

Adapting to Agents' Personalities in Negotiation

Shavit Talman
Dept. of Computer Science
Bar Ilan University
Ramat Gan 52900, Israel
talman@macs.biu.ac.il

Ya'akov Gal
Division of Engineering and
Applied Sciences
Harvard University
Cambridge, MA 02138
gal@eecs.harvard.edu

Meirav Hadad
Caesarea Rothschild Institute
University of Haifa
Haifa 31905, Israel
meirav@cs.haifa.ac.il

Sarit Kraus^{*}
Dept. of Computer Science
Bar Ilan University
Ramat Gan 52900, Israel
sarit@macs.biu.ac.il

Categories and Subject Descriptors

I.2 [Artificial Intelligence]: I.2.11 Distributed Artificial Intelligence
Multiagent systems

General Terms

Design, Experimentation

Keywords

Negotiation, Social decision-making, Opponent modeling, Game formalisms

ABSTRACT

To establish cooperative relationships, agents must be willing to engage in helpful behavior and to keep their commitments to other agents. However, in uncertain and dynamic environments, it is difficult to identify the degree of helpfulness of other agents. This paper describes a model in which agents' helpfulness is characterized in terms of cooperation and reliability. An agent chooses an action based on its estimate of others' degree of helpfulness given the dependency relationships that hold between the agent and others. This model was evaluated in a negotiation game in which players needed to exchange resources to reach their goals, but did not have information about each others' resources. Results showed that agents using the model could identify and adapt to others' varying degree of helpfulness even while the other agents were constantly changing their strategy. Moreover, agents that varied their degree of helpfulness depending on their estimate of others' helpfulness

^{*}Sarit Kraus is also affiliated with the Institute for Advanced Computer Studies at the University of Maryland (UMIACS).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'05, July 25-29, 2005, Utrecht, Netherlands.
Copyright 2005 ACM 1-59593-094-9/05/0007 ...\$5.00.

outperformed agents who did not, as well as increased the social welfare of the group.

1. INTRODUCTION

When making decisions in strategic environments, self-interested agents need to reason about the behavior of other agents as well as the effects of their own actions. In particular, when agents depend on each other to achieve their individual goals, agents' success is contingent on their ability to cooperate with others and their willingness to perform actions that mutually benefit each other. Cooperation can be established by negotiating to trade resources which agents require from each other. However, deceitful agents can take advantage of agents that are helpful to them and others, for example, by renegeing on a commitment after having received their share of an agreed upon exchange. To succeed, agents must be able to identify and negotiate with those who are cooperative while avoiding those who are exploiters. Identifying agents' degree of helpfulness is particularly important in open systems, which consist of multiple agents with no central control of design or behavior.

One approach to this challenge is to use normative models of decision making such as Game Theory to guide agents' behavior [8]. These models prescribe optimal strategies for agents that take into account the effects of their decisions on each other, and they guarantee that no agent is exploited by another. However, in many settings self-interested agents who follow these strategies do not cooperate. In addition, these models assume all agents deliberate in the same fashion about the game and completely adhere to their prescribed strategies, a difficult assumption to make in open systems.

In behavioral economics and social psychology, personality models have been used to explain and predict human behavior across different environments and contexts. In particular, it has been shown that people's personality traits affect their level of cooperation in negotiation in games such as the ultimatum game [1]. This paper presents a decision-theoretic model which explicitly represents and reasons about agents' personalities in environments in which agents are uncertain about each others' resources.

The model characterizes personality along two dimensions: cooperation (the tendency to propose mutually beneficial exchanges

of resources) and reliability (the tendency to fulfill commitments). When deliberating the action to take next, agents reason about their own personality as well as their beliefs about the personality of others. They alternate between updating their model of others and using the model to come up with the best action to perform.

We show that agents that use this model are able to identify those that are helpful and reciprocate their behavior, while staying clear of exploiters. We also show that the model is robust; it can adapt to agents that change their own behavior over time as well as to varying types of environments.

We tested the model on a negotiation game played by agents which were created by the experimenters as well as by agents which were created by others, who were not required to design their agents in any specific manner. The experiments varied the complexity of the game as well as the number of players and the dependency relationship between players. Results show that agents that reasoned about personality were able to identify the cooperation and reliability measures of others; they reciprocated the behavior of helpful agents, while avoiding agents that were deceitful or unhelpful. Agents that adapted their own personality, based on their estimate of the personality of others outperformed agents that did not, and improved the social welfare of all agents in the system.

