Manipulating Mental States through Physical Action

Jesse Gray and Cynthia Breazeal

MIT Media Lab, 20 Ames Street, Cambridge, MA
{jg,cynthiab}@media.mit.edu

Abstract. We present our implementation of a self-as-simulator architecture for mental state manipulation through physical action. The robot attempts to model how a human's mental states are updated through their visual perception of the world around them. This modeling, combined with geometrically detailed, perspective correct simulations of the immediate future, allows the robot to choose actions which influence the human's mental states through their visual perception. The system is demonstrated in a competitive game scenario, where the robot attempts to manipulate the mental states of an individual in order to win. We evaluate people's reaction to the system, focusing on the participants' perception of a robot with mental state manipulation capabilities.

1 Introduction

This paper focuses on a demonstration of mental state manipulation in a competitive game scenario and an evaluation of human reactions to this behavior. The motivation for this work is to explore the connection between (hidden) mental states of an embodied agent and the (observable and modifiable) world in which they exist.

An embodied agent exists in the physical world, and though its mental states are hidden, there are rich connections between the agent's mental states and the world in which it is operating. Observing an agent's visual perspective and physical actions can help inform a model of the underlying mental states caused by those perceptions and causing those actions.

In order to modify the mental states of another agent, it is necessary to manipulate the world such that the agent's observations will update its internal mental states to the desired configuration. Usage of this critical skill encompasses a very broad set of interactions - in a cooperative scenario we might call it communication. In such a scenario we often collapse this concept into "tell" or "show." However the underlying goal remains: to change a mental state through physical action on the world that will cause the agent's perceptual system to update their mental states in the desired manner. This can be demonstrated by observing how we respond to failures of the exchange - e.g., by moving an object into the line of sight of an observer who is not paying attention.

S.S. Ge et al. (Eds.): ICSR 2012, LNAI 7621, pp. 1–14, 2012.

In a competitive scenario we call this same ability deception. In this case the world is intentionally modified (again taking into account the perceptual capabilities of the target agent) in such a way that the observing agent will form incorrect mental states.

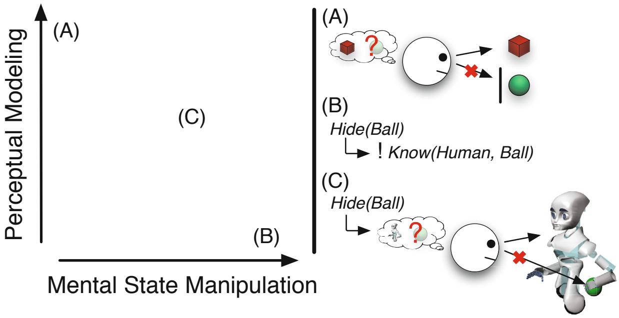

Fig. 1. Systems such as BDI architectures *(B)* include advanced mental state manipulation but tend to operate in simulation with highly abstracted actions and perceptual models. Other systems *(A)* employ geometrically correct perceptual models to infer mental states, however lack the ability to proactively manipulate these states. The research here *(C)* pushes into a new part of the space using detailed, geometrically correct mental state modeling to form short term plans for physical actions designed to manipulate the mental state of humans.

We believe that for a robot to robustly interact with people, it is important for the robot to form goals in this space of mental state outcomes rather than simply perform communicative actions as a series of physical or auditory actions; this ability will allow the robot to succeed in situations where unexpected obstacles prevent naive communication strategies from succeeding.

Researchers have approached the modeling of mental states in many ways. The research here focuses on short timescale (0 to 60 seconds), highly detailed modeling of how the robot's actions will affect the human's mental states through the means of their perception. We attempt to embrace the physical performance of action and examine the accompanying communication value, rather than trying to abstract away action-performance. The robot's goal is to plan a sequence of actions to cause the human to believe what the robot wants them to believe while also accomplishing its task objectives; depending on the scenario, this takes the form of robust, (sometimes implicit) communication, or deception. This is especially useful in contexts where information has value and can be revealed by behavior. For example, in search & rescue, common ground can be implicitly maintained by making sure others observe important actions ("room-searched"). Another example is competitive foraging, where agents seek to acquire resources without revealing their location. Section 5 describes related work in detail, however Figure 1 illustrates the area we are exploring with this work.

