

An Adaptive Agent for Negotiating with People in Different Cultures

YA'AKOV GAL

School of Engineering and Applied Sciences

Harvard University, U.S.A

and

SARIT KRAUS

Institute for Advanced Computer Studies

University of Maryland, U.S.A

and

MICHELE GELFAND

Department of Psychology

University of Maryland, U.S.A

and

HILAL KHASHAN

Department of Political Studies and Public Administration

American University at Beirut, Lebanon

and

ELIZABETH SALMON

Department of Psychology

University of Maryland, U.S.A

The rapid dissemination of technology such as the Internet across geographical and ethnic lines is opening up opportunities for computer agents to negotiate with people of diverse cultural and organizational affiliations. To negotiate proficiently with people in different cultures, agents need to be able to adapt to the way behavioral traits of other participants change over time. This paper describes a new agent for repeated bilateral negotiation that was designed to model and adapt its behavior to the individual traits exhibited by its negotiation partner. The agent's decision-making model combined a social utility function that represented the behavioral traits of the other participant, as well as a rule-based mechanism that used the utility function to make decisions in the negotiation process. The agent was deployed in a strategic setting in which both participants needed to complete their individual tasks by reaching agreements and exchanging resources, the number of negotiation rounds were not fixed in advance and agreements were not binding. The agent negotiated with human subjects in the U.S. and Lebanon in situations that varied the dependency relationships between participants at the onset of negotiation. There was no prior data available about the way people would respond to different negotiation strategies in these two countries. Results showed that the agent was able to adopt a different negotiation strategy to each country. Its average performance across both countries was equal to that of people. However, the agent outperformed people in the United States, because it learned to make offers that were likely to be accepted by people, while being more beneficial to the agent than to people. In contrast, the agent was outperformed by people in Lebanon, because it adopted a high reliability measure which allowed people to take advantage of it. These results provide insight for human-computer agent designers in the types of multi-cultural settings which we considered, showing that adaptation is a viable approach towards the design of computer agents to negotiate with people when there is no prior data of their behavior.

Categories and Subject Descriptors: I.2.11 [**Distributed Artificial Intelligence**]: Intelligent Agents

General Terms: Experimentation

Additional Key Words and Phrases: Human-Agent Decision-Making, Cultural Modeling

1. INTRODUCTION

Negotiation is a tool widely used by humans to resolve disputes in settings as diverse as business transactions, diplomacy and personal relationships. Developing technology is enabling computers to negotiate with people in a growing number of applications. Computer agents may negotiate on behalf of individual people or organizations (e.g., bidders in online auctions [Kamar et al. 2008; Rajarshi et al. 2001]); they can act as training tools for people to practice and evaluate different negotiation strategies prior to embarking on negotiation in the real world (e.g., agents for negotiating a simulated diplomatic crisis [Lin et al. 2009]), or work autonomously to reach agreements for which they are responsible (e.g., computer games, systems for natural disaster relief [Schurr et al. 2006; Murphy 2004]).

The prevalence of the Internet and the mobility and accessibility of technology such as netbooks and smart phones has transcended ethnic and geographic boundaries. As a result, computer agents are increasingly interacting with people of diverse nationalities and cultures. For example, a computer agent may train diplomats to represent their respective countries in international negotiations over climate control policies. These negotiations may consist of multi-attribute issues such as commitments to regulate the emissions of greenhouse gasses, contributions to a relief fund to assist developing nations in their efforts, and enforcement mechanisms to ensure compliance with the agreement. The human participants of this negotiation process would represent multiple nationalities and might differ in their ethnic and cultural affiliations.¹

Work in artificial intelligence has recently began to address the need for computer agents to negotiate proficiently with people [Lin and Kraus 2010]. However, there are particular challenges facing agent design when negotiating in environments that include people from different cultures. Most importantly, there is scarce data available about how people would respond to computational strategies in different cultures. For example, digital technologies were the first telecommunication infrastructure to be introduced in some developing African nations [Businessweek 2007]. These countries bypassed the age of analogue telecommunication, and there was great uncertainty about how people would relate to cell phones and other digital technology. In addition, it is difficult to derive analytically the behavioral and social factors that may be related to cultural differences. Thus it is not possible to adopt equilibrium-based paradigms such as those used in behavioral economics [Camerer 2003; Charness and Rabin 2002] for guiding agents' behaviors in multi-cultural settings in which there is no prior data about the way people may react to different

¹An agent for training to manage diplomatic crises has recently been designed by Kraus et al. [2008], but it was not evaluated by interacting with people of diverse cultural backgrounds.

computational strategies.

This paper presents a new agent design for meeting these challenges by modeling and adapting to behavioral traits of its negotiation partners. People’s cultural affiliation is a key driving force of the way they make decisions in social contexts [Gelfand and Brett 2004; Gelfand and Christakopoulou 1999]. In our example, a representative of a country that values altruism and generosity may agree to invest more in a relief fund for developing nations than a representative from another country with similar economic standing that values self-preservation or competition. The agent uses a decision-theoretic approach for selecting the best action for the agent given its model of others’ behavioral traits, combined with heuristics for narrowing the search space of possible strategies. It is designed to negotiate in strategic bilateral negotiation settings which require participants to reach agreement about the procurement of resources in order to reach their goals. Such settings characterize many types of negotiation scenarios in the real-world in which participants may renege on past agreements during the negotiation process. For example, one of the countries in the above diplomatic simulation may fail to meet its agreed target of emission reduction in favour of sustaining its economic growth.

This agent, referred to as the Personality Utility Rule Based (PURB) agent, used a decision-making paradigm that was composed of the following interleaving components:

- a model of the behavioral traits of its negotiation partners, consisting of the extent to which they are helpful and reliable over time.
- a utility function that is a combination of (1) the behavioral traits of the negotiation partners of the agent, (2) the future ramification of potential agreements (the effect that potential actions in negotiation may have on the agent in future rounds), and (3) the extent to which the negotiation partner is reasoning about the behavioral traits of the agent itself.
- heuristics that use domain knowledge to guide the agent in its search of the best possible strategy. These heuristics help to reduce the set of possible offers to consider when making proposals and the extent to which the agent should make agreements and whether to fulfil them. These heuristics depended on the agent’s model of its negotiation partners as well as the dependency relationships that held between participants in the game.

The PURB agent updated its model of people’s reliability and helpfulness over time by incorporating observations consisting of their actions in negotiation. The proposed agent is evaluated by its performance when interacting with people from different cultures. Our definition of “cultural differences” in the context of this paper refers to variance in behavior that arises between participants that negotiate with our computer agent in different countries.

The evaluation of the agent was extensive, and included two-hundred students from Lebanon and the United States. Subjects interacted with the PURB agent using a computer board game in which players needed to procure and exchange resources in order to fulfil their goals. Three different types of games were used that varied the dependency relationships between participants. The number of negotiation rounds was not fixed in advance, agreements were not binding, and

there was no prior data consisting of people's behavior to guide the agent's behavior in the game.

Our results show empirically that the PURB agent was able to adapt to the different negotiation behaviors that were exhibited by people in the two different countries. In terms of average performance, the PURB agent was able to negotiate as well as people across both countries. In the U.S. alone, the PURB agent outperformed people, while in Lebanon alone, the PURB agent was outperformed by people. More, we found distinct differences in the behavioral traits that were exhibited by people across the two countries. In the U.S., people exhibited low reliability measures. The PURB agent adapted to this behavior by accepting and making proposals that were more beneficial to the PURB agent than to people, and exhibiting slightly lower reliability measures than people. In Lebanon, people exhibited significantly higher reliability measures. The PURB agent adapted to this behavior by making proposals that were more beneficial to people than to the PURB agent, and exhibited slightly higher reliability measures than people. The reason for this was that the adaptation mechanism used by the PURB agent was not fine-grained, and overestimated the reliability measures of people in Lebanon to be higher than there. People in Lebanon were able to exploit the PURB agent's reliability, which explains their better performance.

The study reported in this paper is the first to suggest and evaluate an agent that is able to adapt to the negotiation strategies of people in different cultures. Our goal was not to present a comparison between the way people negotiate with the agent across cultures, but to establish that adaptation is a viable approach towards agent-design for human-computer negotiation when there is no prior data of people's behavior. The contributions of this paper include (1) the presentation of a novel agent design that explicitly reasons about the behavioral traits of others in settings of repeated negotiation and non-binding agreements; (2) the use of a novel empirical paradigm for studying people's negotiation behavior, using a computer agent as a base-line with which to compare the way psychological factors affect people's negotiation behavior across different cultures; (3) A first experimental studies of negotiation in Lebanon, and the first comparison between the way people respond to computational strategies in the U.S. and the Middle East.