2. RELATED WORK

Previous research in the multi-agent systems literature has proposed models for social reasoning, which take into account others' preferences when deliberating about their actions. Hogg and Jennings [6] proposed a model in which agents' utilities were a weighted summation of each others' expected outcomes. By learning these weights from observations, agents changed their measure of helpfulness over time. When all agents in the system were adaptive, high exploration rates led agents to seek out new negotiation opportunities and increased the overall social welfare of the group. Sen and Dutta [9] investigated the effect of agents' helpfulness, and testimonies of others' level of helpfulness on the performance of the group. Agents accepted or declined others' requests depending on their past experiences as well as the cost of helping. Agents that weighed the testimony of others based on their reputation were more likely to succeed than those that did not consider the reputation of others. They were also able to avoid deceitful agents and to cooperate with helpful agents. Zhang *et al.* [7] explored the trade-off between selfishness and benevolence in environments in which agents were uncertain about the helpful nature of others in the system. They showed that although realizing every opportunity for cooperation was impossible, selfish agents do better than helpful agents as the rate of uncertainty in the system grows.

All these models allowed agents to change their measure of helpfulness over time as a function of their model of others, and investigated the effects of this behavior on agents' cooperation and system performance. However, the utility function of all agents in the system was common knowledge. In addition, the size of the domain used to evaluate these models was limited in the strategy space for players as well in the number of possible moves.

The model presented in this paper does not assume or control the form of others' utility functions and is thus suitable for open systems, in which there is no knowledge about any aspect of agents' design mechanisms, including how agents benefit from potential actions of other agents. Our aim was to construct a model which agents could use to outperform other agents, regardless of their helpfulness measure. In our domain, agents needed to choose among hundreds of potential actions at each move of the game and the total number of moves in the game depended on the agents' performance. Because we did not control the design of other agents, there

was uncertainty not only over others' resources, but also over their measure of helpfulness.

Castelfranchi *et al.* [3, 2] proposed a programming language for comparing the behavior of agents with different personality traits where agents must cooperate to achieve their goals. The focus of this work was to investigate which personality combinations facilitate cooperation in a blocks-world domain. Agents' personalities were represented as logic clauses, stating the conditions under which agents help each other, and request help from others. Agents could attempt to deceive others by declaring a false personality, but the strategy of each agent was constant and fully determined by that agent's personality. In contrast, our agents strategies depended on their estimate of the personality of others as well as their own. Also, agents could vary their personality depending on their estimate of others. Lastly, in our environment agents were constantly changing their strategies, as a function of their model.

3. THE COLORED TRAILS FORMALISM

This study used the Colored Trails (CT) game, designed by Grosz and Kraus [5]. CT is a framework for investigating decision-making processes of agents in contexts in which their outcome depends on each others' actions. The game parameters may be set to vary such environmental features as task complexity, the resources available to agents, agents' capabilities, and the dependency relationships between agents. CT provides a clear analog between the properties of the game and real-world task and resources, making it reasonable to assume that results obtained using the CT framework will generalize to other domains.

CT is played on an $N \times M$ board of colored squares. One square is designated as the "goal square" and each player has a piece on the board, initially located in one of the non-goal squares. Each player also has a set of colored chips, whose colors are chosen from the same palette as the squares. To move a piece into an adjacent square a player must turn in a chip of the same color as the square. Chips may be exchanged by the players, and the conditions of exchange may be varied to model different decision-making situations.

A player's performance in CT is determined by a scoring function, which is computed when the player is declared "out-of-game". This function may depend on many factors, such as the player's distance from the goal-square, the number of moves made, and the number of chips the player possesses at the end of the game. In addition, a player's performance in the game can be made to depend on the performance of other players, by including the score of other players in its scoring function.

For our study, we used a version of CT in which two or four players played on boards of varying sizes. Each player had knowledge of the scoring function and full view of the board but could not see the other player's chips.

The game protocol comprised two phases, a communication phase and a movement phase. During the communication phase, new exchanges could be proposed, pending proposals could be replied to, and chips could be sent from player to player. Agreements reached during the communication phase were not enforced by the game controller, allowing agents to deceive each other. During the movement phase, the game controller automatically advanced each player one square closer along the shortest path to the goal given its chips and the board configuration.

A player was declared "out-of-game" if it reached the goal state or if it stayed dormant for 3 moves, at which point its score, denoted $score_i$, was computed. Each player's outcome depended solely on its own performance. The scoring rule used a multi-attribute linear function which incorporated three factors: (1) whether the player reached the goal square; (2) the distance of the player from the

goal square, measured by the Manhattan distance; (3) the number of chips the player possessed at the end of the game.

4. MODEL CONSTRUCTION

We wanted our model to be able to generalize to games of varying complexity measured by the board game size, the number of players, and the dependency relationships between players. In addition, we wanted the model to be able to perform well in systems characterized by uncertainty over others' resources and utility functions. Our approach was to explicitly represent agents' helpfulness in our model. We described agents' helpfulness in terms of personality traits along two dimensions.

