Science   Science Research & Publications

2005-09-21

Sheep and Wolves   Testbed for Interaction and Collaboration between Humans and Robots

Xin, Min; Sharlin, Ehud

http://hdl.handle.net/1880/45623
unknown

Sheep and Wolves—Testbed for Interaction and Collaboration between Humans and Robots

Min Xin   Ehud Sharlin

Interactions Laboratory, Department of Computer Science, University of Calgary

2500 University Drive NW, Calgary, Alberta, Canada, T2N 1N4; Phone: +1.403.210.9404

{xinm;ehud}@cpsc.ucalgary.ca

ABSTRACT

This paper presents the first prototype of *Sheep and Wolves*, a system for testing interaction and collaboration paradigms between humans and robots. The paper contributions are twofold: a mixed reality interface for human-robot interaction, and a practical experimental tool for assessing how different robotic behavioral patterns affect interaction and collaboration with users. *Sheep and Wolves* places humans, robots and virtual entities in a game environment where they have to collaborate and compete. The system is designed around the classic *Sheep and Wolves* board game, played on a large physical checkerboard. In the prototype presented here the user is playing a single wolf in a pack of four autonomous robotic wolves trying to hunt a single virtual sheep. The human interacts with the rest of the wolf pack using a mixed reality video stream, a graphical interface and a text chat tool that enables discussion and planning of future moves within the pack. In preliminary testing *Sheep and wolves* was sensitive to differences in the robots' behavioral patterns and suggested that robotic assertiveness (or *robotic chutzpah*) might enhance the quality and trustfulness of the interaction.

Categories and Subject Descriptors

General Terms

Keywords

1. INTRODUCTION

How will humans, intelligent computers and robots coexist and collaborate? This question motivated thinkers and writers for a long time, with visions 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 scholars and designers of human-robot interaction (HRI) paradigms no longer see robots as fully-controlled subordinates but rather as peers and colleagues with a spectrum of social and emotional abilities (see for example [4,5]). It is logical that humans will find future autonomous robots more useful if the robots act according to behavioral and emotional patterns that humans can recognize and relate to.

For example, Norman suggests that a housemaid robot will relate to cleaning the top of the stairs area by expressing fear of height, helping its owner to intuitively understand the problems of cleaning this location [4]. Similar parallels can be thought of when considering robots acting in search and rescue operations, a battlefield or in space missions. However, a basic question arises from this line of thought: which emotions would humans expect a robot to express in different scenarios and tasks? Would users always want the robot to be an obedient subordinate even in

situations where the robot is more knowledgeable about the task at hand? How would different robotic emotional and behavioral patterns affect the resulting HRI quality and effectiveness?

We designed *Sheep and Wolves* following what we see as a practical near future scenario. Think for example of a search and rescue operation where robots and humans enter a disaster zone looking for survivors. While the humans are expected to have a high level of cognitive understanding of their task as well as moral and ethical values that the robots lack, the robots can instantly access digital information which the users cannot approach directly. For example the robots can view the scene in various non-visible spectral wavelengths. They can extract information using elaborate computationally-demanding algorithms, and access vast online information. Following this vision, the *Sheep and Wolves* system (Figure 1) was designed as a relatively simple indoor experimental testbed that allows rich interaction between humans and robots and enables investigation of a range of social, behavioral and emotional patterns and scenarios.

Our *Sheep and Wolves* system, based on the classic board game, is played on a large physical checkerboard where a pack of wolves needs to hunt a sheep. The wolves can only move forward while the sheep can also move backward, hence the wolves must play as a group if they want to circle and hunt the sheep. Humans, robots as well as virtual entities can play different roles in the game enabling a large variety of scenarios. Virtual entities were included in the game, using mixed reality technology, in order to highlight one of the robots main advantages over the humans: their ability to function in both the physical and virtual realms. The 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.

In the first prototype of the *Sheep and Wolves* system a group of four Sony Aibo[TM] robot dogs are playing the game against a single virtual sheep on a large physical board (Figure 1). Three of the robots are completely autonomous and one is being controlled by a remote user using a multimodal mixed reality interface which allows interaction with both the physical and virtual components of the task. Before each move the wolves, human and robots, must use a textual chat mechanism to discuss their next move against the escaping sheep in a timely fashion.