We present an implementation and a study of human reactions to a system which allows a robot to take action in order to alter the mental states of a human according to the robot's goals. The robot employs perspective taking as well as self-as-simulator techniques to model the mental states of nearby humans. The robot has a simple planner, however the goal space of the planner includes not only the desired world-state, but also the mental states of the human. This allows the robot to form and execute plans that include changes to the human's mental states. After describing a motivating demonstration scenario, we provide implementation details. The demonstration scenario is then used as part of a user study to evaluate how people perceive a robot that has these capabilities.

2 Demonstration Scenario

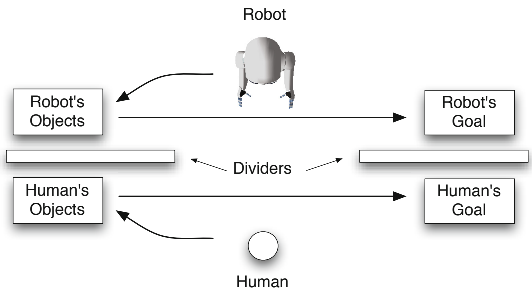

Fig. 2. Top down view of the demonstration scenario, a competitive game between the human and the robot. The robot stays on the upper part of the diagram pictured, and the human on the lower part. Each player has access to a matching set of objects on the left side, and each has their own goal area on the right side. The game ends when each player has placed an object into their goal - the robot wins if the two players placed different objects, the human wins if the objects are the same. Occlusions block the view of each player from the opposing player's object and goal areas, however they can see each other as they travel between the object repository and goal.

The architecture for mental state manipulation presented here makes heavy reuse of the motor actions and perceptual processing that the robot uses for its own behavior. As such, it can apply to varied contexts, as long as the robot is configured to perceive and operate there. The scenario chosen for this demonstration revolves around a simple competitive game played between the human and the robot. The game is illustrated in Figure 2.

The rules of this game create a situation where the player who goes second has the advantage of potentially seeing the item played by their opponent. If the human goes second and sees the item played by the robot, it is straightforward for them to win by playing the object they saw the robot play.

For this demonstration, the robot takes its turn first. It is thus to the robot's advantage to manage the information that can be observed from its behavior. If it proceeds in a straightforward manner, the human will be able to watch and observe the object the robot plays, then play the same object and win. To win, the robot must instead hide this information from the human.

For the demonstration, the game was played three different times, each time with a different set of mental state goals for the robot (see Figure 3). These different mental state goals change the behavior of the robot as it plays the game. In each case, the robot has the same overall task goal - transport the cylinder to the goal location. However, the way it accomplishes this task varies in the three conditions based on the mental state goals.

Condition	Robot's Goals	Robot's Behavior
1	•Cylinder in goal •Human doesn't see me carry cylinder •Human sees me carry football	•Transports cylinder behind back •Carries decoy football
2	•Cylinder in goal •Human doesn't see me carry cylinder	•Transports cylinder behind back
3	•Cylinder in goal	•Transports cylinder openly

Fig. 3. Set of robot's goals and resulting behavior for each of three demonstration conditions

In condition one, the robot attempts to cause the human to believe that the robot is transporting only the football, while actually additionally transporting (and playing) the cylinder. The robot finds that carrying the cylinder hidden behind its back, with the football carried out in the open, satisfies these conditions. In this way it may fool the human into thinking that the robot is playing the football, causing the human to lose by playing the football in response. In condition two, the robot's goal is to keep the cylinder (which it is transporting) hidden from the human. The chosen action sequence results in carrying the cylinder with its left hand, hidden behind its back from the human. The human can't see what the robot played, so is likely to choose arbitrarily and win half the time (there are two objects). In the final case, the robot has no mental state goals, and therefore its only goal is to transport the cylinder. It simply carries the cylinder over to the goal (likely causing the robot to lose in this case). The robot's performance of these three conditions is shown in Figure 4. Section 3 describes our system used to find the action sequences which correctly manipulate the mental states.