- Cooperation (c): a measure of agents' willingness to share resources with others in the game by initiating proposals and agreeing to proposals of others.
- Reliability (r): a measure of agents' willingness to keep their commitments in the game by delivering the chips they had agreed to.

Agents that cooperated more than 50% of the time were regarded as highly cooperative, and agents that reneged on their commitments less than 20% of the time were regarded as highly reliable. We defined three discrete types for measuring reliability and cooperation along the range $[0, 1)$.

- low-cooperation: $[0, 0.3)$; low-reliability: $[0, 0.4)$
- medium-cooperation: $[0.3, 0.6)$;
medium-reliability: $[0.4, 0.8)$
- high-cooperation: $[0.6, 1.0)$; high-reliability: $[0.8, 1.0)$

Agents' personality traits were referred to by a pair, representing their cooperation and reliability levels respectively. For example, an agent whose behavior exhibited low cooperation and medium reliability was referred to as a low-cooperation, medium-reliability (or LM) type agent.

According to the model, agent i 's expected utility of taking action a directed at agent j depended on the following features:

- the personality of agent i , denoted P_i ;
- agent i 's estimate of the personality of agent j , denoted P_j ;
- the expected value of taking action a given the state of the game s , denoted $EV_i(a | s)$;
- the expected cost of future ramifications of taking action a , denoted $EC_i(a)$.

Computing the terms $EV_i(a | s)$ and $EC_i(a)$ required an estimate of the likelihood of agent i reaching the goal in future moves of the game. This was difficult to compute because reaching the goal depended on the future actions of players in the game which in turn depend on their resources which were unknown to agent i . We approximated the probability $P(RG | s)$ of reaching the goal, while at state s , simply as $1 - \frac{\#cn}{M+N}$ where $\#cn$ is the number of chips the agent lacks to get the goal at state s and M and N are the game board dimensions. We then computed

$$EV_i(a | s) = P(RG | s) \cdot RGwt + score_i - e^{\#np}$$

where $RGwt$ is a constant representing a bonus for reaching the goal, $score_i$ is the score of player i at state s and the term $e^{\#np}$ incurs a punishment for remaining idle, where $\#np$ equals the number of consecutive idle turns the player has had in the game.

To compute the expected ramification cost EC_i , we assumed that selfish actions of a player were punished by others, and that considerate actions were rewarded by others. For any proposal made by agent i to agent j , we defined a *selfish exchange* of i to be any exchange that is more advantageous to i than to j and a *considerate exchange* of i to be any exchange that is more advantageous to j than to i . We defined a *feasible exchange* of i to be any exchange for which i possesses the chips to complete. All exchanges were assumed to be beneficial to both agents if accepted by agent j .

We estimated $EC_i(a)$ to equal some constant integer t when a was either (1) a considerate exchange of agent i ; (2) an agreement of i to a feasible exchange; (3) a transfer of chips to fulfill a promise of exchange. We estimated $EC_i(a)$ to equal $-t$ when a was either (1) a selfish exchange for agent i ; (2) an agreement of i to an infeasible exchange; (3) reneging on a promise to send chips.

Agents used the following multi-attribute linear utility function with weights w_1, \dots, w_4 to make their decisions, where we used object-oriented notation to denote the measures of personality traits.

$$u_i(a, j, s) = w_1 \cdot (P_j.c + P_i.c) + w_2 \cdot (P_j.r + P_i.r) + w_3 \cdot EV_i(a | s) + w_4 \cdot EC_i(a)$$

For example, $P_j.c$ referred to agent i 's estimate of the cooperation measure for agent j .

The value of the weights of the components of the utility function depended on agents' personality traits. LC and LR type personalities assigned a higher weight to the expected value EV_i , leading them to adopt behavior that was unhelpful. For example, LC agent types only proposed selfish exchanges and LR type agents never fulfilled their commitments to send chips after having received their share. For medium- and high-cooperation (and reliability) type agents, the weights were tuned empirically.

Given a model that represents the personalities of each opponent, the next step is to show how an agent who used this model behaved in the game. For each opponent j , an optimal action a_j^* maximized the utility of agent i at state s if

$$a_j^* \in \operatorname{argmax}_a u_i(a, j, s)$$

This equation uses set membership rather than equality to indicate that there may be more than one action that maximizes the agent's utility. An optimal action set for i , denoted A_s is the set of all optimal actions of agent i at state s . It includes at least one action for every opponent. Two actions a_1 and a_2 are said to contradict if performing a_1 would require at least one chip that was needed to perform a_2 . For each subset of A_s in which no two actions contradict, agent i summed up its utility for performing every action. Agent i chose the subset of actions that together yielded maximum utility, and performed every action in that subset at the onset of each communication phase of the game.

