robotic behaviours. From our user evaluation, we demonstrated the utility of the test bed and discovered interesting results that may be solidified through further experimentation.

References

1. Licklider, J. C. R.: Man-Computer Symbiosis. IRE Transactions on Human Factors in Electronics, Vol. HFE-1 (1960) 4-11
2. Moravec, H. P.: *Robot Mere Machine to Transcendent Mind*. Oxford University Press Inc. (1998)
3. Dick, P. K.: *Do Androids Dream of Electric Sheep?* (1968)
4. Norman, D. A.: *Emotional Design: Why We Love (Or Hate) Everyday Things*. Basic Books (2003)
5. Breazeal, C. L.: *Designing Sociable Robots (Intelligent Robotics and Autonomous Agents)*. The MIT Press (2002)
6. Kiesler, S., Hinds, P.: Introduction to the special issue on Human-Robot Interaction. Human-Computer Interaction. Vol. 19 (2004) 1-8
7. Brooks, A., Gray, J., Hoffman, G., Lockerd, A., Lee, H., Breazeal, C.: Robot's Play: Interactive Games with Sociable Machines. ACM Computers in Entertainment, Vol. 2 (2004)
8. Robonaut, online: http://robonaut.jsc.nasa.gov/
9. Schaeffer, J.: *One Jump Ahead Challenging Human Supremacy in Checkers*. Springer-Verlag, New York (1997)
10. Tira-Thompson, E., Halelamien, N., Wales, J., Touretzky, D., Tekkotsu: Cognitive Robotics on the Sony AIBO. Abstract for poster at International Conference on Cognitive Modeling, Pittsburgh, PA (2004)
11. Web Interface for Telescience (WITS), online: http://www.ia-tech.com/wits/

12. The Telegarden, online: http://www.usc.edu/dept/garden/
13. Billinghurst, M., Kato, H., Poupyrev, I.: The MagicBook: Moving Seamlessly between Reality and Virtuality. IEEE Computer Graphics and Applications (2001) 2-4
14. Cheok, A., Goh, K., Liu, W., Farbiz, F., Fong, S., Teo, S., Li, Y., Yang, X.: Human Pacman: a mobile, wide-area entertainment system based on physical, social, and ubiquitous computing. Personal and Ubiquitous Computing, Vol. 8 (2004) 71-81