2. RELATED WORK

The field of human-computer negotiation is a budding field within artificial intelligence research. We refer the reader to Lin and Kraus [2010] for a comprehensive review. The approaches used by these works can be divided into three main themes: game theoretic, opponent modeling and logic. Game theoretic approaches to human-computer decision-making assume that agents and people optimize utility functions that may depend on a variety of attributes. These may include their performance in the negotiation, social preferences such as altruism or competitiveness, and features that relate to the negotiation protocol, such as time. The game is then analyzed to find the optimal negotiation strategies for the agent, assuming that people are also optimal with respect to their own utility function. Examples of this approach include Kraus et al. [2008] who modeled human bilateral negotiations in a simulated diplomatic crisis characterized by time constraints and deadlines in

settings of complete information. They adapted equilibrium strategies to people's behavior using simple heuristics, such as considering certain non-optimal actions. In an auction setting, Rajarshi et al. [2001] has shown that computer agents using market equilibrium can outperform human bidders. Jonker et al. [2007] designed computer strategies that involve the use of concession strategies to avoid impasses in the negotiation.

Opponent modeling approaches do not assume that people optimize their respective utility functions. These approaches use learning and adaptation techniques to discover people's negotiation strategies, whether directly, or by defining utility functions that allow humans to deviate from optimality in a controlled fashion. Representative works for human-computer negotiation include Byde et al. [2003], who constructed agents that bargain with people in a market setting by modeling the likelihood of acceptance of a deal. Oshrat et al. [2009] used density estimation techniques to estimate people's acceptance of offers in a repeated multi-attribute negotiation scenario, and show that such an agent was able to outperform people. Learning techniques have also been applied to model gender differences [Katz and Kraus 2006] and the belief hierarchies that people use when they make decisions in one-shot interaction scenarios [Gal and Pfeffer 2007; Ficici and Pfeffer 2008]. Zuckerman et al. [2007] used machine learning to allow agents to predict which strategy will be chosen by people in settings which demand coordination among several possible strategies.

To date, there are scant computational models of human negotiation behavior that reason about cultural differences. More, all work on human-computer negotiation assumes that agreements are binding, and have relied on prior data of people's negotiation behavior. A notable exception is work by Kraus and Lehmann [1995] that proposed a logic for an agent that negotiates with multiple participants and may renege on agreements, but this work was restricted to a specific domain, that of the game of Diplomacy. This research extends the state-of-the-art of human-computer negotiation in its focus on human-computer negotiation in situations where agreements are not binding, and in its extensive empirical study that spanned over two-hundred subjects in two countries.

Within the behavioral economics literature, Henrich [2000] and Roth et al. [1991] have identified a stark difference in the way people play simple negotiation games such as the ultimatum game across cultures and attributed this to people's varying perception of fairness in these communities. (See Section 5.3 for a comparison between our results and these studies). More recent work by Bohnet et al. [2008] has investigated trust games in six countries. This study found that people in the Middle East are less willing to enter a lottery when its source of the risk is another person rather than nature, as compared to people in western countries.

There is a large body of work in the psychological and social sciences that investigate behavior among human decision-makers from different cultures [De Dreu and Van Lange 1995; Kashima et al. 2005; Gelfand et al. 2002] (inter alia), ranging back to the classic studies of Hofstede [1980] on the dimensions characterizing cross-cultural differences among people. Most of the psychological studies on negotiation across cultures have been comparative, in the sense that it points to differences in people's negotiation behavior that are motivated by cultural factors, such as

individualism and collectivism [Gelfand and Realo 1999]. Other works have investigated cross-cultural negotiation in settings in which parties communicate over time. Weingart and Olekalns [2004] have examined communication patterns among negotiators and show that the extent to which they reveal salient information to each other over time depends on the negotiation context. Adair and Brett [2004] have shown that U.S. negotiators are more efficient, and achieve higher joint gains when they share information directly with other parties directly, and that Japanese negotiators are more efficient with information is shared indirectly.

Our work builds on these important studies in the social sciences in that the computational model we propose depends on people’s measures of trust and generosity, which have been shown by these works to be culturally determined. While most work in this realm has compared negotiation behavior in mono-cultural contexts (that is, people negotiate among members of their own cultures), more recent works have begun to study more diverse settings in which negotiators from different cultures need to reach agreements [Imai and Gelfand 2010]. Lastly, we note that almost all of the cross-cultural studies on negotiation have been in the US and Europe and Asia; this is of the first experimental studies of negotiation in Lebanon and the first study comparing between the way people react to computational strategies in the United States and in the Middle East. A version of this agent has recently been used to negotiate with people in Israel [Gal et al. 2010].

3. IMPLEMENTATION USING COLORED TRAILS

Our study was based on the Colored Trails (CT) game, a test-bed proposed by Grosz et al. [2004] for investigating decision-making in groups comprising people and computer agents.²

3.1 The Game Description

The CT configuration we used consisted of a game played on a 7x5 board of colored squares with a set of chips. One square on the board was designated as the goal square. Each player’s icon was initially located in one of the non-goal positions, eight steps away from the goal square. To move to an adjacent square a player needed to surrender a chip in the color of that square. Each player was issued 24 colored chips at the onset of the game.

Snapshots of one of the CT games used in our study are shown in Figure 1. Figure 1a shows the board game, in which there are two players, “me” and “O”. The board game is shown from the point of view of the “me” player. The relevant path from the point of view of the “me” player is outlined. This figure also shows the chips that both players possess at the onset of the game. Both “me” and “O” players are missing three chips to get to the goal. The “me” player is lacking three yellow chips; while the “O” player is lacking three grey chips. In addition, each player has the chips the other player needs in order to get to the goal. For example, the “me” player has ten grey chips. Figure 1b shows an example of a proposal made by the “me” player to give two grey chips to the “O” player in return for two of its

²Colored Trails is Free Software and is available for download at <http://www.eecs.harvard.edu/ai/ct>

yellow chips. The Message History Panel (Figure 1c), includes the response of the “sun” player, which accepted the offer.

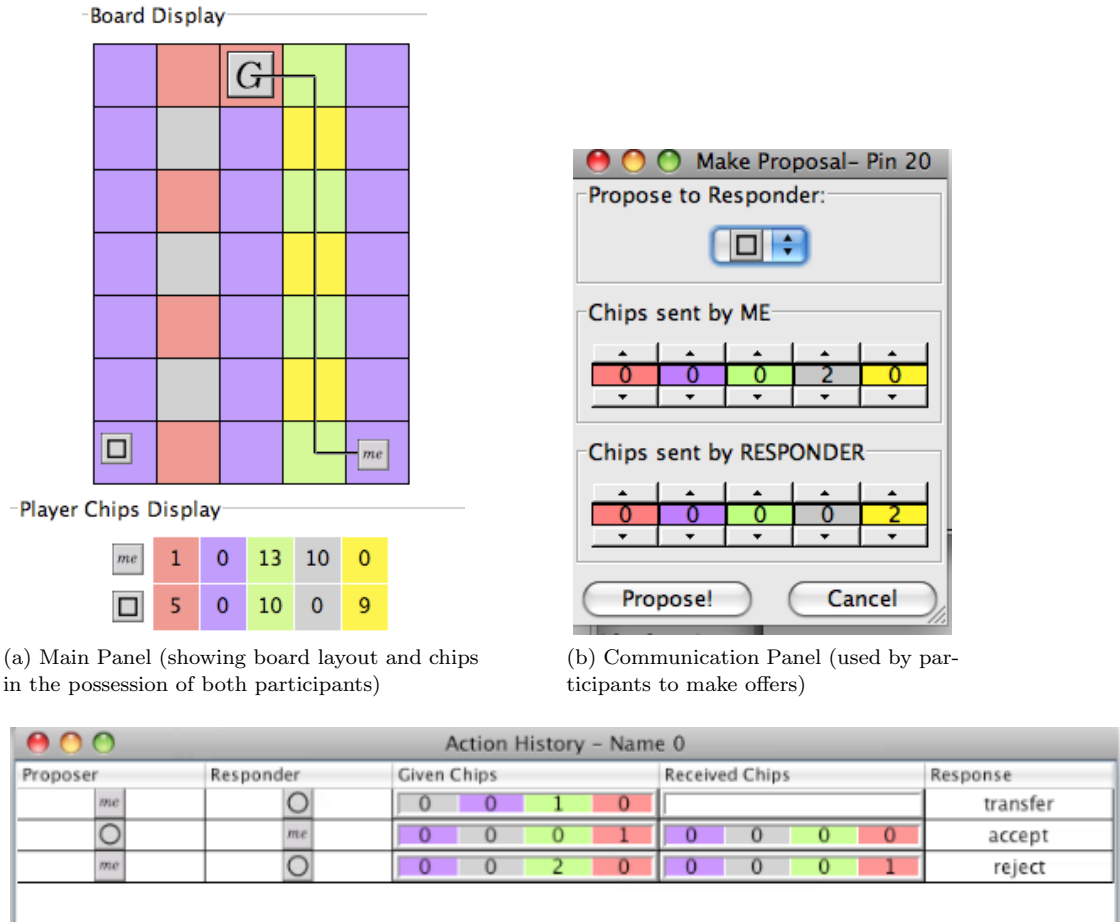


Fig. 1: Snapshots of CT GUI

At the onset of the game, one of the players was given the role of proposer, while the other was given the role of responder. The interaction proceeded in a recurring sequence of phases. In the *communication* the proposer could make an offer to the responder, who could accept or reject the offer. In the *transfer phase*, both players could choose chips to transfer to each other. The transfer action was done simultaneously, such that neither player could see what the other player transferred until the end of the phase. In particular, players were not required to fulfil their commitments to an agreement reached in the communication phase. A player could choose to transfer more chips than it agreed to, or any subset of the chips it agreed to, including transferring no chips at all. In the *movement phase*, players could manually move their icons on the board across one square by surrendering a chip in

the color of that square. At the end of the movement phase, a new communication phase began. The players alternated their roles, such that the previous proposer was designated as a responder, and vice versa. These phases repeated until the game ended, which occurred when one of the following conditions held: (1) at least one of the participants reached the goal square; or, (2) at least one of the participants remained dormant and did not move for three movement phases. Note that players had full view of the board and each others' chips, and thus they had complete knowledge of the game situation at all times during the negotiation process.