In the current experimental setup the users were told that a decision will be reached according to a democratic and fair voting process where both they and the other three wolves can cast a vote towards the next move. In practice however the robots function according to two extreme conditions: in the *human-centric control* mode the robotic wolves will always be supportive and approving, confirming the human suggested move even if it contradicts their own thoughts and logic; in the *robot-centric*

Figure 1. *Sheep and Wolves*

control mode the robotic wolves will always be negative and unenthusiastic regarding the human suggested move (for example, even if the move is inline with their thoughts they will write to the user messages like "*But I thought of this move earlier*"…).

In the following we review HRI efforts related to our work. We then describe the *Sheep and Wolves* interface and our current experimental setup in detail. We conclude with results from a preliminary user study we performed suggesting that *Sheep and Wolves* is sensitive to the robotic behavioral pattern conditions.

2. RELATED WORK

2.1 Human Robot Interaction

Although the fields of Human-Computer Interaction (HCI) and Robotics have generated significant interests and made substantial progress, 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 HCI. However, as increasingly intelligent and capable autonomous robots come into existence, traditional HCI approaches become less applicable and special attention is given to the unique requirements and advantages embodied in autonomous robots [6]. People intuitively perceive robots to be more human-like (anthropomorphic) than they do when relating to other computing systems [6]. This hints at the potential for more intimate and natural forms of interaction amongst humans and robots. Breazeal's Kismet [5], an expressive anthropomorphic robot able to convey emotions and provide feedback, allows 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 behavior 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 HCI paradigms of clicking on icons or opening windows. Furthermore, with mobile autonomous robots, interaction occurs within 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 robotic astronaut equivalents that are able to collaborate with humans in order to perform tasks in space.

Works such Breazeal's Kismet, Robotic Life Group's Leonardo, and NASA's Robonaut are the prelude to a fascinating future for the field of HRI. Attitudes towards HRI are already shifting from a "robots as tools" approach to a "robots as partners" outlook.

2.2 Telerobotics and Mixed Reality

Robert Heinlein's fictional character Waldo [9] invented a series of remote manipulators, WALDOs, that enabled him to cope with his severe muscular weakness. NASA developed the Web Interface for Telescience (WITS) [10] 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 [11] 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 telerobotics in encouraging remote collaboration, active and assisted learning, and developing a sense of virtual partnership.

With autonomous robots, the benefits of telerobotics can be extended further. Arguments can be made that most current telerobotics 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 telerobotics, remote users can collaborate with a team of remote robots as a participating member rather than a superior operator having to control the entire team. This methodology allows HRI to make the first steps towards involving the general public in testing novel interaction techniques.

In order to provide a rich interface for telerobotics, basic sensory elements of the remote environment is essential. Visual information is often provided as live video depicting the remote scene. Augmented/mixed reality is a practical method for supplementing the visual environment with graphics which serve a specific purpose or express indirect information. Applications such as the Magic Book [12] or the Human Packman [13] use mixed reality to superimposes graphics onto physical scenes displaying and allowing rich interaction with virtual entities. Mixed reality can also be effectively used in telerobotics applications to compensate for the loss of direct physical context.

For example, graphics can be displayed to indicate the parameters, thoughts or emotions of other remote robotic teammates.

2.3 Game-playing Robots

Robots or simulated computer agents playing games against humans are a familiar concept. However, interaction amongst humans and such competing machines are often trivial or non-existent. Applications such as IBM's chess-playing Deep Blue[TM] or various other machines playing checkers or air hockey do not require natural and intuitive interaction amongst humans and the robotic opponents. Carnegie Mellon University's Cognitive Robotics [14] suggests means of implementing more evolved physical interaction between robots and games, for example the paper outlines an Aibo-based tic-tac-toe game where the Aibo can move game pieces on a physical board.

Interaction amongst robots and humans within a game application can be further improved by requiring humans and robots to play on the same team instead of against each other. The concept originates from using robots for search and rescue operations where performing collaborative task can be critical. Since human ability, artificial intelligence and computational ability can be fairly balanced within a limited game environment, it is conceivable to implement meaningful human-robot interfaces where the robots and humans collaborate as equals.