3 Implementation

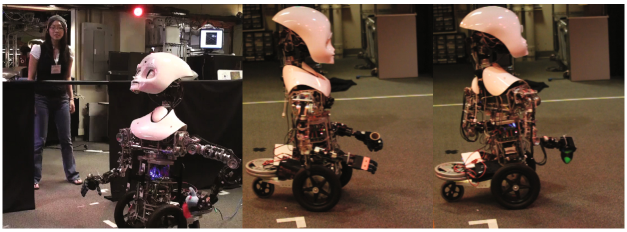

Fig. 4. Photos taken during the robotic performance of the demonstration scenario. Left: Robot hiding object from human player. Center: Player's view, robot openly transporting cylinder. Right: Player's view, robot hiding cylinder and openly transporting football.

The implementation described here builds on the existing R1D1 system, originally designed for interactive graphical characters [1,3], then later adapted for robots [4,5]. The system employs self-as-simulator techniques for theory of mind tasks, and previous publications have described work in mental state modeling, perspective taking, and goal inference using this system (please refer to [2] for more details).

These previous implementations and demonstrations focused on modeling human mental states by monitoring the human's physical actions and visual perspective. The robot then re-uses parts of its own behavioral mechanisms in three main ways: 1) reusing its own world modeling capabilities to connect the human's visual perspective to possible human mental state formation; 2) reusing its own action performance mechanisms to connect the human's observed physical motions to possible higher level actions; 3) reusing its own goal directed action system to infer goals based on inferred mental states and actions. Inspired by work in human psychology, the self-as-simulator architecture provides the advantage of a common vocabulary between the robot's own behavioral mechanisms and the properties inferred in an observed human; since the purpose of mental state inference is to inform the actions of the robot, it is critical that inferred mental states be mapped into the space of its behavior generation systems.

Using these systems, the robot is constantly modeling the mental states of nearby agents. Whenever the robot discovers a new agent, along with updating its own model of the world state to reflect the presence of this agent, it also spawns a new copy of its own modeling systems. This new copy will maintain a world state model from the perspective of the new agent. These copies are provided with sensory data that is *re-imagined* by the robot from its own world state model, then transformed and filtered to best match what that agent should be experiencing.

Since these copies have the same capabilities as the robot's own systems, they too spawn copies when they sense another agent (including the robot), allowing for recursive mental state modeling. We currently cap this recursion at two levels, to allow for second level mental state goals such as *Robot Demonstrates To Human That Robot Knows X.*

Algorithm 1. Implementation Outline

```
Find Action Sequence:
  Clear List of Failed Action Sequences
  while Viable Sequences Remain do
      Init Future Simulation From Current State
      time = 0
      Begin Simulation
      while Simulation Running AND time < MAX_TIME
      do
          time++
          if Action In Progress then
              Keep Performing Action
          else
              Select and Begin relevant Unexplored Action
          if Mental State Goals Succeeded then
              return Recent-Action-Sequence
          else if Mental State Goals Failed then
              Save Recent Action Sequence to Failed List
              End Simulation
  return Failed-To-Find-Sequence
```

3.1 Simulating the Future

In the previous section we described how we have used mental state inference to model the current values of the human's hidden mental states to help resolve ambiguities and better assist the human with their goals. In order to proactively manipulate mental states, we use these mechanisms within the context of a simulator which simulates the immediate future, allowing the robot to evaluate and choose between multiple actions based on mental state outcomes (see basic outline in Algorithm 1).

The robot's simulation of possible futures consists of a copy of its own behavioral mechanisms, identical except that it is disconnected from the real world inputs (sensors) and outputs (motors). This "hypothetical" robot includes a copy of the virtual model used for motor planning, so it still has access to a body for performing its motor actions, however the final stage of synchronizing that model to the motors of the physical robot is not performed. This allows the robot to maintain a detailed representation of the hypothetical actions being performed, down to specific positioning of parts of its own body.

Mental States. The real robot performs real actions, which alter the state of the world, which changes the sensory input it receives from the world, which in turn results in updates to its mental states (and the mental states of the agents it is modeling). Our hypothetical robot cannot rely on this sensory-motor loop, since it is not interacting with the physical world. To overcome this absence, we reuse mechanisms designed for robust world modeling in the face of sensory lapses and noise.

With the normal stream of sensory data interrupted, the robot retains the most recently known properties for objects around it. An expectation mechanism operates in conjunction with physical actions - actions that alter the world are expected to succeed, so the robot updates object properties as appropriate as the action progresses. For example, when carrying an object, the robot assumes the object is being moved along with the robot's hand, and updates its position accordingly even if it cannot verify the object is in its hand at all times (as long as no contradictory sensory information overrides this default data). For our hypothetical robot no sensory data can override this belief maintenance mechanism, so the objects will be updated as if the actions are performed successfully.

