Learning Human Negotiation Behavior Across Cultures

Galit Haim[1], Ya'akov Gal[2], Sarit Kraus[1,4], and Yael Blumberg[3]

[1]Computer Science Department, Bar Ilan University

[2]Department of Information Systems Engineering, Ben-Gurion University of the Negev

[3]Gonda Brain Research Center, Bar Ilan University

[4]Institute for Advanced Computer Studies University of Maryland, College Park 20742

ABSTRACT

The ability to negotiate successfully is critical in many social interactions. The dissemination of applications such as the Internet across geographical and ethnic borders are opening up opportunities for computer agents to negotiate with people of diverse cultural and organizational affiliation. These automated negotiators should be able to proficiently interact and collaborate with their human partners. In this paper we compare several techniques for modeling the negotiation behavior of people across three different countries. Culture plays an important role in people's decision making and people differ in the way make offers and fulfil their commitments in negotiation across cultures. We consider a setting that included multiple rounds of negotiation with non-binding agreements. Participants in each of the countries interacted with a computer agent that used a baseline negotiation strategy that adapted to the extent to which participants were helpful and reliable. The models considered various features of the negotiation task, such as the extent to which proposals are generous and helpful to participants and whether participants fulfil their agreements. Our models achieved high accuracy rates in predicting important decisions that are made in negotiation such as whether participants reach agreement and the extent to which they are reliable in different situations. We show that the features that best vary the prediction accuracy depend on people's cultural affiliation in addition to their actual negotiation behavior. These models for predicting human behavior in negotiation will form the basis of a computer agent that can successfully negotiate with people across cultures.

1. INTRODUCTION

Negotiation is a tool widely used by humans to resolve disputes in settings as diverse as business transactions, diplomacy and personal relationships. Computer agents that negotiate successfully with people have profound implications: They can negotiate on behalf of individual people or organizations (e.g., bidders in on-line auctions (Kamar et al., 2008; Rajarshi et al., 2001)); they can act as training tools for people to practice and evaluate different negotiation strategies in a lab setting 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 purpose of this paper was to investigate the role of culture in predicting people's negotiation behavior with computer agents. Our long-term goal is to be able to build computer agents that can learn to negotiate proficiently with people across different cultures. Culture is a key determinant of the way people interact and reach agreements in different social settings. Advancing technology such as the Internet requires that computer systems negotiate proficiently with people across geographical and ethnic boundaries. It is thus important to understand the decision-making strategies that people of different cultures deploy when computer systems are among the members of the groups in which they work and to determine their responses to different kinds of decision-making behavior of others.

This paper investigates the hypothesis that explicitly representing behavioral traits that vary across cultures will improve the ability of computer agents to predict human negotiation behavior, and in turn, improve the performance of computer agents when negotiating with people. To evaluate this hypothesis, we collected data of human negotiation behavior in laboratory conditions from three different countries, including Israel, Lebanon and the U.S. We used an identical negotiation scenario in each country that required people to complete a task by engaging in bilateral negotiation rounds with non-binding agreements. The negotiation protocol included alternating take-it-or-leave-it offers for the exchange of resources. Agreements were not binding, and participants were free to choose the extent to which they fulfilled their commitments.

We collected data of people playing a computer agent referred to as the Personality, Utility and Rules Based agent (PURB), that was designed to adapt to the extent to which the other participant was helpful and kept commitments over time. We included various features relating to people's negotiation behavior in our prediction models. These included the extent to which their proposals were generous and selfish, the extent to which they exhibit reliable behavior in past behavior, and their current performance in the game. We used standard machine-learning techniques to build predictive models of the types of offers people make, whether they are able to complete their tasks, and the extent to which they fulfilled their commitments.

Our results showed that the optimal features to use for the various learning tasks varied according to the cultural affiliation of the participants. In general, models that learned from data that was restricted to a single country were as well as, or better than models that learned from combining the data for all countries. In addition, there were considerable differences in people's negotiation behavior across the different countries, and these differences affected the prediction accuracy of the models.

The contributions of this work are twofold. First, it shows that people's cultural diversity affects the accuracy of prediction models of human negotiation behavior. Second, it suggests a new paradigm for constructing automatic agents that can negotiate with people across culture, by combining the individual learner models into an agent's decision-making model.