3. SYSTEM DESIGN

Sheep and wolves is intended to explore interaction issues amongst humans and robots. The focus is on discovering novel methods of interaction, effective robot behavior, and corresponding human reactions. To accomplish our goals, we have devised a collaborative game where humans and robots must play together as a team in order to win the game. We see this game as a metaphor for future human-robot collaborative tasks such as search and rescue operations. By performing a collaborative task in a controlled game environment instead of the complex physical world, we are able to simplify the task and focus on interaction. Also, as mentioned previously, interaction becomes more meaningful and believable when robot and human intelligence are fairly balanced which is the case in our game.

3.1 Sheep and Wolves

The application we have developed is based on a classic board game called *Sheep and Wolves*. This turn-based game is being played on a checkerboard, and game pieces can only occupy and move on squares of the same color. The game involves fives 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 each turn. The team's objective is to surround the sheep so it cannot make any legal moves. Meanwhile, the sheep is allowed to move forward and backward diagonally by one square during each 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' strengths are in their numbers and ability to move as a pack.

3.2 Conceptual Description

In order to provide our application with a physical environment, the game is being played on a large physical checkerboard. The wolves are represented by Sony's Aibo robotic dogs, and the sheep is a virtual entity. The Aibos physically move and sit down on the physical checkerboard to indicate movement of the wolves in the game. Human players play the game at remotely located computers by logging into the application. Using telerobotics, a human player is able to control an Aibo wolf, personifying the robotic entity within the game. Other uncontrolled Aibo wolves are autonomous robotic teammates which the human player must collaborate with. Live video of the physical game environment is provided to the remote human player, and augmented/mixed reality is utilized for visualizing the virtual sheep.

Winning the game as wolves requires excellent teamwork. The human player has to provide suggestions to the team and consider propositions made by other teammates in order to help the team reach intelligent decisions on the moves the team should make.

3.3 Two Game Conditions

We have designed two extreme robotic behaviors for the autonomous Aibo wolves to test their effect on human-robot collaboration within the game. The two conditions will portray significantly different personalities for the human player's robotic teammates.

3.3.1 Human-Centric Condition

The robot behavior which humans are most accustomed to is obedience. After all, the "robots as tools" approach has been adopted in many robotic applications, and people often expect robots to perform the tasks they are asked to do. The game's *human-centric condition* is designed with that human perception in mind. When playing the game with *human-centric control*, the human player's robotic teammates always follow advice given by the human player. To further invoke a feeling of superiority, we direct the autonomous Aibo wolves to praise the human player for his/her input, and all comments provided are communicated in a supportive manner.

3.3.2 Robot-Centric Condition

The opposite of obedience is defiance, and this is reflected in our *robot-centric condition*. We attempt to agitate the human player by placing him/her in a position of inferiority. In essence, the game will be completely controlled by the three autonomous Aibo wolves, all following the same game algorithm, thinking alike and neglecting any advice from their human teammate. To make the situation worse, we direct the autonomous Aibos to mock the human player for any mistakes and move suggestions that do not match their own. Even when the human player suggests a move that corresponds with the opinion of the rest of the team, he/she is greeted with contempt.

4. SYSTEM IMPLEMENTATION
4.1 System Hardware

To provide a physical environment for the game, we elected to use a 104'' by 104'' RolaBoard[TM] with the standard black and white checkerboard pattern. Each square measures 13'' by 13'', providing sufficient room for an Aibo wolf to sit on or humans to stand on. We chose four Aibos as our robotic wolves because they have a playful appearance, the ability to physically walk on the checkerboard, vision capabilities, and wireless connectivity which is essential for our telerobotics application. The Aibo provides 208 pixels by 160 pixels color streaming video and

connects wirelessly using the 802.11b standard. The remote computer used for experiments is a Pentium IV 3.4 GHz machine with 2 GB of RAM.