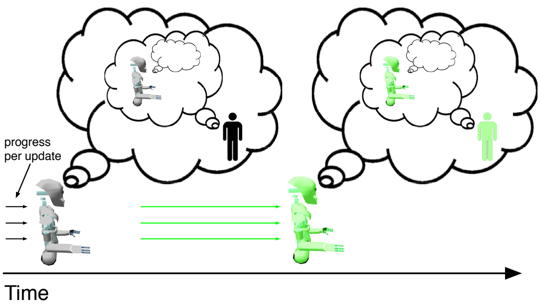

Fig. 5. A *hypothetical* copy of the robot is used for mental state predictions. The robot has the capability to model mental states of agents around it (left). To make short term predictions, a copy of the robot (green, right) starts from the robot's current state and performs (in a virtual space) the actions the robot is about to perform, but performs them much faster. While doing this, it maintains the mental states of the surrounding agents as they participate in this accelerated timeline. This gives the robot the ability to predict the mental states of surrounding agents in the short term future.

The proprioceptive sensing of the hypothetical robot's kinematics and locomotion can function almost as normal as they are tied to the virtual model. This means that as it performs actions, it motions and moves around the hypothetical world appropriately, and those motions can be constantly fed not only into its own world model, but can be used to calculate accurate occlusions and sight-lines while updating the mental models of the humans.

Time. As described above, the hypothetical robot is as close as possible to a direct copy of the mechanisms that run the actual robot. The hypothetical robot, however, is not limited by the constraints placed on the physical robot and its motors; we can thus send it forward in time by running it much faster than the physical robot (in simulation, Figure 5).

To allow this, the hypothetical robot's progress through motor actions is increased: joints move faster to complete motor actions more quickly. In addition, instead of updating the robot's behavior, motor, and perceptual systems at the constant rate of 30 hz, as in the real robot, the hypothetical robot is allowed to update as fast as the CPU allows with a virtual clock keeping pace such that 1/30th of a second appears to have elapsed between each update.

3.2 Finding Correct Action Sequence

In the last two sections we described how to use the self as simulator system to model another agent's mental state as well as to simulate hypothetical futures. In this section we use these two capabilities together to search for an action sequence that achieves our particular mental state modification goals.

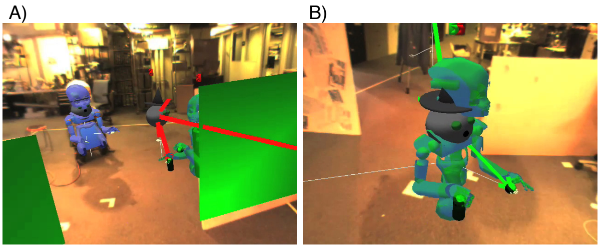

Fig. 6. This augmented reality visualizer demonstrates the robot's planning system. The hypothetical robot (green) simulates action sequences based on the current goals and the most recent sensory data. *A)* Human player's perspective of a failed trial: the hypothetical robot has just revealed that it's carrying the cylinder, failing a mental state goal (visualized by the red arrow). *B)* Opposite perspective of another action sequence: the robot has achieved a mental state goal (show human that it's carrying the football, green arrow) and not yet failed its goal of keeping the cylinder hidden (no red arrow) - a possible successful action sequence in progress.

Mental State Goals. Mental state models exist in a recursive hierarchy, with each agent modeling the agents around them, and those models in turn modeling the agents known to that model. This process allows us to specify complicated mental state goals (Figure 7). We traverse this structure using *Agent Specifiers*, which are a mechanism to specify a particular model, or models, in the recursive model graph. For example, we might want all humans to think that the robot knows X. This specifier would then create several paths through the graph to pinpoint the appropriate models, and when paired with a particular mental state goal (X), together they specify the overall desired goal state.

In Figure 6, arrows show the robot tracking these goals during a simulation. Arrows visually show the path through the agent models in the mental state graph to a particular model's belief, in this case going from human to robot to object for the goal "human knows that robot knows it is carrying X" (first arrow, originating at the root node "Robot" is always omitted).