3.2 The Scoring Function

When the game ended, both participants were automatically moved as close as possible to the goal square, and their score was computed as follows:

- 100 points bonus for getting to the goal square,
- 5 points bonus for any chip left in a player's possession,
- 10 points penalty for each square left in the path from a player's final possession and the goal square.

These parameters were chosen so that getting to the goal was by far the most important component, but if a player could not get to the goal it was preferable to get as close to the goal as possible. Note that the score in CT depended not only on whether a player could reach the goal square, but also on the number of chips the player had left over in its possession at the end of the game.

Although performance for each player is measured by applying the scoring function at the end of the game, the scoring function can also be applied at intermediate stages of the game. This allows a player to hypothesize about the benefit it would incur for carrying out actions. For example, a player can reason about whether a proposed exchange will enable it to reach the goal by computing its score in the game given that it agreed to and realized the exchange.

3.3 Dependency Relationships

We used two different types of boards in the study. In both of these boards, there was a single distinct path from each participant's initial location to its goal square. One of the board types exhibited a symmetric dependency relationship between players: Neither player could reach the goal given its initial chip allocation, and there existed at least one exchange such that both players could reach the goal. We referred to players in this game as task co-dependent. The other board type exhibited an asymmetric task dependency relationship between players: One of the players, referred to *task independent*, possessed the chips it needed to reach the goal, while the other player, referred to as *task dependent*, required chips from the task-independent player to get to the goal. An example of the co-dependent board is shown in Figure 1a. In this game both "me" and "O" players were missing three chips to get to the goal: The "me" player was missing three yellow chips whereas the "O" player was missing three grey chips. The relevant path from the point of view of the "me" player is outlined. We note that in all the dependency conditions we used in the study, there were no cases in which the same colors were needed for both players to get to the goal, and thus participants did not directly compete over each other's resources (i.e., the game was not zero sum.)

3.4 Task Analogy

One of the advantages of using CT for cross cultural studies is that it provides a realistic analog to task settings, highlighting the interaction among goals, tasks required to achieve these goals, and resources needed for completing tasks. In CT, chips correspond to players' capabilities and skills required to fulfill tasks. Different squares on the board represent different types of tasks. A player's possession of a chip of a certain color corresponds to having the skill available for use at a time. Not all players possess chips in all colors, much as different players vary in their capabilities. Traversing a path through the board corresponds to performing a complex task whose constituents are the individual tasks represented by the colors of each square.

The abstraction provided by CT is especially appropriate for investigating negotiation in different cultures because it avoids culturally-loaded contexts and settings that may confound people's behavior (e.g., negotiating over items that may have religious connotations). CT is thus the right kind of test-bed to use to study negotiation that occurs between people of different cultures, in which negotiation processes are conducted within task contexts, and involve the exchange of resources (for example, within diplomatic negotiations for trade agreements or peace treaties).

To illustrate the task analogy, we present an example of the way CT corresponds to the diplomatic negotiation domain described in Section 1. Players in CT correspond to diplomats representing their respective countries. A goal square might represent a goal to accomplish, such as a protocol on climate change or signing a peace accord. Paths on the board represent the completion of tasks, such as completing phases in a program for greenhouse gas reduction. Chips represent resources, such as oil reserves, natural gas and technological capabilities. Negotiation over these resources by the diplomats is necessary to reach agreement. For example, one country may offer another the benefit of solar technology in return for rights to use part of the generated energy. Players' scores in this game may be set to depend solely on their individual performance as negotiators or may also include the score of other diplomats to capture coalitions among countries.

4. THE PERSONALITY UTILITY RULE BASED AGENT

In this section we present an agent design for negotiating with people in the CT scenario described in Section 3. In order to succeed, the agent needs to make the right trade-off between adopting cooperative negotiation strategies (e.g., making generous proposals, fulfilling agreements) and selfish strategies (e.g. making selfish proposals, fulfilling agreements in part, or not fulfilling them at all). To this end the agent needs to reason about the cumulative effect that its potential actions may have on its performance given the perceived strategy of its negotiation partner. For example, if the agent reneges on an agreement it reached with the other participant to gain more chips (i.e, it does not transfer all of the chips in the agreement), the other participant may renege on a future agreement, and this may harm the agent's performance compared to the case in which it had decided to transfer.³ Thus it is

³This is true even for the case in which the agent can get to the goal independently, because the CT scoring function depends on the number of chips in the agent's possession that are left at the

necessary to reason about the ramifications of potential negotiation actions in the game, and these in turn depend on the negotiation strategy of the other participant.

The agent that we constructed for this setting, called the Personality Based (PURB) agent, modeled other participants in terms of two behavioral traits: helpfulness, and reliability.⁴ The *helpfulness* measure of a participant represented the extent to which the participant shared resources with its negotiation partner through initiating and agreeing to proposals. The *reliability* measure of a participant described the degree to which the participant kept commitments to its negotiation partner.

Without loss of generality, we will use “agent i ” to refer to the PURB agent, and “agent j ” to refer to its negotiation partner. Given a proposal $O = (O_i, O_j)$ made by j at a round in the game, we refer to the set of chips offered by j to i as O_j , and the set of chips that j requested from i as O_i . We use O_j^* and O_i^* to denote the chips that were *actually* transferred by j and i respectively during the transfer phase that followed the agreement. Using these terms, we can define the helpfulness of agent j , denoted h , as the percentage of proposals in the game in which $|O_j| > |O_i|$ (that is, when the negotiation partner of PURB offered more chips to PURB than it requested for itself.) We can also define the reliability of agent j , denoted r , as $\frac{|O_j^*|}{|O_j|}$ (the ratio of the number of chips actually transferred to the PURB agent in an agreement by its negotiation partner and the number of chips promised). Together, we refer to the pair (h, r) as the *cooperativeness* measure of j .

4.1 Social Utility Function

The PURB agent used a social utility function u_i to make decisions which was a weighted combination of the following features:

- (1) The expected future score for agent i . This score was computed using a heuristic function that estimated the benefit to the agent from a potential exchange. It depended on the probability that i will get to the goal at some state s given that proposal O is fulfilled. We denote this probability as $p_i(G | s, O)$, and define the expected future score to i as

$$(p_i(G | s, O) \cdot 100) + ((1 - p_i(G | s, O)) \cdot 10 \cdot d) + (c \cdot 5)$$

where 100 is the number of bonus points to get to the goal according to the CT scoring function; d is the Manhattan distance of i from its final position on the board and the goal square, given that the agreement was fulfilled by j ; 10 is the number of penalty points for each square in the distance from the final position of PURB and the goal square; c is the number of chips left in the player’s possession after it advances to the goal using the shortest possible path, and 5 is the number of points awarded to the player for each chips left in its possession at the end of the game. The probability $p_i(G | s, O)$ to get to the

end of the game, and the independent player can still improve its performance if it receives chips from the other.

⁴PURB is based in part on an agent-design proposed by Talman et al. [2005] for repeated negotiation among artificial agents that adapted to others’ varying level of cooperativeness. We use the term “participants” to refer to both people and computer agents.

goal at state s given proposal O was estimated as the ratio between the number of chips that j delivered to i , and the number of chips that i was missing to get to the goal at state s given that O was fulfilled.

- (2) The expected future score for agent j (computed in the same way as for i).
- (3) The cooperativeness of agent i . This term is representing using helpfulness and reliability, as defined above.
- (4) The *perceived* cooperativeness of agent j . This feature represented i 's model of j 's beliefs about the reliability and helpfulness of i . This term is represented using helpfulness and reliability as defined above.

The weightings for the feature of u were set by hand, and depended on the dependency relationships between participants as well as their cooperativeness measures. Generally, as agent j increased its cooperativeness measures, the weighting in i 's social utility function that was associated with j 's score was increased. This was to incentivize i to be more generous when j was cooperative. Also, when agent j was task independent, the weighting in i 's social utility function that was associated with its score in the game decreased. The purpose for this was to provide an incentive to i to make less generous offers to an independent participant j . Each time an agreement was reached and transfers were made in the game, i updated the helpfulness and reliability measures of both agents (these values were aggregated over time using a discounting rate of 0.1). Using this social utility allows the PURB agent to vary its strategy based on its estimate of the other participant's cooperativeness measure. For example, if the reliability of the other participant was high, this would increase the social utility of actions that favour the other participant.