4.2 System Software

The software required for running the game is categorized into local software and remote software. Local software describes processes running natively on the Aibos such as checkerboard traversal and the game algorithm, and remote software involves the remote user interface application. Communication occurs between local and remote software using wireless networking. Carnegie Mellon University's Aibo programming platform, Tekkotsu [15], was used extensively for the local software, offering access to image data, wireless network administration, and easy walking commands.

4.2.1 Low Level Checkerboard Traversal

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 algorithm, we decided to use lines and corners as means of localization and determining orientation. Working only with low resolution grayscale 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 Figure 2. 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 neighbors.

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.
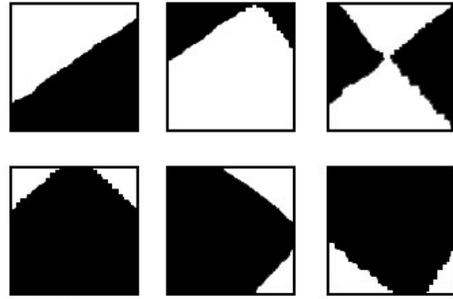

Figure 2. Binary image representations of snap shots of the Aibo's view of the board

With corner points and lines as our vision cues, we then define the three stages of a move: pre-move, move, and post-move. For each stage, a set of states are outlined. In the pre-move stage, the Aibo is programmed to walk forward, find a line, orientate itself so the line appears at a certain angle to the vertical axis depending on the direction of the target square, and follow the line until a corner point is seen. This stage positions the Aibo to face the target square. In the move stage, the Aibo crosses the corner point and tries to find the next corner point by moving forward. If a line segment is seen before the next corner is found, the Aibo uses the line segment to adjust its orientation, trying to maintain a 45 degree angle between the line segment and the vertical axis. The Aibo then follows the line and eventually finds the next corner point. This stage moves the Aibo from its current square to the target square. In the post-move stage, the Aibo uses the position of the found corner point to either turn left or right of the corner point depending on the direction of the move. After the corner point is lost due to turning, the Aibo sees a single line segment case. For the final adjustment, the Aibo attempts to maintain a 90 degree angle between this line segment and the vertical axis while backing up. This stage ensures the Aibo is facing directly forward on the target square after the move. Lastly, the Aibo sits on the square to indicate the completion of a move.

4.2.2 Augmenting the Physical Scene

In order to visualize the virtual sheep and demonstrate the application of augmented/mixed reality, we enhance the live video provided by the Aibo's camera by superimposing a computer generated 3D sheep onto the scene (Figure 1). To achieve this, we set up an OpenGL^{TM} 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 board, we know approximately the position of the camera we are calibrating. Using this information and measurements of the camera's height and tilt, we set the camera in the world coordinate system based on these parameters. Next, we fine-tune the calibration using high resolution image data from the camera.

Figure 3. Extracted corner points

First, we extract the checkerboard corner points from the image (Figure 3). This is accomplished using a corner detection algorithm. The effectiveness and accuracy of the algorithm depends on the closeness of the corner points to the camera and the amount of perspective distortion present. In most cases, we can extract at least three accurate points close to the camera which we use to perform a simplified camera calibration. There are several challenging cases where calibration is difficult or impossible using only the information described. These will be explained later in detail.

After obtaining the corner points, we then inverse project these 2D points into our 3D world coordinate system. This is possible because we know the y values of these potential 3D points are all supposed to be 0. With the inverse projected 3D corner points, we pair them 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 amount to make the adjustment. This corrects misalignment issues caused by the Aibo not always facing exactly forward. To make adjustments for possible changes in the tilt of the Aibo's camera, we rotate the scene around the x-axis by a ratio determined from the difference between the length of an actual edge and the length of an edge extracted from the image. We make the assumptions that the Aibo's camera does not require roll adjustment, and its height remains the same.

Applying the calibration procedures illustrated, 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 sometimes in the inability to accurately portray the sheep. Another visual annoyance that sometimes occurs is the lack of clipping for the 3D sheep when it is displayed behind a physical Aibo in the video.

4.2.3 Game Play

The game algorithm for both the sheep and the 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 robotic members of the wolf pack will make the move which results in the longest available path or no available path for the sheep.