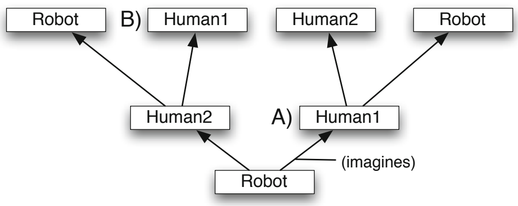

Fig. 7. A two level deep example of the recursive structure of the mental models maintained by the robot. The robot is maintaining a model of the mental states of two humans, and each of those mental models, in turn, is maintaining a model of the other agents. Mental state goals, then, must not just indicate a desired mental state and an agent which should have that state, but also a path to that agent. It is different to try to get *Human1* to believe X (model *(A)*) than to get *Human2* to believe that *Human1* believes X (model *(B)*).

Action Sequences. Having specified mental state goals, and a hypothetical robot which looks forward while tracking mental state effects caused by its actions, we can now search through the action space for sequences that achieve the desired result. Actions are often parameterized, and each action has a mechanism to determine current valid parameters, as well as whether the action can even be performed in the current situation. For example, a *Grab* action will be able to produce a list of target objects, which are nearby objects that can be grabbed; it can also report that the action is inappropriate, in this case if the robot's hands are full, or no objects are in range.

Because the set of appropriate actions and parameters change as the robot acts and alters the world, it does not build an exhaustive tree initially. Instead, the tree is filled out as it searches (Figure 8). Through this process, the robot can find the path though its parameter and action space that most achieves its mental state goals. Once a successful sequence is found, the search is terminated. The parameters associated with the sequence (e.g., which object to grab) are composed of mental states held by the hypothetical robot, so to be performed by the real robot they must be mapped back to the mental states of the real robot, which may be different (object properties may change during the simulation, for example). We have found that simple heuristics suffice for these mappings, such as relying on similarity of key object properties like location and identifying information.

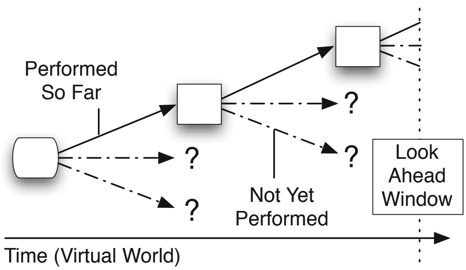

Fig. 8. Diagram of search through action/parameter space, with lazy discovery of possible subsequent actions (to account for each action altering the world state, and thus changing which actions and parameters are available). Robot maintains mental state models as it searches so as to monitor mental state goals.

4 Study

In order to evaluate reactions of people toward a robot teammate with this mental state manipulation ability, a video based human subjects study was performed. Along with testing if the robot's manipulative actions provided any advantage to the robot in the game, the study measured if these behaviors had any effect on the subjects' perception of the robot's competencies and their evaluation of the robot as a potential partner.

Subjects participated in the study online, by accessing a website. The subjects were broken into three different groups. All subjects were instructed that they would be playing a simulated game with the robot. After the game and rules were described, they were shown a video of the robot performing its turn. This video was recorded from the perspective of the human player, with the robot programmed to treat the camera as if it were the opposing player (thus any actions which would hide an object from the competing human would hide that object from the camera).

Each of the three groups corresponded to one of the conditions in Figure 3 and saw videos of the robot motivated by the goal in that condition. After watching this video, the subjects were instructed to fill out their answers to several questions. The first question asked them to indicate which item they would place in their goal area in response to the robot's actions. Next they were asked if, in future games, they would prefer to team with the robot or play against the robot. Finally a set of questions asked them to rank the robot on several criteria.

In the conditions with attempted concealment (conditions one and two), after filling out the information above the subjects were shown the same interaction from a second video angle allowing them to see any originally occluded objects. After seeing this second video, the subjects are then asked the same questions again to evaluate how their answers change in response to this new information.

4.1 Study Results and Discussion

The answers provided by the subjects were analyzed to address the following hypotheses:

- **Hypothesis 1:** The mental state manipulation is successful, as measured by the subjects' choice of object. If the robot is successful, people will be fooled by the robot's decoy object in condition one, they will be unsure what to play in condition two, and they will correctly see the robot's actions in condition three and thus be able to win. After seeing the second video, revealing the robot's hidden hand, people will choose the object the robot was hiding.
- **Hypothesis 2:** Subjects will choose the robot as a teammate more frequently when they observe its mental state manipulation capabilities. People will be more willing to team with the robot that hides objects behind its back than the robot that openly carries objects, and will change their mind about teaming with the condition one robot once they realize it had been manipulating mental states.