4.2 Rules of Behavior

The second component of PURB's decision-making paradigm was a set of rules that narrowed the search space of possible actions to be considered by the agent's utility function. These rules depended on aspects relating to the state of the game (e.g., the number of chips each agent had, whether a participant can independently reach the goal). At each step of the game, the agent used its social utility function to choose the best action out of the set of possible actions that were constrained by the rules. The rules were designed such that the PURB agent begins by acting reliably, and adapts over time to the individual measure of cooperativeness that is exhibited by its negotiation partner. To enable to specify a finite set of rules for different measures of reliability and helpfulness, the possible values that these traits can take were divided into three equal intervals representing low, medium or high measures. For example, low reliability measures ranged from 0 to $\frac{1}{3}$. We then defined the cooperativeness of an agent to depend on the extent to which it was reliable and helpful.

Specifically, we defined the cooperativeness of an agent to be *high* when it exhibited high helpfulness and high reliability measures, or high helpfulness and medium helpfulness measures; the cooperativeness measure of an agent was *medium* when it exhibited medium reliability and medium helpfulness measures, or medium reliability and high helpfulness measures; the cooperativeness measure of an agent was *low* when it exhibited low reliability measures (regardless of its helpfulness measure) or

medium reliability measures and low helpfulness measure. These values were tuned by hand.

We now list the set of rules used by the PURB agent in combination with its social utility function:

- a) *Making Proposals* The PURB agent generated a subset of possible offers that were evaluated by its social utility function. It non-deterministically chose any proposal out of the subset that provided a maximal benefit (within an epsilon interval) according to its social utility function. Before outlining the rules by which the set of possible proposals were generated, we will introduce the following notation: We say an agent i is “stronger” than agent j if i is able to reach the goal independently of j , or if it requires less chips to reach the goal than j . Let $O_{i=j}$ represent the set of proposals in which agent i asks for as many chips as it receives; $O_{i>j}$ represents the set of proposals in which i asks for more chips than it receives; $O_{j>i}$ represents the set of proposals in which i asks for less chips than it receives.

Offers were generated by PURB in a way that considered which participant was stronger than the other. When participants were co-dependent, the set of possible offers i considered included those offers that favoured the stronger agent. If i was stronger than j , then the set $O_{i>j}$ was considered (i.e., i requested from j more chips than i proposed to j) (and conversely for the case in which j was stronger than i). In addition, the set $O_{i=j}$ was also generated and considered (i asks for as many chips as it sends).

In the other dependency roles, the offers that were generated depended on the cooperativeness measure of j :

- 1) When the cooperativeness of j was high or medium, then if i was stronger than j , then the set of possible offers that i considered included $O_{i>j}$. This is because that when the reliability of j was high, there was a higher likelihood that j would keep its commitments, and thus the set of possible exchanges for i included exchanges that were highly favorable to i . However, if j was stronger than i , then offers were chosen from the set $O_{j>i}$. This was because that i wished to minimize the chances that j would reject its offers given that j did not need i to get to the goal.
- 2) When the cooperativeness of j was low, then offers were chosen from the set $O_{i>j}$, regardless of which agent was stronger. This was because i did not expect j to fulfil its agreements, and thus it proposed offers that were less beneficial to j .

Lastly, the degree to which i offered proposals that made j task-independent (i.e., they allowed j to get to the goal) but i would remain task-dependent, occurred only when i was highly certain that j 's reliability was high. This occurred when there was at least one past example of reliable behavior from the negotiation partner.

- b) *Accepting Proposals* As a responder, the PURB agent accepted an offer if it was more advantageous to it than the offer it would make as a proposer in the same game state, or if accepting the offer was necessary to prevent the game from terminating. To state this formally, let $u_i(O, \text{accept} \mid s)$ denote the social utility for i from an offer O made by j at state s . Let O' denote the offer that

agent i would make at state s according to the rules in (a). Agent i accepted an offer O if $u_i(O, \text{accept} | s) \geq u_i(O', \text{accept} | s)$. In addition, i would accept any proposal that prevented the game from ending, which occurs when the following conditions hold: (1) the chips in the possession of agent i do not allow it to move on the board at state s ; (2) the offer O_j allows agent i to move; and (3) if i rejects an offer, the limit for dormant turns will be reached and the game would end.

Lastly, the degree to which i accepted proposals from j that made j task-independent but i would remain task-dependent, occurred only when i was highly certain that j 's reliability was high. This occurred when there was at least one past example of reliable behavior from the negotiation partner.

- c) *Transferring Chips* These rules specify the extent to which the PURB agent fulfilled its agreements in the game. The agent was programmed to be fully reliable (it sent all of the promised chips) after the first agreement made in the game in order to promote a positive reciprocal relationship between PURB and its negotiation partner.⁵ Otherwise, the extent to which PURB transferred chips directly depended on its model of the cooperativeness of its partner:

If the reliability of j was *high*, it was likely that j would fulfil its agreements. Therefore i sent all of its promised chips.

If the reliability of j was *low*, it was likely that j would not fulfil its agreement. Therefore i did not send any of its promised chips.

If the reliability of j was *medium*, then the extent to which i was reliable depended on the dependency relationships in the game:

- 1) If j was task dependent, and the agreement resulted in j becoming task independent, then i sent the largest set of chips such that j remained task dependent.
- 2) If the exchange resulted in i becoming task independent, and j remaining task dependent, then i sent all of the promised chips, or two thirds of its promised chips, depending on its confidence level of j 's reliability measure being medium. This confidence was high when there was at least one past example of reliable behavior from the negotiation partner.
- 3) If both agents were task dependent, and the agreement resulted in *both* agents becoming task independent, then i sent half of the promised chips with a probability of 2/3 and all of the chips with a probability of 1/3, because it was not certain that j would fulfil the agreement.

Combining the PURB agent's social utility function with these rules allows it to adapt its negotiation behavior to that of the other participant.

5. EXPERIMENTAL DESIGN AND RESULTS

The following section describes an experiment for evaluating the ability of the PURB agent to negotiate proficiently with people. We hypothesized that an agent that exhibited high reliability measures would be reciprocated by people and would

⁵We provide additional motivation for this rule in Section 5.

be able to perform well in this setting.⁶ The experiments were conducted in two countries, Lebanon and the U.S. The total number of participants in the study was 209. Basic demographic details of the subjects were as follows: In both countries all subjects were students enrolled in a university or college degree program. The mean age of subjects in the U.S. was 23 (with standard deviation 5.8), whereas the mean age of subjects in Lebanon was 20 (with standard deviation 1.35). In the U.S., 44% of the subject population was male, whereas in Lebanon, 61% of the population was male. The main ethnic groups of subjects in the U.S. consisted of 45% Caucasian, 21% African or African American, 20% Asian or Asian American. The main religious affiliations of subjects in Lebanon were 68% Christian and 31% Muslim.

Each participant was given an identical 30 minute tutorial on CT. This tutorial consisted of a written description of the CT game, as well as a short movie that explained the rules of the game using a different board than those used in the study (See Appendix). Participants were seated in front of terminals for the duration of the study, and could not speak to each other or see their terminals. All participants played a single game with the PURB agent, but were told they would be playing another person. Authorization for this deception was granted by the ethics review board of the institutions that participated in the study. Subjects were given an extensive debriefing at the end of the study which revealed this fact and explained the study (see Appendix). All participants were paid a constant sum (\$20 in the U.S., \$15 in Lebanon) that did not depend on their performance in the game.⁷

Each subject played a single CT game and was randomly assigned one of the following dependency roles: A task co-dependent participant that was paired with another task co-dependent computer agent (played on the symmetric board shown in Figure 1a); a task independent participant that was paired with another task dependent participant or, a task dependent participant that was paired with another task independent participant (both of these conditions were played on a board that represents asymmetric dependencies.) The games corresponding to both of these conditions were played on an asymmetric board in which the task independent participant could get to the goal, while the task dependent participant could not get to the goal.

Our analysis compared between the performance of the PURB agent and people, and distinguished between the behavioral traits exhibited by people in both countries. In Lebanon, we collected 30 game instances for each dependency role, In the U.S., we collected 38 game instances when players were task co-dependent, 32 game instances when people were task dependent, and 38 instances when people were task independent. In all of the CT games we ran, people were designated as first proposers. The reason for this was to be able to compare between the types of offers people make in negotiation in the U.S. and in Lebanon before they react

⁶This hypothesis is supported in a study that shows that CT was able was to bring about cooperative behavior in people [Gal et al. 2007].

⁷Our purpose was to evaluate the ability of PURB to interact with people in different cultures across the same conditions in both countries. The question of whether people's behavior was monetarily induced was not significant to this study. We did not pay people in a way that depended on their performance.

	Co-Dep.	Task Ind.	Task Dep.	Average
PURB agent (Combined)	161	200	137	165
People (Combined)	161	195	128	162
PURB agent (Lebanon)	160 ($n = 30$)	198 ($n = 30$)	120 ($n = 30$)	159
People (Lebanon)	180 ($n = 30$)	206 ($n = 30$)	154 ($n = 30$)	180
PURB agent (U.S.)	161 ($n = 38$)	202 ($n = 32$)	151 ($n = 35$)	170
People (U.S.)	146 ($n = 38$)	185 ($n = 35$)	104 ($n = 32$)	146

Table I: Performance for different Dependency Conditions

to the strategy used by PURB, as we report in Section 5.3.