To update the personality model after each communication phase k , we computed the reliability level of agent j at phase k , denoted $P_j.r^k$, to be the fraction of times that agent j met its commitments at phase k . The cooperation level of agent j at phase k , denoted $P_j.c^k$, was the fraction of times that agent j proposed exchanges in which it offered chips to others. To update $P_j.r$, we computed $\frac{1}{k} \sum_{l=1}^k \delta^{k-l} P_j.r^l$, where δ is a discount factor. We updated $P_j.c$ in a similar manner.

Using this model, agents exhibited distinct behaviors for different personality traits. The effect of an agent's reliability level on its own behavior was as follows: Low-reliability type agents never kept their commitments to others. Medium-reliability type agents were (1) more likely to keep their commitments to medium- and high-reliability type agents than they were to low-reliability type agents; (2) less likely to keep their commitments to medium- and

high-cooperation type agents then they were to low-cooperation type agents. High-reliability type agents always kept their commitments, regardless of the personality type of the other.

The effect of an agent’s cooperation level on its own behavior according to the model was more complex, and depended on the nature of the dependency relationship between players. Player i was said to be *task dependent* on player j when player i lacked some chips it needed to reach its goal and depended on some player j , who possessed these chips, to supply them. A player was said to be *task independent* if it possessed all the chips it needed to get to the goal. Table 1 presents the behavior that was associated with agents’ level of cooperation given their task dependency type. For example, when it was task independent of others, a high-cooperation agent type was helpful; it accepted any exchange. Medium- and high-cooperation agent types proposed exchanges whose benefit to the recipient agent was correlated with their estimate of its level of cooperation; the higher this estimate, the higher the benefit of the exchange to the recipient agent. This is denoted as “cooperation-dependent” behavior in Table 1. Also shown in the table is that when agents were task dependent on each other, medium- and high-cooperation type agents were more helpful to cooperative agents than to others.

5. EXPERIMENTAL DESIGN AND ANALYSIS

We used two classes of agents in our study. The first consisted of Multiple-Personality (MP) and Single-Personality (SP) agents, which were designed by the experimenters. Both MP and SP class agents modeled the personality traits of each opponent as described above. The personality traits of an SP agent included constant cooperation and reliability levels, whereas an MP agent adopted different measures of cooperation and reliability for each personality type of its opponents based on a matching scheme described below.

Both MP and SP agents were adaptive: they changed their behavior as a function of their estimate of others’ measure of helpfulness, given the history of their observations. However, the MP agent adopted a unique personality type for each player. For example, if both MP and SP agents estimated a player to be highly cooperative, a low-cooperation type SP agent would avoid it (see Table 1), while an MP agent could choose to interact with it, provided that it matched a high-cooperation type agent with a high-cooperation type personality.

The second class of agents we used was Peer-Designed (PD) agents. To create a truly open system, the designers of these agents were graduate-level computer science students at Bar Ilan University who were not given any explicit instructions regarding the design of agents’ decision-making strategies. In particular, the utility function that guided PD agents in their play was not known to the experimenters.

We classified PD and SP agents as either “helpful” or “unhelpful”, based on a preliminary evaluation: Helpful PD agents were those who mainly engaged in *reciprocal-type* exchanges, in which chips were both received and sent by the PDs; unhelpful PD agents were those agents who engaged in (1) *take-type* exchanges in which players received chips, but did not give out chips, or (2) were idle and did not engage in any negotiation with others. Helpful SP agents exhibited medium- and high-cooperation and reliability type personalities, while unhelpful SPs exhibited low-cooperation and reliability type personalities.

We expected helpful agents to be able to realize opportunities for exchange with each other more often than unhelpful agents and to exceed them in performance, as measured by the score in

Opponent Personality Type	Personality Matched by MP agent
LL	LL
LM	LM
LH	LM
MM	LM
MH	MM
HM	MM
HH	MM

Table 2: Matching Table for MP agent by Personality Type (cooperation and reliability measure)

the game. We also expected that in some cases, unhelpful agents would be able to take advantage of the vulnerability of those helpful agents who allow themselves to be exploited. We hypothesized that the MP agent would be able to identify and reciprocate helpful agents more quickly than SP or PD agents, while avoiding those who would exploit it. As a result, the MP agent would perform better than all other agents in the game.

