Sheep and Wolves
Test Bed for Human-Robot Interaction

Min Xin
Interactions Laboratory,
Department of Computer Science,
Univeristy of Calgary.
2500 University Drive NW
Calgary, AB T2N 1N4 Canada
xinm@cpsc.ucalgary.ca

Ehud Sharlin
Interactions Laboratory,
Department of Computer Science,
Univeristy of Calgary.
2500 University Drive NW
Calgary, AB T2N 1N4 Canada
ehud@cpsc.ucalgary.ca

Copyright is held by the author/owner(s).
CHI 2006, April 22–27, 2006, Montreal, Canada.
ACM 1-xxxxxxxxxxxxxxxxx.

Abstract
This paper presents a dynamic 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 within the game. With meaningful interaction occurring within our confined setup, various aspects of human-robot interaction can be easily explored and evaluated such as interface methods. We also present the results of a user evaluation which shows the sensitivity of our system in assessing robotic behaviours.

Keywords
Human-robot interaction, augmented reality

ACM Classification Keywords
H.5.2 [Information Interfaces and Presentation]: User Interfaces—Interaction Styles

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 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 behavioural patterns that humans can recognize 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. Our aim is to construct a manageable environment and task that will serve as a metaphor for major human-robot interaction applications and encompass their common interactive qualities. Our idea is a collaborative board game involving humans and robots played in a shared physical environment. For our prototype system following this concept, we implemented *Sheep and Wolves* based on a classic checkerboard based game in which humans and robots play together as a team of four wolves in attempt to hunt or surround a single sheep. Our game is played on a large checkerboard, allowing humans and robots to represent game pieces and engage in physical interaction while collaborating within the shared environment.

Humans, robots as well as virtual entities play as game pieces, enabling a large variety of scenarios. Virtual entities were included in the game, using augmented reality, in order to highlight one of the robots' main advantages over the 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 the prototype system, we have also designed two extreme robot behaviours that are evaluated using *Sheep and Wolves*. We performed user evaluations asking human participants to play two games using the prototype interface where in one game the robots are always supportive and obedient, and in the other game the robots behaved negatively and always ignored input from their human teammate. The responses from the human participants are collected, analyzed, and found to be sensitive to the contrasting robotic behavioural conditions.

Game-Playing Robots

Simulated computer agents playing games with or against humans are a familiar concept (see for example [6]). However, interaction and collaboration between humans and robots within a physical game environment is rare. For example, Carnegie Mellon University's Cognitive Robotics [7] 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, but the potential for human-robot interaction is still limited.

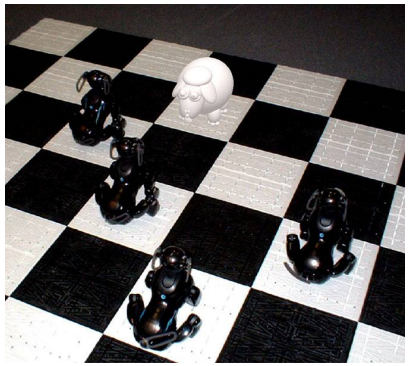

figure 1. *Sheep and Wolves* (wolves closing in on sheep)

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. The concept originates from using robots for search and rescue operations where performing collaborative tasks 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.

Sheep and Wolves

Sheep and Wolves (figure 1) is a human-robot interaction test bed following our goal of constructing a manageable environment and task that will serve as a metaphor for major human-robot interaction applications. The task we have designed is based on a classic board game. This turn-based game is played on a checkerboard, and game pieces can only occupy and move on squares of the same color. 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 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.

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 the complex physical world, we are able to focus on interaction. Also, since implementing artificial intelligence for checkerboard based games is relatively simple, we are able to easily adjust the intelligence of the robots in order to develop varying robotic behaviours.

In our game we have elected to use Sony's Aibo dogs as our robotic participants. These fairly capable commercial robots allow us to rapidly build prototype interfaces for evaluation. For the physical environment of the game, we elected to use a 264cm (104'') by 264cm RolaBoard™ with the standard black and white checkerboard pattern. Each square measures 33cm (13'') by 33cm, providing sufficient room for an Aibo

wolf to sit on or humans to stand on. This confined 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 augmented reality interfaces and localization of humans on the checkerboard.

Implementation

In our prototype interface application, all four wolves are represented by the Aibos and the sheep is a virtual entity. The Aibos physically move and sit down on the checkerboard to indicate 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 robotic teammates 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 augmented reality is utilized for visualizing the virtual sheep. Winning the game as wolves requires 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.

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 physical 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. For our 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 then performing a binary threshold and finally searching for line end points around the perimeters of the images.