All of the reported results were determined to be statistically significant for $p < 0.05$ using parametric statistical tests. In addition, all of the results we report are significant when controlling for subjects' gender.

5.1 Analysis of Performance

We measured performance in a game according to two factors: whether a participant was able to reach the goal, and the score obtained in the game by the participant according to the CT scoring function described in Section 3.2. Table I shows the average score obtained by the PURB agent and people for each country. The "average" column lists the average performance of all the games that were collected. This table also lists the number of game instances that were collected in each dependency role in parentheses.

As shown by the table, there was no significant difference between the average performance of the PURB agent and people when combining the data from games played in both countries (165 points for PURB agent versus 162 points for people).⁸ However, there was a striking difference in performance when separating the analysis by countries. In Lebanon the average score obtained by people was higher than that of the PURB agent (180 points per game versus 159 points per game). In contrast, in the U.S., the average score obtained by the PURB agent was higher than that of people (170 points per game versus 146 points per game). This result was consistent for task co-dependent, task independent and task dependent roles. As shown by Table I, in the U.S., the PURB agent achieved a higher score than people in all three dependency roles. In Lebanon, people achieved a higher score than the PURB agent in all three dependency roles.

This pattern is also apparent when comparing the likelihood of reaching the goal square (not shown in the table). When analyzing the data combining both the U.S. and Lebanon, we found that there was no significant difference between the likelihood to get to the goal between people and the PURB agent. However, people in Lebanon reached the goal square significantly more often than the PURB agent, while in the U.S. the PURB agent reached the goal square significantly more often than people.

When participants were task co-dependent and task independent there was no significant difference in the likelihood to get to the goal square in both the U.S. and

⁸We report the average results to provide a general performance rating for the agent across all conditions. We do not imply that the agent will actually achieve this same result for particular board games.

	Co-Dep.	Task Ind.	Task Dep.	Average
PURB agent (Lebanon)	0.96	0.99	0.99	0.98
People (Lebanon)	0.96	0.94	0.87	0.92
PURB agent (U.S.)	0.59	0.59	0.72	0.62
People (U.S.)	0.64	0.78	0.51	0.65

Table II: Reliability Measures of Participants in Lebanon and the U.S.

in Lebanon. Also, in both countries, participants who were task independent always got to the goal square.⁹ However, the likelihood to get to the goal was different for both countries in the task dependent condition. In the U.S., the PURB agent got to the goal more often than people, while in Lebanon, the converse was true.

5.2 Analysis of Behavior: Reliability

We define the reliability of an agent as the extent to which it fulfilled its commitments in a game. Formally, for any two participants i and j , let C_i^n denote the set of chips in possession of i at round n in the game. Let $O = (O_i, O_j)$ denote a proposal, where $O_i \subseteq C_i$ was the set of chips that i agreed to send to j , and let $O_i^* \subseteq C_i$ be the set of chips actually sent by i following the agreement. (And similarly define C_j, O_j , and O_j^* for player j). Let $r_i(\{C_i \cup O_j\})$ denote the score to player i in the case that j sent all of its promised chips O_j , and i did not send any of its chips. This is computed using the scoring function for the CT game that is described in Section 3. We refer to this value as the score that was *promised by* j in proposal O . The factor $r_i(\{C_i \cup P_j^*\})$ denotes the score to player i given the chips P_j^* that j actually delivered by j . We refer to this as the *actual* score to i given the chips that j transferred. The *reliability* of j given proposal p is the ratio between the promised score to i from the chips O_j in proposal p and the score to i given the chips that j transferred. This is defined as $\frac{r_i(\{C_i \cup O_j^*\})}{r_i(\{C_i \cup O_j\})}$. Note that this measure of reliability represents a continuous, more fine grained description than the discrete reliability measure used by the rule-based component of the PURB agent (Section 4).

Table II lists the reliability measures of participants. As shown by the table, subjects in Lebanon were highly reliable. On average, they delivered 92% of the benefit they had promised to the PURB agent. In contrast, the subjects in the U.S. exhibited strikingly lower reliability measures than subjects in Lebanon. On average, these subjects delivered only 65% of the benefit they had promised the PURB agent.

The PURB agent responded to this behavior by adapting a measure of reliability towards participants in each country that was proportional to the measure of reliability exhibited by people in that country. We found a positive significant correlation of 0.42 between the reliability of the PURB agent and the reliability of people. Specifically, in Lebanon the PURB agent adapted a reliability measure of 98%, while in the U.S. the PURB agent adapted a reliability measure of 62%. Interestingly, the average reliability measure of the PURB agent in Lebanon was higher than the average reliability measure of people, but there was no statistically

⁹It was technically possible for task independent players to give chips away and prevent themselves from reaching the goal. None of the subjects who were task independent chose to give up chips.

significant difference between the reliability measure of the PURB agent and people in the U.S.¹⁰

When separating the analysis by participants' dependency role, we found that in the task dependent role, the PURB agent exhibited a higher reliability than people in both countries: As shown in Table II, in the task dependent role, the reliability of the PURB agent in Lebanon was higher than that of people (99% versus 87%), and this was also the case in the U.S. (72% versus 51%). In both countries there was no significant difference in the reliability of people and the PURB agent in the task independent and task co-dependent role.¹¹

The adaptation capabilities of the PURB agent were also apparent in the extent to which offers were accepted (not shown in table). On average, the PURB agent accepted proposals significantly less often than did people in the U.S. However, this trend was reversed in Lebanon, where the PURB agent accepted more proposals than did people.

5.3 Analysis of Behavior: Proposals

To analyze the types of proposals that are made and accepted in both cultures, we will denote the *actual score* of a proposal to be the score (as computed by the CT scoring function of Section 3.2) in the case that both players fully commit to the chips they agreed to send. Using the notation defined in Section 5.2, we say that the score to player i at round $n + 1$ from a proposal $O = (O_i, O_j)$ is its score in the game after both i delivers all of the promised chips O_j to j and j delivers all of the promised chips O_i to i . Let C_i denote the chips in the possession of player i at round n . The actual score of i for an exchange O is $r_i(\{C_i \cup O_j \setminus O_i\})$. Similarly, the actual score of j for an exchange O is defined as $r_j(\{C_j^n \cup O_i \setminus O_j\})$. For the remainder of this paper, we abbreviate the term “actual score” to “score”.

Figure 2 presents the average score for people and for the PURB agent in Lebanon (top) and the U.S. (bottom).¹² The figure is shown as a plot that describes the average promised score of a proposal to people (x-axis), and to the PURB agent (y-axis). The plot lists the proposals made by people as well as the PURB agent, as well as the score to both participants from the subset of proposals that were actually accepted. The edges that are displayed in the graph highlight the difference between those scores associated with the proposals made by people and the PURB agent and those scores associated with accepted proposals. Intuitively, the score to *people* from offers made by the PURB agent represent the degree to which the PURB agent was generous, whereas the score to the *PURB agent* from offers made by the PURB agent represent the degree to which the PURB agent was selfish (and symmetrically for people).

As shown by the figure, in Lebanon, the PURB agent made offers that were

¹⁰The difference between the reliability measure of the PURB agent (0.98) and people (0.92) in Lebanon was small but statistically significant. The difference between the reliability measure of the PURB agent (0.62) and people (0.65) in U.S. was small and not statistically significant.

¹¹The difference between the reliability of the PURB agent (0.59) and people (0.78) in the U.S., was not statistically significant in the task independent and co-dependent conditions.

¹²Note that these score do not represent the actual performance in the game, because they assume that agents transfer all of the chips in the agreement. We cannot directly compare the average score of accepted proposals with the average performance per game shown in Figure I.

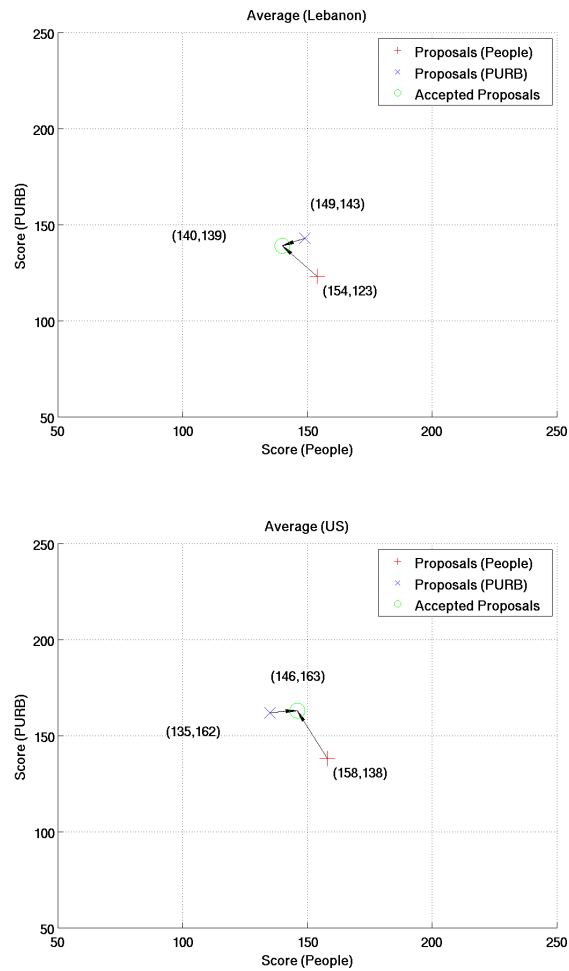


Fig. 2: Proposed offers versus accepted offers (averaged over all dependency roles). Lebanon on top; U.S. on bottom

more generous than selfish: The average score from proposals made by the PURB agent to people (149 points) was higher than the score to the PURB agent (143 points). This difference was small but statistically significant. People in Lebanon were significantly less generous than the PURB agent. The average score from proposals made by people to people (154 points) was higher than the score to the PURB agent (123 points). There was no difference in score to people and the PURB agent from accepted offers (140 points versus 139 points).