We used the following approach to come up with a matching scheme for the MP agent in order to assign a different personality type for each agent type. We ran a series of games in which SP agents played each other using 30 2-player boards, which varied every possible task dependency combination between two players. Each SP agent played multiple games against each of the other possible SP types. We matched each SP agent’s personality with the personality of the opponent that resulted in the highest average score for the SP agent, as described in Table 2. Using the table, the MP agent varied its own personality traits, based on its estimates of the personality traits of others,

This matching scheme adopted the right “balance” between helpful and selfish behavior. Low-cooperation type agents were matched with low-cooperation type personalities. As a result, the MP agent could avoid them and keep from getting taken advantage of. Medium and high-reliability type agents were matched with medium-reliability type personalities, so that the MP agent kept its commitments as long as others kept theirs. Some of the matchings were not intuitive. For example, an high-cooperation high-reliability (HH) type SP agent was matched with a medium-reliability personality.

The experiments used a single MP agent, 7 SP agents with personality traits LL, LM, LH, MM, HM and HH, and 10 PD agents, all of which achieved the highest score in a preliminary evaluation. We report on the performance of agents in the system by comparing their scores and behavior across different settings of the game, which varied the dependency relationships between players as well as the number of players. All results are statistically significant in the 95% confidence interval range unless indicated otherwise.

5.1 Repeated Game Settings

We evaluated the MP agent by playing a series of repeated games with the other agents in the systems. We allowed agents to update their model of others from game to game. Each agent’s final outcome was the aggregate of its scores in all of the games it participated in. We expected the MP agent to score higher than helpful and unhelpful agents in each of the games played, and that the rate of increase in the score of the MP agent from game to game would be significantly higher than for those agents that also increased their score. Also, we expected that helpful SP and PD agents would score higher than unhelpful SP and PD agents in each game, and would improve their performance from game to game. Lastly, we

Cooperation Level	Personality types negotiated with	Types of Exchanges Accepted	Types of Exchanges Proposed
Low	none / all	selfish / selfish	none / selfish
Medium	high-reliability / medium- and high reliability	selfish/ cooperation-dependent	cooperation-dependent / cooperation-dependent
High	all / medium- and high- cooperation and reliability	any beneficial / cooperation-dependent	cooperation-dependent / cooperation-dependent

Table 1: Behavior by Cooperation Level for Task Independent/Dependent Players

expected that both MP and helpful agents, when playing together, would score more than when playing with unhelpful agents, that, as described in Table 1, do not generally negotiate with others. This would prevent all agents from realizing beneficial opportunities for exchanges in the game.

The evaluation used two types of game boards, in which all players were task dependent on each other (allDep board), and in which one player was task independent and the other players were task dependent on it (oneSelf board). In the experiment we executed 5,040 games, played in 1,080 rounds of three consecutive games each. The board games we used in each round alternated between the (oneSelf, allDep, oneSelf) boards and (allDep, oneSelf, allDep) boards. The players in each game included a MP agent, two SP agents, and one of the PD agents. Each group of four players played all possible task dependency roles, to control for any effect brought about by dependency relationships. Table 3 presents the average score for the MP agent when playing against helpful and unhelpful agents across all games. The scores reported in the table sum over the other players in the game.

	MP agent	PD and SP agents
Helpful	170.6	114.8
Unhelpful	142.5	98.2

Table 3: Average performance of MP agent against helpful/unhelpful agents (3 repeated games)

As expected, the average score achieved by the MP agent was significantly higher than all other agents, regardless of their level of helpfulness. Also, the MP agent’s score when playing against helpful agents (170.6) was higher than its score when playing against unhelpful agents (142.5). Helpful agents also benefited from cooperating with the MP agent: their performance was significantly higher than their unhelpful counterparts (114.8 vs. 98.2).

To show that the MP agent established a cooperative relationship with helpful agents, while staying clear of unhelpful agents, we examined the fraction of reciprocal exchanges the MP engaged in with others, as well as the fraction of turns it was idle and did not offer any exchange. Results are shown in Table 4 for both the MP agent and SP agents. They confirm that the percentage of reciprocal exchanges between an MP agent and helpful agents (60%) were significantly higher than that of the MP agent and unhelpful agents (20%), while the fraction of turns the MP agent remained idle when playing with unhelpful agents (39%) was significantly higher than with regard to helpful agents (25%). This also proves that the MP agent avoided agents who were exploitive and kept them from taking advantage of it.