For each turn, the sheep or the wolf pack has 60 seconds to arrive at a decision for the move. The wolves win when the sheep can no longer make another legal move, and the sheep wins if it gets pass the last wolf on its way to the other end of the board.

At the end of the sheep's turn, each autonomous Aibo wolf processes the same game algorithm and comes up with its ideal move for the team to make. Then it makes the suggestion to the rest of the team, and the other autonomous Aibos provide either positive or negative feedback depending on if the suggestion matches their own ideal move. Since all autonomous Aibos process the same game algorithm they will always agree with each other. This behavior may not be realistic for actual collaborative work but is appropriate for our test conditions. Responses generated when processing human suggestions differ depending on the game condition being played. In the *human-centric condition*, autonomous Aibo wolves always provide positive feedback to human suggestions. While in the *robot-centric condition*, human suggestions are compared with the move generated by the game algorithm. Mismatches trigger mocking comments, and matches trigger a reluctant agreement with contempt.

To play a game, we first turn on the Aibo wolves and allow them to stretch. After the Aibos reach a sitting position, they are placed on the white squares of the last row of the checkerboard facing the other end. Next, we show the Aibos where they are on the checkerboard using a simple two-button interface which iterates through the rows and columns until the correct square is located. After this initialization, the Aibos are able to keep track of their positions on the checkerboard as they move in the game. At this point, the remote user interface is launched. The Aibo to be remotely controlled is selected and the game begins.

4.2.4 Remote User Interface

To allow a human player to effectively control an Aibo wolf and naturally interact with the rest of the team, we have devised an intuitive graphical user interface (Figure 4). In the following section, 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 user to see the physical board from the point of view of the controlled Aibo. The virtual sheep is visible to the user 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 radar (Figure 5) indicates the positions of the wolves relative to the edge 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
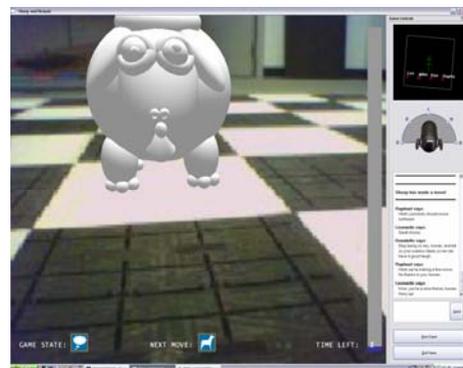

Figure 4. Remote user interface

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 robotic wolf is represented by a red dot. The robotic wolf controlled by the human player is indicated with a blue ring around its dot. When a robotic wolf moves, its dot will flash to indicate the movement. Displayed next to their corresponding dots are the Aibo robotic wolves' nicknames. These along with their names, Leonardo, Michelangelo, Donatello, and Raphael, are used to refer to the 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 towards north.

Underneath the radar is the head-panning device (Figure 5). Since the initial forward-facing view is limited, we allow the human player to pan the head of the controlled Aibo 45º or 90º left or right (east or west). This feature can be used to explore the checkerboard, locate the sheep, observe other Aibo teammates, and watch them move. The radar also rotates to match the orientation of the Aibo's current view to further assist in spatial orientation and awareness.

The most important interface component is the text messaging interface (Figure 6). This allows the human player to communicate with the rest of the Aibo wolves in a familiar interaction paradigm. Although currently the richness of conversation is lacking, we feel this interaction technique has potential in effectively engaging human users in active collaboration with robotic entities especially in telerobotics applications since most human users are already familiar with instant messaging programs.

In our game, conversation occurs amongst four teammates. Due to the lost 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 duration. 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 (Figure 6), the human player is able to make a suggestion using the syntax *"{Aibo's name or nickname} move {the direction of the target square, either northwest, northeast, or nw, ne}"*. Currently messages not following the syntax cannot be interpreted by the autonomous Aibos, but we plan to expand the Aibo's vocabulary by implementing a more advanced language parsing feature. The simple syntax is sufficient for our present test conditions.