In contrast, in the U.S., the PURB agent made offers that were more selfish than generous: The average score from proposals made by the PURB agent to people (135 points) was lower than the score to the PURB agent (162 points). People in the U.S. proposed offers that were more selfish, similarly to Lebanon. The average score from proposals made by people to people (158 points) was higher than the

score to the PURB agent (138 points). The score to the PURB agent from accepted proposals was significantly higher than the score to people (163 versus 146 points).

Lastly, we note that the fact that the results vary significantly across these conditions show that these dependency relationships are significant determinants of performance. We report the average results simply to provide a general performance rating for the agent for all conditions.

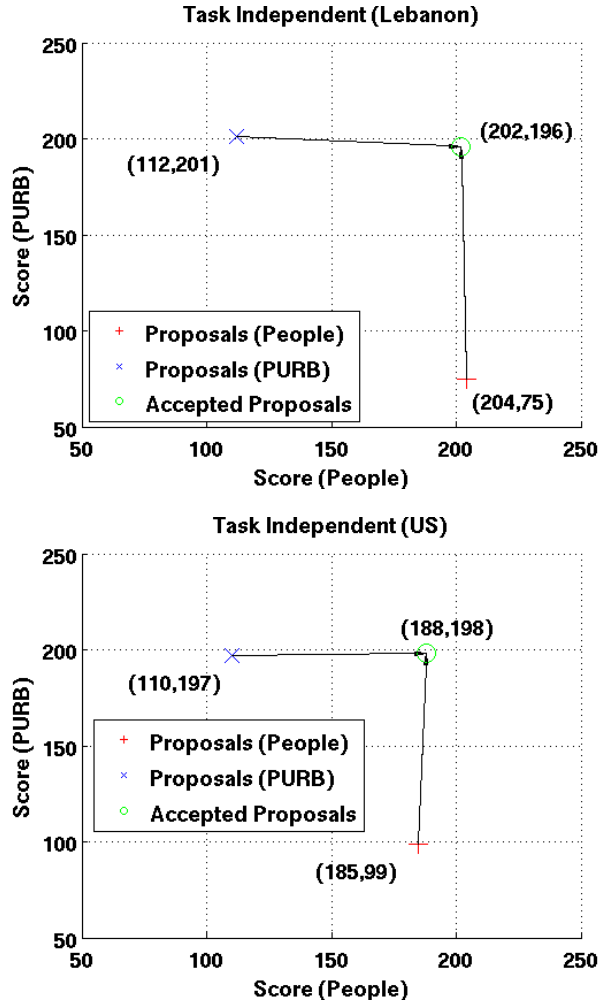


Fig. 3: Proposed offers versus accepted offers for task codependent participants. Lebanon on top; U.S. on bottom

We now turn to compare between the proposals people make before they interacted with the PURB agent. Recall that humans were always designated as first proposers in our scenario. The purpose for this was to be able to compare between their negotiation behavior prior to reacting to the PURB agent to PURB. We can attribute differences in people’s offers to their country, given that we control for

	Co-Dep.	Task Ind.	Task Dep.	Average
Lebanon	(85, 62)	(199, 49)	(92, 186)	(125, 99)
U.S.	(125, 121)	(173, 125)	(150, 174)	(149, 138)

Table III: Proposed scores from first proposals made by people: to themselves (left), and to the PURB agent (right)

gender in our analysis, and that subjects share similar demographic backgrounds. Table III presents the scores associated with first proposals for the human proposer (left), which correspond to the degree to which the proposer is selfish, and for the agent responder (right), which corresponds to the degree the proposer is generous. As can be seen in the table, proposals in the U.S. were more beneficial to both people and to the PURB agent than proposals in Lebanon: The proposed score to human proposers from first proposals in the U.S. was 149 points, which was significantly higher than the score to human proposers from first proposals in Lebanon, which was 125 points. Similarly, the score to the PURB agent from first proposals in the U.S. was 138 points, which was significantly higher than the score to the PURB agent from first proposals in Lebanon, which was 99. The fact that proposers in the U.S. are more generous than their Lebanese counterparts align with studies in behavioral economics that have compared between people’s behavior in the ultimatum game across different cultures [Roth et al. 1991; Henrich 2000]. In all of these studies, proposers in the U.S. made more generous offers than proposers in the Middle East or in developing nations. Because first proposers did not yet interact with the PURB agent, these results support the claim that at least some of the negotiation behavior exhibited by people in the U.S. and Lebanon can be attributed to cultural differences. A possible explanation for the fact that proposers in the U.S. were more generous than proposers in Lebanon could be attributed to the fact they were also less reliable than their Lebanese counterparts (see Table II). Under this interpretation, offers made by human proposers in the U.S. were more generous because the proposers were not prepared to commit to their proposals and transfer all of the promised chips. These results also serve to extend existing studies in investigating more complex dependency relationships than those considered in the ultimatum game. For example, in the task independent condition, the proposed score to human proposers from first proposals in the U.S. was lower than the equivalent score proposed in Lebanon (173 versus 199 points).

6. DISCUSSION

The main hypotheses of this study was that an agent that modeled and adapted to the cooperativeness measures exhibited by people would be able to negotiate proficiently across these different cultures: Our analysis confirmed that negotiators in Lebanon and in the U.S. varied widely in their reliability and helpfulness measures. People in Lebanon were significantly more reliable than people in the U.S. (Table II); People in Lebanon made proposals that were less generous than people in the U.S. (Figure 2). Our second hypothesis was also confirmed, in that the PURB agent negotiated as well as people when combining its performance across the U.S. and Lebanon (Table I).

As shown by the analysis, the PURB agent adapted a different negotiation strat-

egy depending on the cultural affiliation of its negotiation partner, even though this affiliation was unknown to the PURB agent. This was apparent in several factors of its behavior. First, the reliability of the PURB agent was significantly higher in Lebanon than in the U.S. (Table II). Second, the PURB agent made proposals that were significantly more selfish in the U.S. than in Lebanon (Figure 2).

Despite the demonstrated ability of the PURB agent to adapt to different cultures, its performance was not consistent in both countries. As shown in Table I, in the U.S., the PURB agent outperformed people, while in Lebanon, the PURB agent was outperformed by people. To explain this phenomenon, we will further investigate the interaction between performance and cooperativeness of participants in the different dependency roles.

In the co-dependent condition, there was no significant difference in the likelihood to get to the goal of the PURB agent and people in both countries. Therefore the difference in performance in this condition must be attributed to the discrepancy in the scores that are associated with agreements. Figure 4 plots the scores associated with proposals made by both participants and the scores associated with accepted proposals in the co-dependent condition. In Lebanon the score from accepted proposals was more beneficial to people than the PURB Agent (127 versus 118 points), whereas in the U.S., the score from accepted proposals was more beneficial to the PURB agent (151 versus 136 points).

There were two reasons for this discrepancy, both of which are concerned with the difference in people's reliability measures in Lebanon and the U.S. First, the agent's social utility function, as described in Section 4 directly depended on its estimate of the cooperativeness measure of the other participant: Its social utility increased as the reliability and helpfulness measures of the other participant increased (and conversely for the case when these measures decreased). People in the U.S. exhibited medium reliability (64%, as shown in Table II) in the co-dependent role. In contrast, in Lebanon, people exhibited high reliability (92%, as shown in Table II) when they were co-dependent. Therefore the proposals that were made by the PURB agent in Lebanon were likely to be generous to people, and conversely for offers that were made by the PURB agent in the U.S.

Second, according to the rules of behavior for the PURB agent that are described in Section 4, the PURB agent accepted a proposal if it provided it with a higher social utility than the utility from the proposal the PURB agent would propose itself. Because people in the U.S. exhibited low reliability measures, it was more likely that the PURB agent would choose to make selfish offers (a 2:1 ratio in favour of the PURB agent), in order to increase the number of chips it would receive. This is shown in Figure 4, where the promised score from proposals made by the PURB agent provide 132 points for people and 148 points for the PURB agent. Therefore, the PURB agent was not likely to accept offers that would provide more benefit to people than to the PURB agent. Because people were highly reliable when they were co-dependent, the PURB agent generated more proposals that were egalitarian in this condition (a 1:1 ratio). This can be seen in Table 4, where the average is no difference in score from proposals made by the PURB agent to people (141 points) and the PURB agent (140 points). Therefore, in Lebanon the PURB agent was more likely to accept offers that favored people.