Exchange Type	Helpful agents	Unhelpful agents
Reciprocal	60%	25%
Idle	20%	39%

Table 4: Percentage of exchange types proposed by MP agent

	Game 1	Game 2	Game 3
MP agent	149.87	153.06	154.74
Helpful PDs	115.53	115.92	112.67
Unhelpful PDs	116.85	107.29	102.94
Helpful SPs	117.26	107.93	116.28
Unhelpful SPs	95.33	85.96	91.46

Table 5: Agent performance by game

In contrast, helpful SP agents were more likely to engage in reciprocal exchanges (65%) and far less likely to remain idle towards others (1%), regardless of their class, indicating that they were more vulnerable to exploitation. Note that percentages do not add up to 100% because we have left out exchanges which were not reciprocal.

Table 5 shows the average performance for each game. As expected, the performance of the MP agent increased from game to game, while the performance of unhelpful PD agents decreased from game to game. This result is supported by the fact that the MP agent avoided interacting with unhelpful agents, as shown in Table 4. Unhelpful PDs do worse from game to game, and are always worse off than helpful PDs, indicating that the MP agent successfully adapted to these agents. Although the performance of the MP agent was consistently better than the performance of all agents in every game, we were surprised that the performance of both helpful and unhelpful SP agents increased from game 2 to game 3. We hypothesized that some of the SP agents were exploiting the MP agent and as a result, increasing their average score.

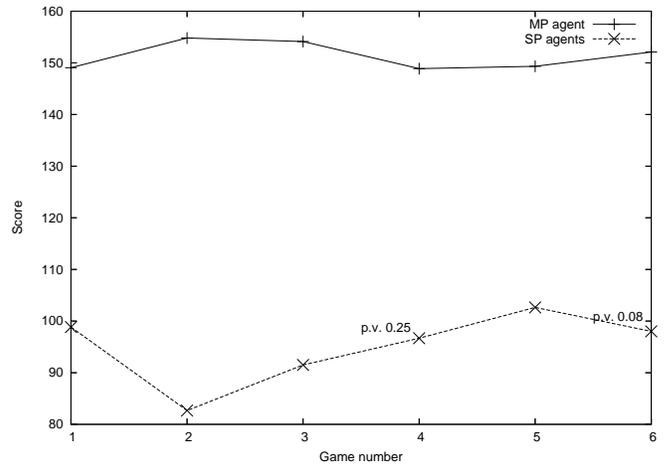


Figure 1: Performance of MP agent vs. SP agent (6 repeated games)

To evaluate this hypothesis, we ran six repeated games played on 4-player board games in which the MP agent played against three other SP agent types. Figure 1 shows average performance of MP vs. SP agents game across the six game series. The significant differences in score between games is labeled whenever the difference

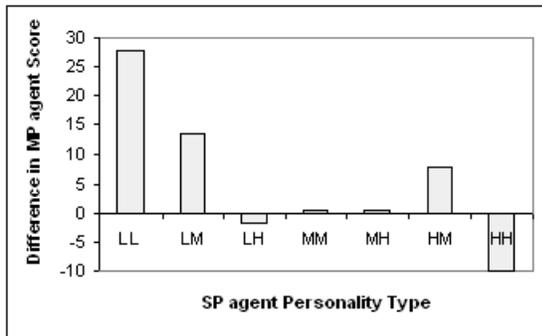
was not in the 95% confidence interval. As shown in the figure, the score of the MP agent in game 2 was higher than its score in game 1, while the score of the SP agents in game 2 was lower than their score in game 1. This is consistent with the results described in Table 5. However, from game 2 to game 4, the score of the MP agent decreased while the score of the SP agents increased monotonically. It seemed that some SP agents avoided detection and were able to exploit the MP agent in games 2 to 4.

To find out which of the SP agent types exploited the MP agent, we analyzed agents' performance according to their types. Figure 2 depicts the difference in score for the MP agent between game 2 and 4 when playing against each SP personality type. As shown, the MP agent suffered a significant decrease in score when playing against unhelpful SP agent types LL and LM. This was to be expected, because these agents do not communicate with others, and it is difficult to identify their true nature. However, the MP agent also decreased performance when playing the *helpful* agent type HM, leading to the conclusion that the MP agent could not identify its personality in this game.

This finding was supported by a separate experiment, in which we ran a series of games in which an HH SP agent type played against every other SP agent type on each of the boards. We used a HH type agent for this experiment because such a player communicates a lot with other player. Therefore, these results represent an upper bound on the MP agent's ability to identify its opponents. We recorded the HH SP agent's success rate for every personality type of SP agent. The highest error rate in identification (35%) occurred when the HH type agent tried to identify an HM type agent, perhaps because the behavior of this agent is somewhat contradictory. It agrees to and proposes considerate exchanges while fulfilling some, but not all, of its commitments.

To conclude, in repeated game settings, the MP agent, which conditioned its own personality based on its estimate of others, outperformed all of the other agents in the system. Some types of agents (LL, LM and MH) escaped identification in intermediate rounds, resulting in an increase in their scores. However, the general performance of the MP agent was not affected. It performed better than any other agent in each game, and increased its score from game to game during the final rounds of the experiment.