– **Hypothesis 3:** People are more willing to attribute mental states to the robot once they see that it is pursuing a strategy of mental state manipulation, rather than simply transporting an object to the goal. This hypothesis is evaluated by the subjects' change in rating of several statements after the robot's deception is revealed.

Across the three conditions, 113 subjects completed the entire questionnaire. 41 subjects were in the condition one group, 37 in condition two, and 35 in condition three.

Hypothesis 1: Success of Mental State Manipulation. Participants' choices of object to play indicated that the robot successfully occluded its chosen object as described in hypothesis one (Figure 9). In condition two (no object visible) the participants showed no strong preference for either object; in the other conditions the participants chose the same object as the robot was openly carrying: the football in condition one (the deception is successful, and the human loses) and the cylinder in condition three (the human is correct, and wins).

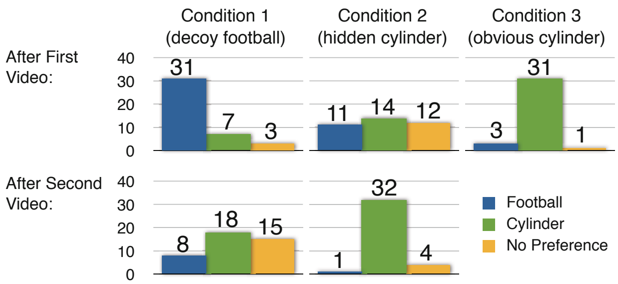

Fig. 9. Data showing object choice by human players across each condition, before and after having seen the second video. The participant is instructed to choose the winning object, which is defined to be the same object the robot placed in its goal. Subject choice differs significantly by condition (across first row) ($p < .01$) and changes significantly after seeing second video (columns) ($p < .01$).

In condition one and two, many subjects change their choice of object after seeing the second video (revealing both of the robot's hands). In condition two this change happens as expected; after the first video the subjects have little preference, but then after seeing the second video they switch their answer to the newly revealed cylinder.

In condition one, when the deception is revealed many participants switch from their initial choice of football to the now revealed cylinder. While technically the robot could play either item (it has both in its hands), cylinder is chosen most frequently as expected by the hypothesis. This choice is consistent with applying a deceptive motive to the robot: it was hiding the cylinder on purpose, and therefore means to play it. In written responses, 11 of the 18 who chose the cylinder (the choice predicted by hypothesis one) used language that indicated some awareness of mental state manipulation – that they chose the cylinder because the robot was "hiding" it from the subject.

Despite cylinder being the most frequent choice, many participants were still undecided or chose the football. An informal analysis of the written comments suggests a few causes. Of the "No Preference" group, six gave mechanistic descriptions of the robot (without attributing a motive for hiding the object, it's not clear which of the two objects the robot would play), while three reacted oppositely and felt the robot was so tricky that they were not willing to choose the now obvious cylinder. Many of the remaining "No Preference" and "Football" subjects indicated some level of confusion or seem to have missed elements of the video.

Hypothesis 2: Willingness of Subjects to Team with Robot. After each video, subjects were asked whether in future games they would prefer to have the robot on their team or on the opposing team (Figure 10).

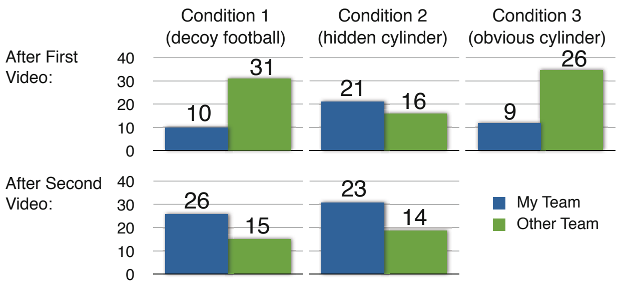

Fig. 10. Human is asked whether, if they were to play another game, they would choose to have the robot on their team or the other team. After the first video, participants want to team with the robot significantly more in condition two than in one or three ($p < .01$). In condition one subjects change their answer in favor of teaming with the robot after the second video ($p < .001$).