Augmenting the Physical Scene

In order to visualize the virtual sheep and demonstrate the application of augmented 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 viewing frustum based on the camera's field of view and focal length. Frames of video received from the Aibo's camera are then texture mapped onto a distant rectangle to provide a video background for the virtual 3D sheep in the scene. As the Aibo moves on the checkerboard, the exact position of the Aibo's camera is 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 high resolution image data from the camera, we are able to fine-tune the calibration by extracting and using the corner points of the checkerboard.

Game Play

For each turn, the sheep or the team of wolves has sixty seconds to arrive at a decision for the next move. 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 (figure 2), and other autonomous Aibos provide either positive or negative feedback depending on if 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 uncertainties and lack of information. Although we provide the user 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 other Aibos, allowing us to explore issues of trust between humans and robots.

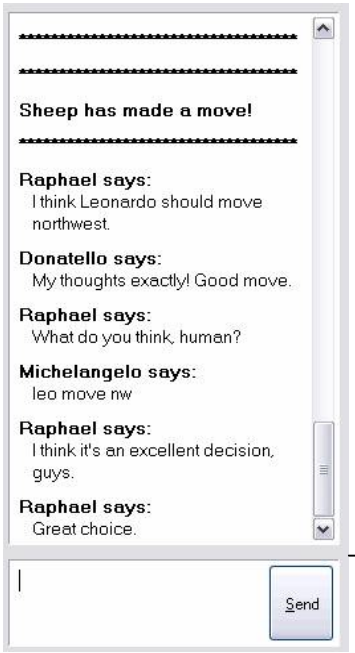

Conversation between the team of wolves (Michelangelo is the human player)

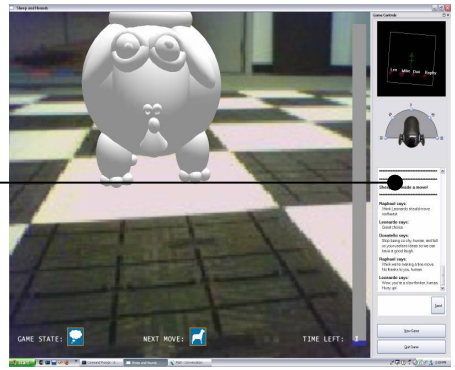

figure 2. Telepresence interface using augmented reality

Two Robotic Behavioural Conditions

We have designed two extreme contrasting robotic behavioural conditions for the autonomous Aibo wolves to test their effect on human-robot collaboration within the game.

Human-Centric Condition

The robot behavior which humans are most accustomed to is obedience. The game's *human-centric condition* 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 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.

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, 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 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.

User Evaluation

We performed a user evaluation of two robotic behaviours using our telepresence interface to demonstrate the utility of *Sheep and Wolves*. The

evaluation was conducted by asking participants to play two games: one in the human-centric condition and one in the robot-centric condition. We introduced participants to the purpose of our study, showed them the rules and concepts of the game, and familiarized them with the user interface. They were made to believe that the game supported a democratic decision-making process for the team of wolves with the decision receiving the majority of votes or suggestions being selected by the team. Participants were encouraged to actively collaborate with their robotic teammates, either trying to convince the Aibos to support a decision or trusting the Aibos' decision when they are unsure about the next move. After the games participants were asked to complete a questionnaire, and casual interviews were conducted. We first performed a pilot study with five participants from our laboratory. This evaluation resulted in one interesting finding. Four out of the five participants indicated that they trusted suggestions made by the robots in the robot-centric condition more than the human-centric condition. This led us to believe that assertive robotic behaviour may increase trust. In our actual study we recruited students and faculty from around campus, and with fourteen participants we were able to perform statistical analysis of collected results. Using ANOVA we were pleased to find that participants felt they collaborated better and had more control in the human-centric condition than the robot-centric condition. This is what we expected, and results were consistent with the pilot study. However, upon re-examination of the trust condition, the statistics 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 actual study most of them lost the game.

Conclusion and Future Work

We presented the idea of constructing an effective test bed for human-robot interaction. With our prototype test bed, *Sheep and Wolves*, we were able to explore and evaluate a telepresence interface using augmented 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. We intend to improve the interface, implement and test more interesting robotic behaviours, explore other simple collaborative board games such as an end-game condition in chess, and investigate the interaction issues that arise when humans and robots play together on the physical checkerboard.

References

[1] Licklider, J.C.R., Man-Computer Symbiosis. *IRE Transactions on Human Factors in Electronics*, HFE-1, 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] Schaeffer, J., *One Jump Ahead Challenging Human Supremacy in Checkers*. Springer-Verlag New York, Inc., 1997.

[7] 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.