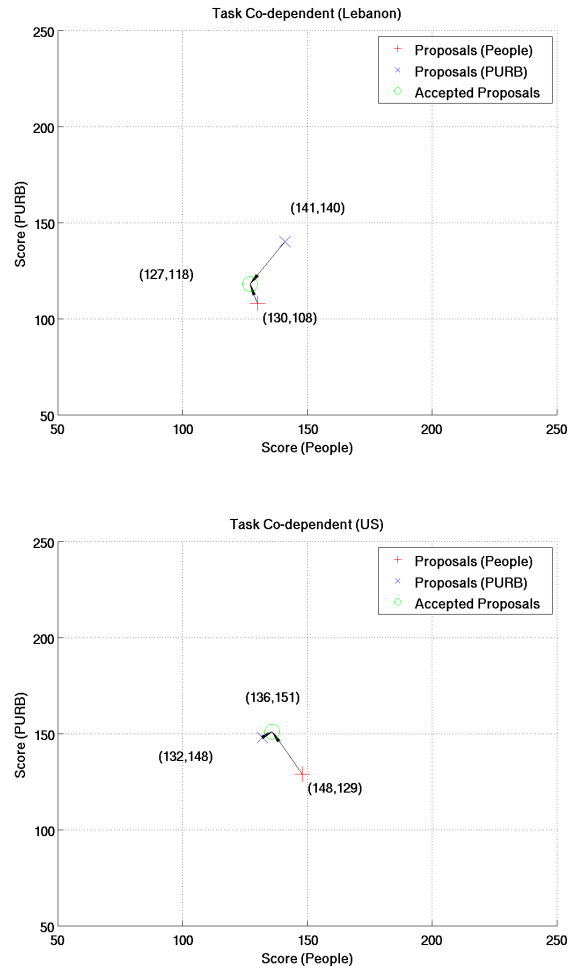


Fig. 4: Proposed offers versus accepted offers for task codependent participants. Lebanon on top; U.S. on bottom

When participants were allocated the task independent role, there was no significant difference in the likelihood of people and the PURB agent to get to the goal. (In fact, all of the subjects in this condition got to the goal.) Again, the difference in performance in this condition has to be attributed to the discrepancy in the scores that are associated with agreements.

To explain this, we show Figure 3 that plots the scores associated with proposals made by both participants and the scores associated with accepted proposals. As shown in the figure, the score from accepted proposals in Lebanon favored people (202 versus 196 points), whereas the score from accepted proposals in the U.S. favored the PURB agent (198 versus 188 points, the differences are small, but statistically significant). Using similar arguments to the ones we specified for the task co-dependent condition, we can claim that (1) in the U.S., the PURB agent

generated and accepted proposals that tended to favor the PURB agent, and (2), in Lebanon, the PURB agent generated and accepted proposals that tended to favour people.

In addition, note that when the PURB agent was task independent, then people were task dependent. The reliability of people when they were task dependent was significantly lower than the reliability of the PURB agent when it was task independent (87% versus 99%). Therefore, in the task dependent role, people were more successful than the PURB agent despite being less reliable than the PURB agent.

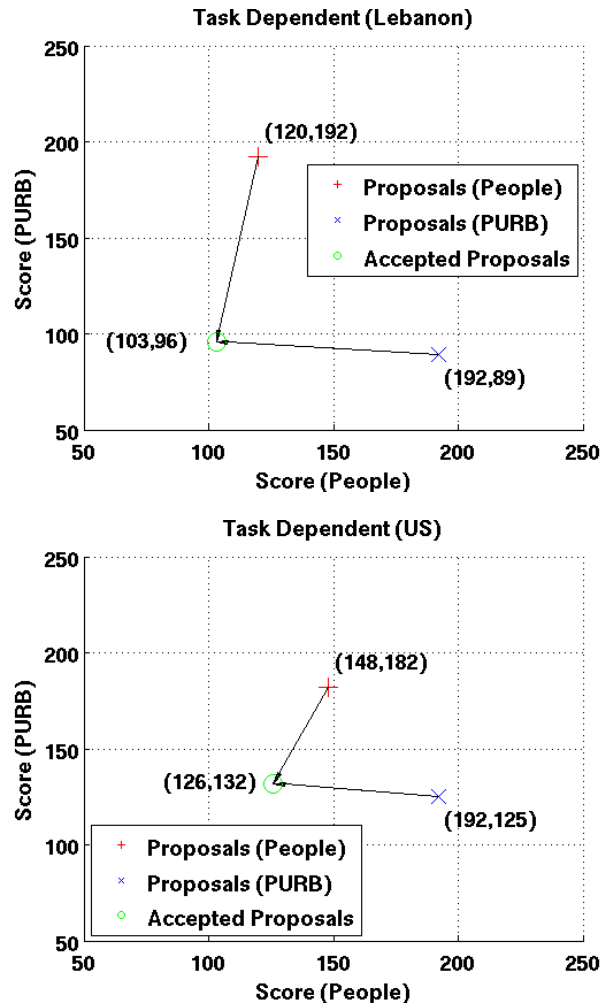


Fig. 5: Proposed offers versus accepted offers for task dependent participants. Lebanon on top; U.S. on bottom

In contrast to the findings in the task co-dependent and task independent roles, in Lebanon there was no significant difference in the score from accepted offers

in the task dependent role. Figure 5 plots the scores associated with proposals made by both participants and the scores associated with accepted proposals. As shown in the figure, the score to the PURB agent from accepted was 196 points, and the score to people from accepted proposals was 202 points. This difference was not statistically significant. Therefore, the difference in performance in this condition has to be attributed to the discrepancies in getting to the goal square. In the task dependent condition people in Lebanon were more likely to get the goal than the PURB agent, and the PURB agent was more likely to get to the goal than people. To see why, consider that when the PURB agent plays the dependent role, then people play the independent role. In the U.S., the reliability of people in the independent role (78%) was higher than the reliability of the PURB agent in the dependent role (72%). In Lebanon, the reliability of the people in the independent role (94%) was lower than that of the PURB agent in the dependent role (99%). Thus people were more likely to get the chips they needed to get to the goal than the PURB agent. Another contributor to this result is that the PURB agent was pre-programmed to be fully reliable after the first agreement and always transfer its promised chips. This was done for two reasons. First, because the agent could not adapt to the reliability of its partner without any experience. Second, because of recent results from the psychological literature showing that early lapses in reliability at the onset of relationships lead to irreversible damage to cooperation [Lount Jr et al. 2008]. Together, these aspects explain the performance of the PURB agent in Lebanon.

To summarize our findings, when people’s reliability measure was clear-cut (such as in the U.S., where people’s reliability is consistently low), the PURB agent was able to adapt a negotiation strategy that allowed it to outperform people. In contrast, when people’s reliability measure was less certain, (such as in Lebanon, where people exhibited medium or high reliability measures depending on the dependency relationships) the PURB agent adopted a high reliability measure, and this allowed people to outperform the PURB agent in this condition. These results have several implications for the design of agents that negotiate with people. First, they show that adaptation to the behavioral traits exhibited by people is a viable approach towards creating agents that are able to negotiate proficiently across cultures. Additionally, they show that in some cases, people may be able to take advantage of adaptive agents by adopting behavior that is ambiguous. In our study, the PURB agent misclassified people in Lebanon to be highly reliable in all dependency roles, and thus accepted proposals that favoured people.

We conclude with an example of the way the PURB agent was able to adapt to the different behavioral traits that were exhibited in both countries. The example is taken from two games, one in Lebanon and one in the U.S., in which the PURB agent was task independent, and its negotiation partner was task dependent. In both of these games, there was an identical offer made by the human participant that would allow both players to reach the goal. This example consisted of the person asking the PURB agent to exchange 3 green chips in return for 2 grey chips. The proposal was accepted by the PURB agent. In the game played in Lebanon, the PURB agent chose to fulfil the agreement and sent all of the promised chips. However, in the game played in the U.S., the PURB agent chose not to fulfil the

agreement and did not send any of its promised chips.

7. CONCLUSION AND FUTURE WORK

This paper described a new agent-design that uses adaptation techniques to negotiate with people across different cultures. It focused on a repeated negotiation setting in which participants need to accrue and exchange resources in order to complete their individual goals, and agreements are not binding. This setting was implemented using a test-bed that consists of a computer board game that provided a task analogy to the types of interactions that occur in the real world. The decision-making model for the agent consisted of two components: a social utility function that represented the extent to which its negotiation partner was helpful and reliable traits of other participants, as well as a rule-based mechanism that used the utility function to make decisions in the negotiation process. This agent was able to negotiate proficiently with people across cultures. This provides an empirical proof of the benefit towards using adaptation mechanisms in cross cultural settings in which there may not exist prior data consisting of people's play. The decision-making mechanism that we described in Section 4 can be extended to enable PURB to negotiate successfully with different types of people or new types of settings. For example, additional social factors such as competitiveness or risk aversion may be added to PURB's social utility function, and the rules guiding its behavior can be modified to account for different negotiation protocols, for example, ones that limit the number of repeat negotiation rounds. We showed that the rules are general enough to allow PURB to negotiate proficiently in environments that vary the dependency relationships between participants.