Figure 2: Difference in Score of MP agent between games 2 and 4



5.2 The Influence of Task Dependency

When agents need each others' resources, helpful CT agents are more likely to agree to an exchange of chips than unhelpful agents. However, they take a chance of getting exploited by players who are deceitful and renege on their commitments. Assuming that the MP agent can identify exploiters more quickly than other agents, we expected it to perform better than both helpful and unhelpful agents,

no matter what their task dependency role. Also, we hypothesized that when helpful SP and PD agents are task independent, they will be taken advantage by unhelpful players. Therefore we expected their performance to be lower than their unhelpful counterparts in this setting.

Agent Type	Task Dependent	task Independent
MP agent	160.32	231.2
Helpful PDs	137.23	221.16
Unhelpful PDs	150.58	253.54
Helpful SPs	113.26	231.93
Unhelpful SPs	93.69	241.2

Table 6: Agent Performance by Task Dependency

We analyzed the performance of all agents for each type of task dependency on the game boards, as presented in Table 6. Recall that we used two types of game boards in our experiment: One in which all agents were task dependent (allDep board) and one in which one agent was task independent and all other agents task dependent on it (oneSelf board).

Results confirmed part of our hypotheses: When they were task independent, unhelpful SP and PD agents were more successful than helpful SP and PD agents; when they were task dependent, helpful SPs were more successful than unhelpful SPs. A likely explanation is that when task independent, unhelpful agents did not negotiate with others since they did not need their help, while helpful agents that negotiated with others agents were taken advantage by unhelpful agents that did not fulfil their commitments. Therefore, unhelpful agents were more successful in this setting. When task dependent, helpful agents realized opportunities for exchange when fulfilling their commitments, which unhelpful agents did not, making helpful agents more successful.

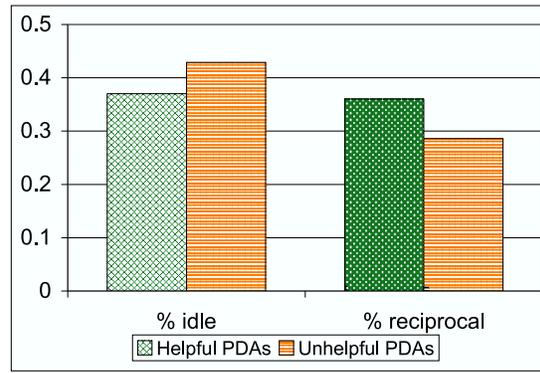
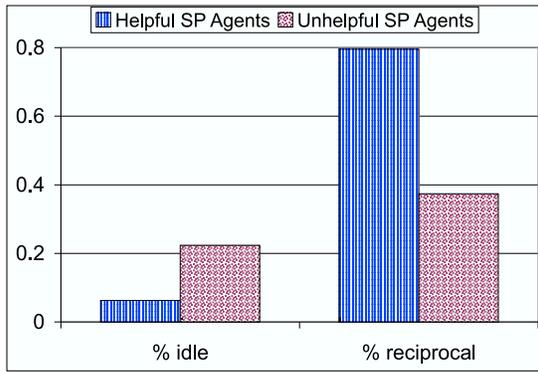
Regarding the performance of the MP agent, its score was significantly higher than both helpful and unhelpful agents when it was task dependent. The premise that the MP agent captured the other players was supported by Figure 3, which describes the fraction of times the MP agent engaged in each exchange type when it was task dependent. As shown in the figure, the MP agent engaged in idle behavior more often with unhelpful SP and PD agents than it did with helpful SP and PD agents. Interestingly, the figure also shows that the MP agent was idle much more often when dealing with unhelpful PDs (42%) than when dealing with unhelpful SPs (21%).

When it was task independent, the MP agent performed better than unhelpful SPs, but not better than unhelpful PDs. They were able to take advantage of its willingness to negotiate and renege on their commitments. It took longer for the MP agent to identify PD agents who were exploiters, and as a result they performed better in this setting. However, averaging over task dependencies, the MP agent performed much better than both PD and SP agents, supporting our findings in Section 5.1.

5.3 MP agents and Social Welfare

Our hypothesis was that any group of agents would increase its overall social welfare when playing with an MP agent. This is because MP agents engage in helpful exchanges that would not be realized when other agents are playing. To evaluate this claim, we ran a series of 2-player repeated games which included SP and PD type agents, but did not include MP agents, and compared it to the performance of each agent type after including an MP agent in the group. The results are described in Figure 4 and are statisti-

Figure 3: MP Exchange Behavior when task dependent. Left: with SP agents Right: with PD agents



cally significant with $p\text{-value} < 0.08$. The performance of helpful and unhelpful agents increased significantly when interacting with the MP agent. As expected, this increase was more profound for helpful SP and PD agents.