5. Preliminary Testing of *Sheep and Wolves*
In order to evaluate the sensitivity to behavioral conditions, usefulness, and playability of *Sheep and Wolves*, we have performed preliminary evaluations of our application. In this section, we will outline our experimental approach and present the current results.
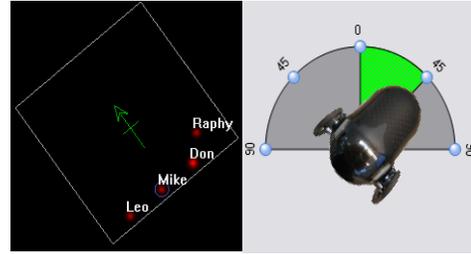
Figure 5. Radar and head-panning

5.1 General Methodology
Our goal for the project is to explore collaboration issues between humans and robots. We want to measure the human response to robotic teammates with different behavior conditions when immersed in a collaborative task. As mentioned previously, we have constructed two conditions: the *human-centric condition* and *the robot-centric condition*.

The participants played the game at a remote computer where the physical board was not visible. The experiment was conducted following a written protocol to make sure each participant received the same information. 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 robotic counterparts, either trying to convince the Aibos to support a decision or trusting the Aibo's decision when they are unsure about the next move.

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 robotic teammates if they are not sure where the virtual sheep is and cannot derive an intelligent move. We also asked participants to keep in mind questions such as "*Can you accept an advice from robots that might know more about the situation than you do?*" and "*Will you be able to convince your robotic teammates in the logic of your plans and actions?*".

Each participant played one game with human-centric control and another with robot-centric control. The condition order was counterbalanced between participants. We recorded the outcome
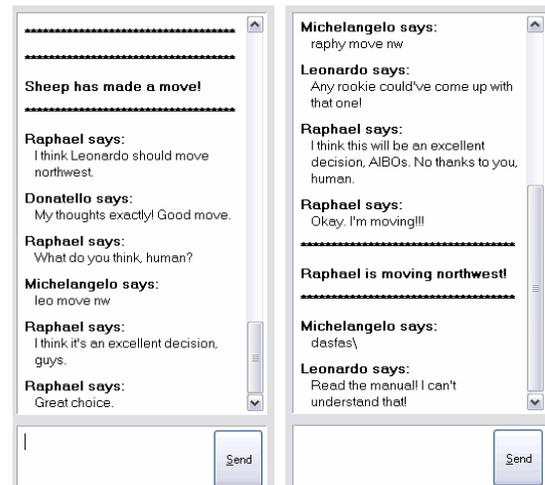
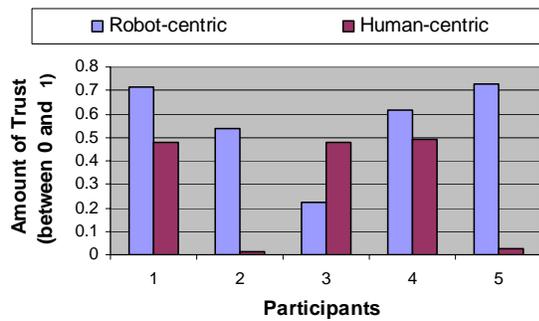**Figure 6. Chat interface showing conversions from the two game conditions**

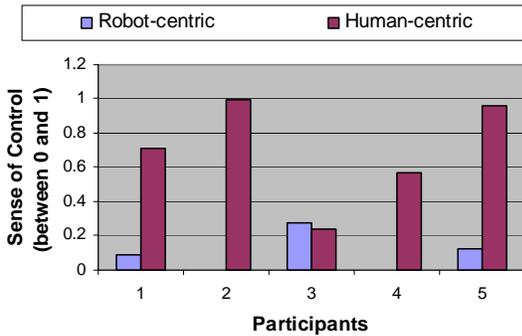**Figure 7. Human trust for robots' suggestions**