We are investigating several future directions for this work. First, we are using the data we collected for this study to build predictive models of human negotiation behavior, including the types of offers they make and the extent to which they follow their commitments. Second, we are designing computational models of the way collectivism affects the evolution of personal disputes into large-scale intercultural conflict. Lastly, we are using CT as a highly configurable tool that allows for the specification of games that reflect different task environments that may affect the way people of different cultures make decisions, such as group solidarity and in-group biases.

8. ACKNOWLEDGEMENTS

This study was partially funded by MURI grant number W911NF-08-1-0144, the National Science Foundation under grant 0705587 and the University of Maryland. Thanks very much to Yael Ejgenberg and Yael Blumberg for their help in programming the PURB agent and assistance with data analysis, and to Swapna Reddy, Leandra King and Konstantin Pozin for their valuable assistance in administering experiments.

REFERENCES

- ADAIR, W. AND BRETT, J. 2004. Culture and negotiation processes. *The handbook of negotiation and culture*, 158–176.
- BOHNET, I., GREIG, F., HERRMANN, B., AND ZECKHAUSER, R. 2008. Betrayal aversion: Evidence
ACM Transactions on Intelligent Systems and Technology, Vol. V, No. N, Month 20YY.

- from Brazil, China, Oman, Switzerland, Turkey, and the United States. *The American Economic Review* 98, 1, 294–310.
- BUSINESWEEK. 2007. Upwardly mobile in Africa. *Bloomberg Businessweek*.
- BYDE, A., YEARWORTH, M., CHEN, K., BARTOLINI, C., AND VULKAN, N. 2003. Autona: A system for automated multiple 1-1 negotiation. In *Proceedings of the 4th ACM conference on Electronic commerce*. ACM New York, NY, USA, 198–199.
- CAMERER, C. F. 2003. *Behavioral Game Theory. Experiments in Strategic Interaction*. Princeton University Press.
- CHARNESS, G. AND RABIN, M. 2002. Understanding social preferences with simple tests. *Quarterly Journal of Economics* (117), 817–869.
- DE DREU, C. AND VAN LANGE, P. 1995. The impact of social value orientations on negotiator cognition and behavior. *Personality and Social Psychology Bulletin* 21, 1178–1188.
- FICICI, S. G. AND PFEFFER, A. 2008. Simultaneously modeling humans' preferences and their beliefs about others' preferences. In *Proc. 7th International Joint Conference on Autonomous Agents and Multi-agent Systems (AAMAS)*.
- GAL, Y., GROSZ, B., KRAUS, S., PFEFFER, A., AND SHIEBER, S. 2010. Agent decision-making in open mixed networks. *Artificial Intelligence* 174, 18, 1460–1480.
- GAL, Y., GROSZ, B., SHIEBER, S., PFEFFER, A., AND ALLAIN, A. 2007. The effects of task contexts on the decision-making of people and computers. In *Proc. of the Sixth International Interdisciplinary Conference on Modeling and Using Context, Roskilde University, Denmark*.
- GAL, Y. AND PFEFFER, A. 2007. Modeling reciprocity in human bilateral negotiation. In *Proc. 22nd National Conference on Artificial Intelligence (AAAI)*.
- GELFAND, M. J. AND BRETT, J. M. 2004. *Handbook of Negotiation and Culture*. Stanford University Press.
- GELFAND, M. J. AND CHRISTAKOPOULOU, S. 1999. Culture and negotiator cognition: Judgment accuracy and negotiation processes in individualistic and collectivistic cultures. *Organizational Behavior and Human Decision Processes* 79, 3, 248–269.
- GELFAND, M. J., HIGGINS, M., NISHII, L. H., RAVER, J. L., DOMINGUEZ, A., MURAKAMI, F., YAMAGUCHI, S., AND TOYAMA, M. 2002. Culture and egocentric perceptions of fairness in conflict and negotiation. *Journal of Applied Psychology* 87, 5, 833–845.
- GELFAND, M. J. AND REALO, A. 1999. Individualism-Collectivism and Accountability in Intergroup Negotiations. *Journal of Applied Psychology* 84, 5, 721–736.
- GROSZ, B., KRAUS, S., TALMAN, S., AND STOSSEL, B. 2004. The influence of social dependencies on decision-making. Initial investigations with a new game. In *Proc. 3rd International Joint Conference on Autonomous Agents and Multi-agent Systems (AAMAS)*.
- HENRICH, J. 2000. Does culture matter in economic behavior? Ultimatum game bargaining among the Machiguenga of the Peruvian Amazon. *American Economic Review* 90, 4, 973–979.
- HOFSTEDE, G. 1980. *Culture's consequences: Comparing values, behaviors, institutions, and organizations across nations*. Sage Publications, Inc.
- IMAI, L. AND GELFAND, M. J. 2010. The culturally intelligent negotiator: The impact of cultural intelligence (CQ) on negotiation sequences and outcomes. *Organizational Behavior and Human Decision Processes*.
- JONKER, C., ROBU, V., AND TREUR, J. 2007. An agent architecture for multi-attribute negotiation using incomplete preference information. *Autonomous Agents and Multi-Agent Systems* 15, 2, 221–252.
- KAMAR, E., HORVITZ, E., AND MEEK, C. 2008. Mobile opportunistic commerce: Mechanisms, architecture, and applications.
- KASHIMA, Y., KASHIMA, E., CHIU, C., FARSIDES, T., GELFAND, M., HONG, Y., KIM, U., STRACK, F., WERTH, L., YUKI, M., ET AL. 2005. Culture, essentialism, and agency: Are individuals universally believed to be more real entities than groups? *European Journal of Social Psychology* 35, 2, 147–169.
- KATZ, R. AND KRAUS, S. 2006. Efficient agents for cliff edge environments with a large set of decision options. In *Proc. 5th International Joint Conference on Autonomous Agents and Multi-agent Systems (AAMAS)*.

- KRAUS, S., HOZ-WEISS, P., WILKENFELD, J., ANDERSEN, D., AND PATE, A. 2008. Resolving crises through automated bilateral negotiations. *Artificial Intelligence* 172, 1, 1–18.
- KRAUS, S. AND LEHMANN, D. J. 1995. Designing and building a negotiating automated agent. *Computational Intelligence* 11, 132–171.
- LIN, R. AND KRAUS, S. 2010. Can automated agents proficiently negotiate with humans? *Communications of the ACM* 53, 1, 78–88.
- LIN, R., OSHRAT, Y., AND KRAUS, S. 2009. Investigating the benefits of automated negotiations in enhancing people’s negotiation skills. In *Proc. 8th International Joint Conference on Autonomous Agents and Multi-agent Systems (AAMAS)*. 345–352.
- LOUNT JR, R., ZHONG, C., SIVANATHAN, N., AND MURNIGHAN, J. 2008. Getting off on the wrong foot: The timing of a breach and the restoration of trust. *Personality and Social Psychology Bulletin* 34, 12, 1601.
- MURPHY, R. 2004. Humanrobot interaction in rescue robotics. *IEEE Transactions On Systems, Man, and Cybernetics Part C: Applications and Reviews* 34, 2 (May).
- OSHRAT, Y., LIN, R., AND KRAUS, S. 2009. Facing the challenge of human-agent negotiations via effective general opponent modeling. In *Proc. 8th International Joint Conference on Autonomous Agents and Multi-agent Systems (AAMAS)*. 377–384.
- RAJARSHI, D., HANSON, J. E., KEPHART, J. O., AND TESAURO, G. 2001. Agent-human interactions in the continuous double auction. In *Proc. 17th International Joint Conference on Artificial Intelligence (IJCAI)*.
- ROTH, A., PRASNIKAR, V., OKUNO-FUJIWARA, M., AND ZAMIR, S. 1991. Bargaining and market behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An experimental study. *The American Economic Review* 81, 5, 1068–1095.
- SCHURR, N., PATIL, P., PIGHIN, F., AND TAMBE, M. 2006. Using multiagent teams to improve the training of incident commanders. In *Proc. 5th International Joint Conference on Autonomous Agents and Multi-agent Systems (AAMAS)*.
- TALMAN, S., GAL, Y., HADAD, M., AND KRAUS, S. 2005. Adapting to agents’ personalities in negotiation. In *Proc. 4th International Joint Conference on Autonomous Agents and Multi-agent Systems (AAMAS)*.
- WEINGART, L. AND OLEKALNS, M. 2004. Communication processes in negotiation: Frequencies, sequences, and phases. *The handbook of negotiation and culture*, 143–157.
- ZUCKERMAN, I., KRAUS, S., AND ROSENSCHEIN, J. 2007. Using focal point learning to improve tactic coordination in human-machine interactions. In *Proc. 20th International Joint Conference on Artificial Intelligence (IJCAI)*. 1563–1568.

Received March 2010; revised July 2010; accepted October 2010