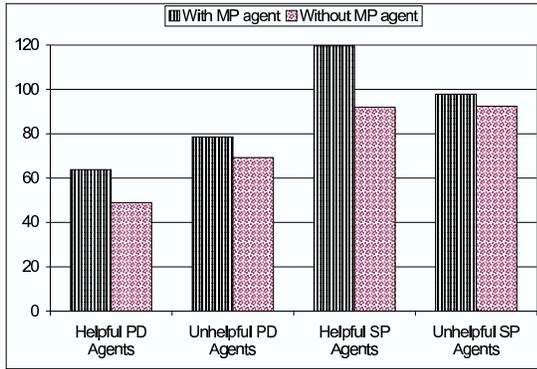


Figure 4: SP and PD agent Performance with/without MP agent

6. CONCLUSIONS AND FUTURE WORK

We have presented a model of negotiation which explicitly represents and reasons about agents' level of helpfulness. We evaluated this model in an open system for which there was no central design for the control of agents and in a domain characterized by uncertainty over agents' resources, as well as their helpfulness level. We have shown that agents that adopt a different cooperativeness and reliability measure, depending on who they interact with, could outperform agents that did not adapt. They could identify others personality, and adopt the right balance of behavior towards them more quickly and accurately than other agents. This enabled them to reciprocate helpful behavior while punishing deceitful behavior. Also, they improved the performance of all agents in the system, including unhelpful agents. We showed that when helpful agents are task independent, they engage in benevolent behavior and are taken advantage by unhelpful agents. However, when they are task independent, helpful players do better than unhelpful players, who do not realize the full potential of cooperating with others.

One future goal is to use personality models to describe team formation in agent systems. Are agents which exhibit similar personality traits more likely to form teams? How will this affect the performance of the system? Also, we intend to use this model to

build computers which interact with people. It has been shown that people's behavior in negotiation is affected by their preferences towards others outcomes, as well as their own [4]. We are interested to see if modeling people's measures of helpfulness can lead to better computer negotiators and explain people's behavioral tendencies outside of the game. We are also planning to tailor the model to deal with higher levels of uncertainty. For instance, by limiting the visibility of the board, it will be difficult for agents to assess their own chances of getting to the goal.

7. ACKNOWLEDGMENTS

The research reported in this paper has been supported in part by National Science Foundation grant no. 0222914. We thank Barbara Grosz for helpful discussions and suggestions.

8. ADDITIONAL AUTHORS

9. REFERENCES

- [1] H. Brandsttter and M. Knigstein. Personality influences on ultimatum bargaining decisions. In *European Journal of Personality*, volume 15, pages 53–70, 2000.
- [2] C. Castelfranchi, R. Falcone, and F. de Rosis. Deceiving in golem: How to strategically pilfer help. In *Autonomous Agent '98: Working notes of the Workshop on Deception, Fraud and Trust in Agent Societies*, 1998.
- [3] C. Castelfranchi, D. Fiorella, F. Rino, and P. Salvatore. A testbed for investigating personality-based multiagent cooperation. In *Proceedings of ESSLLI'97 Symposium on Logical Approaches to Agent Modelling and Design*, pages 23–35, Aix-en-Provence, France, August 1997.
- [4] Y. Gal, A. Pfeffer, F. Marzo, and B.J Grosz. Learning social preferences in games. In *Proc. 19th National Conference on Artificial Intelligence (AAAI)*, 2004.
- [5] B. Grosz, S. Kraus, S. Talman, and B. Stossel. The influence of social dependencies on decision-making. Initial investigations with a new game. In *Proc. 3rd International Joint Conference on Multi-agent systems (AAMAS)*, 2004.
- [6] L.M.J. Hogg and N.R. Jennings. Socially intelligent reasoning for autonomous agents. *IEEE Trans on Systems, Man and Cybernetics - Part A* 31(5), pages 381–399, 2001.
- [7] X. Zhang V. Lesser and T. Wagner. Integrative negotiation in complex organizational agent systems (research abstract). In *Proc. 1st International Joint Conference on Multi-agent systems (AAMAS)*, 2002.

- [8] R. Myerson. *Game Theory*. Harvard University Press, 1991.
- [9] S. Sen and P.S. Dutta. The evolution and stability of cooperative traits. In *Proc. 1st International Joint Conference on Multi-agent systems (AAMAS)*, pages 1114–1120, 2002.