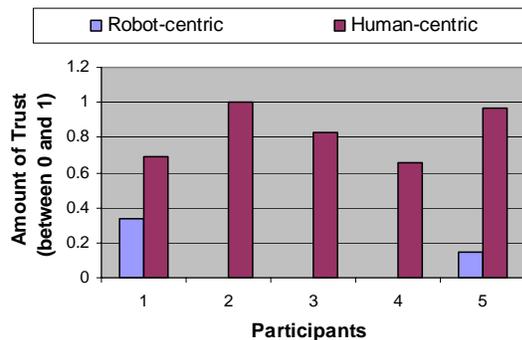

Figure 8. Human sense of control

Figure 9. Perceive trust of robots' trust for humans

of each game and any program errors. After each game the participants were given a short questionnaire. The following is a list of some of the questions asked:

- *How much trust did you have for your robotic teammates' suggestions?*
- *How strong was the sense of control you had over your robotic teammates?*
- *How much trust do you think your robotic teammates had for your suggestions?*

Participants were asked 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 a value was calculated and normalized between 0 and 1.

5.2 Experiments

We started our studies with a limited pilot experiment performed by one male student. Several interface issues were discovered and corrected such as disabling the radar display of the position of the sheep which can lead to the participant playing the game using only the radar and not the live video of the physical environment.

We also learned the participant wanted to pan the head of the Aibo and allow it to stay at the new location, where in our initial design the head always panned back to the forward-facing position after a certain amount of time.

After refining the system according to the pilot experiment we invited five graduate students from our lab to play the game. Three participants played one game with *human-centric control* first and another with *robot-centric control* second. The other two participants played one game with *robot-centric control* first and another with *human-centric control* second. Each experiment, consisting of two games, lasted approximately one hour, and all participants completed the evaluation. Roughly once per game there was a traversal error, and the robot position had to be corrected, however, this did not affect game play. Four participants won the game played with robot-centric control and lost the game played with human-centric control. One participant lost both games.

Figure 8 shows results from the questionnaire indicating the amount of trust the human player had for the robots in each behavioral condition. All but one participant had more trust for the robots in the *robot-centric condition* than in the *human-centric condition*. Figure 9 indicates the sense of control the human player had during the game. It is clear that all participants except #3 had a stronger sense of control in the *human-centric condition*. Figure 10 points to the participant's perception of the amount of trust the robots had for their human suggestions. We can see that all participants thought the robots trusted them more in the *human-centric condition*.

6. DISCUSSION

Here we discuss the outcomes of our preliminary experimental examination of *Sheep and Wolves*. We review the limitations of the current implementation, results we were expecting as well as surprises. We discuss implications for the use and improvements of the *Sheep and Wolves* system.

6.1 Limitations

The *Sheep and Wolves* study presented here was an early and limited experiment and its results should be considered with caution. It is hard to derive solid conclusions from the limited number of users and from the current measures that, other than the game final outcome, are qualitative and subjective in nature.

We use the game final outcome (wolves won or sheep won) as an objective outcome of the experiment. However, we are currently not taking into account the full complex nature of the experiment, something that could be rectified by measuring the finer factors that influence the game progress, such as overall number of moves, and objective correctness or incorrectness of the moves taken.

6.2 Confirmations

Overall, the *Sheep and Wolves* system, hardware and software, performed quite well. Although we had the odd traversal error in each game this were fixed quickly and did not affect the game experience. Users manage to interact with the application and play the two games in full, usually enjoying the experience. Expectedly, all the users but one won all their *robot-centric control* condition games (the one loss was due to an error in the robotic wolves algorithm).

An expected confirmation was that users felt more control over their robotic teammates in the *human-centric control* condition relatively to the *robot-centric control* condition experiments. We were also glad to find that all users felt that their robotic teammates trusted them (the humans) more in the *human-centric control* condition experiments.

6.3 Surprises

We were surprised to find that all users lost their *human-centric control* condition games. We suspect this was caused by lack of experience in the classic board game "Sheep and Wolves" which was played by some of the users for the first time during the experiment, as well as by lack of experience in the *Sheep and Wolves* interface. The classic game of "Sheep and Wolves" can be won by the wolves every time if they maintain an uninspiring "keep your column" game plan. The robotic wolves kept up to this plan and continued to suggest the right moves to the human wolf. However, we noticed that users followed the uninspiring suggestions from the robots till they reached close contact with the sheep. At which point it seems to us that users tended to follow their intuition and deviate from the robots' good (but boring) advised moves. This being *human-centric control* condition, the robot were quick to confirm and support the human suggested move, and the game was lost.