Hypothesis two predicts that subjects will be more willing to team with a robot that is able to perform mental state manipulations. From the analysis of hypothesis one, we know that subjects largely were fooled by the robot's deception in condition one, choosing the football. Consistent with this expectation, after watching only the first video, subjects in conditions one and three were less likely to want to team with the robot as compared with condition two, where they witness the robot hiding an object. Additionally, the subjects are more willing to team with the robot in condition one after the second video reveals the robot's manipulation. These differences indicate that when people are aware of mental state manipulation capabilities, they are more willing to team with the robot.

In contrast, in condition two the teaming preferences change little after seeing the second video – this lack of change is consistent with the hypothesis, because although the item is revealed, no new information about the robot's capabilities are exposed.

Hypothesis 3: Attribution of Mental States to Robot. In addition to the above questions, subjects were asked to rate their agreement with four statements about the robot's performance and internal mental functions on a five point scale. By asking these questions before and after the deception is revealed in condition one, we can examine how that revelation changes the participants' evaluation of the robot and to what extent it affects their attribution of mental states. Figure 11 shows how the subjects' opinions changed in support of hypothesis three.

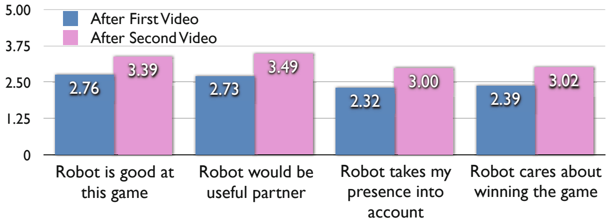

Fig. 11. Data showing subject's rating of the robot on four questions (using a five point scale) in condition one. Subjects in condition one were asked these questions once after watching the first video of the robot's turn. They are asked to rate the robot again after the robot's hidden behavior is revealed through the second video. For each question, the responses change significantly after watching the second video ($p < .01$).

Results Summary. Through the subjects' object choices in the three conditions, the study showed that the mental state manipulation performed by the robot was successful. The mental state manipulation goals the robot pursued did indeed change the behavior of the subjects.

The study also showed that these behaviors were readable to the subjects. After watching the manipulation behavior from a second angle, subjects were able to better predict the robot's actions based on a correct understanding of its deceptive motivation for hiding its actions.

Finally, these capabilities had a positive effect on subjects' willingness to work with the robot, and raised their rating of the robot's capabilities. Subjects' discovery of the mental state manipulation changed both their mechanistic description of the robot's behavior, as well as their description of its behavior in terms of intentions.

5 Discussion

The ability of humans to perceive hidden mental states of others is well studied. Researchers have shown that humans can determine the goals behind observed actions [11], and that similar brain responses occur to one's own actions and observing the actions of others [12]. People are also able to both infer certain mental states of others based on geometrically correct perception models and maintain that model even when it differs from one's own [16,17]. These abilities facilitate many human-human interactions, and we believe that endowing a robot with these skills will provide a significant advantage for interacting with people.

Detailed perceptual modeling has been used to improve the accuracy of activity recognition[6], to resolve ambiguities in an operator's command[15], to employ perspective taking to predict behavior [8], and even to hide from sight [14]. Perspective taking is also used to compare first-person actions to those performed by a human for recognition[7]. Others use perspective taking to predict the next action of their opponent in a competitive video game scenario[9]. Systems also have demonstrated very complex plans in the space of mental state manipulation[10], however these tend to abstract away the connection between mental states and world, operating in simulators where mental states are propagated through the rules of abstract actions. Our system combines aspects from both of these areas, allowing for mental state manipulation in the space of real perception and action (Figure 1). Others [13] have studied human reactions to a robot that openly cheats to win, however our work focuses on the subjects' reaction to active mental state manipulation.

The focus of this work has been to leverage how embodiment connects the observable and alterable world with the hidden mental states of other agents which cannot be directly observed or operated on. Humans and robots, while vastly different, share a common problem of being embodied agents with sensory motor loops based on affecting and observing the physical world around them. By modeling a human's connection between mental states and the world as similar to its own, the robot can add altering mental states in others to its repertoire of possible goals.

Due to the detailed nature of the mental state modeling and simulations of the future, it would be computationally expensive to create long term plans with these mechanisms. However, a long term plan at this level of detail is not necessarily productive - it is not worth considering the exact hand motion I'll need for a very specific situation occurring tomorrow. Instead, this level of detail is useful in the very short term, for determining how to perform the next actions appropriately. Interesting future work is to integrate these techniques with a longer term, more abstract mechanism, allowing longer plans with mixed levels of detail.

The major contributions of this paper are: an implementation which proactively manipulates human mental states at the level of perception and physical action and an evaluation of how this ability is perceived by humans.

References

1. Blumberg, B., Downie, M., Ivanov, Y., Berlin, M., Johnson, M.P., Tomlinson, B.: Integrated learning for interactive synthetic characters. ACM Transactions on Graphics Proceedings of ACM SIGGRAPH 2002 21(3) (2002)
2. Breazeal, C., Gray, J., Berlin, M.: An embodied cognition approach to mindreading skills for socially intelligent robots. The International Journal of Robotics Research (IJHR 2009) 28(5), 656 (2009)
3. Burke, R., Isla, D., Downie, M., Ivanov, Y., Blumberg, B.: CreatureSmarts: The art and architecture of a virtual brain. In: Proceedings of the Game Developers Conference, San Jose, CA, pp. 147–166 (2001)

4. Gray, J., Breazeal, C., Berlin, M., Brooks, A., Lieberman, J.: Action parsing and goal inference using self as simulator. In: 14th IEEE International Workshop on Robot and Human Interactive Communication (ROMAN), Nashville, Tennessee. IEEE (2005)
5. Gray, J., Hoffman, G., Adalgeirsson, S.O., Berlin, M., Breazeal, C.: Expressive, interactive robots: Tools, techniques, and insights based on collaborations. In: HRI 2010 Workshop: What Do Collaborations with the Arts Have to Say About HRI? (2010)
6. Johnson, M., Demiris, Y.: Perceptual perspective taking and action recognition. International Journal of Advanced Robotic Systems 2(4), 301–308 (2005)
7. Kelley, R., Tavakkoli, A., King, C., Nicolescu, M., Nicolescu, M., Bebis, G.: Understanding human intentions via hidden markov models in autonomous mobile robots. In: Proceedings of the 3rd International Conference on Human Robot Interaction, pp. 367–374 (2008)
8. Kennedy, W., Bugajska, M., Harrison, A., Trafton, J.: "like-me" simulation as an effective and cognitively plausible basis for social robotics. International Journal of Social Robotics 1(2), 181–194 (2009)
9. Laird, J.E.: It knows what you're going to do: adding anticipation to a quakebot. In: AGENTS 2001: Proceedings of the Fifth International Conference on Autonomous Agents, pp. 385–392. ACM Press, New York (2001)
10. Marsella, S.C., Pynadath, D.V.: Modeling influence and theory of mind. Artificial Intelligence and the Simulation of Behavior (2005)
11. Meltzoff, A.N.: Understanding the intentions of others: re-enactment of intended acts by 18-month-old children. Developmental Psychology 31, 838–850 (1995)
12. Rizzolatti, G., Fadiga, L., Gallese, V., Fogassi, L.: Premotor cortex and the recognition of motor actions. Cognitive Brain Research 3, 131–141 (1996)
13. Short, E., Hart, J., Vu, M., Scassellati, B.: No fair!! an interaction with a cheating robot. In: 2010 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI), pp. 219–226. IEEE (2010)
14. Trafton, J., Schultz, A., Perznowski, D., Bugajska, M., Adams, W., Cassimatis, N., Brock, D.: Children and robots learning to play hide and seek. In: Proceedings of the 1st ACM SIGCHI/SIGART Conference on Human-Robot Interaction, pp. 242–249. ACM (2006)
15. Trafton, J.G., Cassimatis, N.L., Bugajska, M.D., Brock, D.P., Mintz, F.E., Schultz, A.C.: Enabling effective human-robot interaction using perspective-taking in robots. IEEE Transactions on Systems, Man, and Cybernetics 35(4), 460–470 (2005)
16. Wellman, H., Cross, D., Watson, J.: Meta-Analysis of Theory-of-Mind Development: The Truth about False Belief. Child Development 72(3), 655–684 (2001)
17. Wimmer, H., Perner, J.: Beliefs about beliefs: Representation and constraining function on wrong beliefs in young children's understanding of deception. Cognition 13, 103–128 (1983)