We were pleasantly surprised that most users reported trusting the suggestions coming from their robotic teammates more in the *robot-centric control* condition than they did in the *human-centric control* condition. This finding is surprising since the *robot-centric control* condition move suggestions were forceful, less polite and even aggressive in tone, and we were expecting them to be generally annoying. The current results suggest users translated assertiveness to credibility, and trusted their robotic teammates more when their move suggestions had an added quality of *robotic chutzpah*.

6.4 Implications

Is Sheep and Wolves a useful tool for assessment of HRI paradigms? We believe it is a promising tool. The hardware and software we used and developed were reliable, and are replicable and relatively affordable, allowing studies of elaborate HRI paradigms in lab conditions. We think the use of a mixed reality interface between the robots and human highlights the unique nature of HRI 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? In the short term we plan to polish the current interface and solve the remaining technical issues (Section 4). We are planning to further enhance the chat mechanism allowing richer textual interaction between the user and the robots. Later, we are planning to develop versions of *Sheep and Wolves* which allow several humans to play with several robots and possibly with multiple virtual entities. Another direction we are pursuing is allowing the humans to play physically on the game board, interacting directly with the robots using gestures and speech.

What does Sheep and Wolves signify to the domain of HRI? 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 lab setting.

7. CONCLUSION

In this paper we presented the design, implementation and early testing of *Sheep and Wolves*, a mixed reality-based interactive testbed for assessment of HRI paradigms. *Sheep and Wolves* is based on a large board game where robots, humans and virtual entities play together. Our current implementation involves a virtual sheep and a pack of four wolves composed of three autonomous robotic wolves and one played by a human user. The system was evaluated in two extreme robotic behavioral conditions: *human-centric control* where robots tend to confirm and follow moves suggested by the human and *robot-centric control* where the robots tend to disregard and disrespect the human suggested moves. In preliminary testing the system demonstrated sensitivity to the robotic behavioral patterns, and pointed to the possibility that adding assertiveness to robotic reactions might enhance their overall perceived reliability and effectiveness. In the future we are planning to expand *Sheep and Wolves* capability to support collaboration of several users with the robotic wolves as well as to allow humans to play the game while standing on the physical game board.

8. REFERENCES

[1] Licklider, J. C. R., "Man-Computer Symbiosis", IRE Transactions on Human Factors in Electronics, volume HFE-1, pages 4-11, 1960.

[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., and Hinds, P., "Introduction to the special issue on Human-Robot Interaction", Human-Computer Interaction. Volume 19, pages 1-8, 2004.

[7] Brooks, A., Gray, J., Hoffman, G., Lockerd, A., Lee, H., and Breazeal, C., "Robot's Play: Interactive Games with Sociable Machines.", ACM Computers in Entertainment, volume 2, 2004.

[8] Robonaut, online: http://robonaut.jsc.nasa.gov/

[9] Heinlein, R. A., "Waldo," in *Waldo and Magic, Inc.* NY: Garden City: Doubleday, 1950.

[10] Web Interface for Telescience (WITS), online: http://www.ia-tech.com/wits/

[11] The Telegarden, online: http://www.usc.edu/dept/garden/

[12] Billinghurst, M., Kato, H.,Poupyrev, I.: The MagicBook: Moving Seamlessly between Reality and Virtuality In IEEE Computer Graphics and Applications, pages. 2-4, 2001.

[13] Cheok, A., Goh, K., Liu, W., Farbiz, F., Fong, S., Teo, S., Li, Y., and Yang, X., "Human Pacman: a mobile, wide-area entertainment system based on physical, social, and ubiquitous computing", Personal and Ubiquitous Computing, volume 8, pages 71-81, 2004.

[14] Tira-Thompson, E., Halelamien, N., Wales, J., and Touretzky, D., Tekkotsu: Cognitive Robotics on the Sony AIBO, Abstract for poster at International Conference on Cognitive Modeling, Pittsburgh, PA, 2004.

[15] Tekkotsu Framework, online: http://www.cs.cmu.edu/~tekkotsu/