2. RELATED WORK

There is a body of work in the psychological and social sciences that investigate cross-cultural behavior among human negotiators (De Dreu and Van Lange, 1995; Gelfand and Dyer, 2001; Gelfand et al., 2002). However, there are scant computational models of human negotiation behavior that reason about cultural differences. In artificial intelligence, past works have used heuristics, equilibrium strategies, and opponent modeling approaches toward building computer agents that negotiate with people. For a recent comprehensive review, see Lin and Kraus (2010). Within repeated negotiation scenarios, Kraus et al. (2008) 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. Jonker et al. (2007) designed computer strategies that involve the use of concession strategies to avoid impasses in the negotiation. Byde et al. (2003) constructed agents that bargain with people in a market setting by modeling the likelihood of acceptance of a deal and allowing agents to renege on their offers. Kenny et al. (2007) constructed agents for the training of individuals to develop leadership qualities and interviewing capabilities.

Recent approaches have used learning techniques to model the extent to which people exhibit different social preferences when they accept offers in one-shot and multiple interaction scenarios (Gal et al., 2009; Oshrat et al., 2009; Lin et al., 2008). 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).

To date, 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 an agent for negotiating with multiple participants that 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 and twenty subjects in three countries.

3. IMPLEMENTATION USING THE COLORED TRAILS TEST-BED

Our study was based on the Colored Trails (CT) game (Grosz et al. (2004)), a test-bed for investigating decision-making in groups comprising people and computer agents. Colored Trails is Free Software and is available for download at `http://www.eecs.harvard.edu/ai/ct`. 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. Players were issued 24 colored chips at the onset of the game.

Figure 1 shows the CT 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. Figure 2 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 3 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 yellow chips.

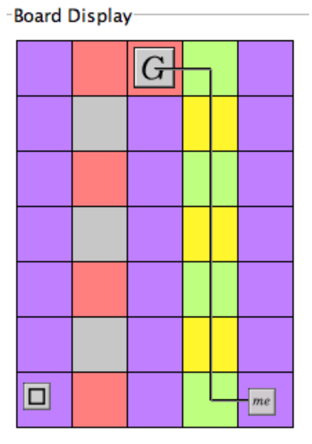

Figure 1: An example of a CT Board

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* phase, one of the players was designated the role of a proposer, and could make an offer to the other player, who was designated the responder. In turn, the responder could accept or reject the offer. If the offer was rejected, then players switched roles: the responder

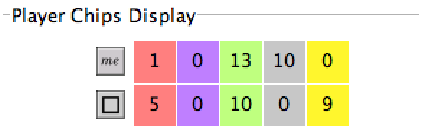

Figure 2: Chip Display Panel (showing the chips in the possession of both participants)

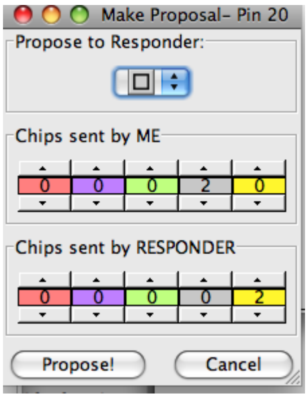

Figure 3: Communication Panel (used by participants to make offers)

became the proposer and the proposer became the responder. This sequence of alternating offers continued until an offer was accepted, or the time limit for the communication phase was up. 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 fulfill 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. Players alternated their roles, such that the first proposer in the previous communication phase was designated as a responder in the next communication phase, 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. 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 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.

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 agent 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 agents 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.

CT is thus particularly suitable for modeling 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). In addition, it has been shown that people that use CT generally display more cooperative behavior than identical decision-making scenarios that involve more abstract representations such as payoff matrices or decision-trees (Gal et al., 2007). This incentive for cooperation may allow both parties in negotiation to reach agreements more quickly. Both of these are important qualities to multi-cultural disputes that are often volatile.

4. DATA COLLECTION

We collected data from three countries, Lebanon, the U.S. and Israel. Two hundred and twenty two subjects participated the study. Ninety subjects were students studying in the Beirut area, 119 subjects were students enrolled in college degree programs at institutions in the greater Boston area, and 32 subjects were students from Bar-Ilan University. The demographics of the subjects were as follows: In the U.S., 44.5% of subjects identified as Caucasian, 21% identified as African American, 16% of subjects identified as Asian Americans and 4.2% of subjects identified as Latin or Hispanic. The mean age of subjects in the U.S. was 24. In Lebanon, 80% of the subjects identified themselves as Arab or Lebanese, while 10% of the subjects identified themselves a Phoenician. The mean age of subjects in Lebanon was 22. In Israel, all subjects identified themselves as Israeli Jews. The mean age of subjects in Israel was 24.

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. 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 one or two games with the PURB agent, but were told they would be playing with different people. Authorization for this slight 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.

The study included the CT scenario that was described in Section 5. We used two different types of boards. 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 1. In this game both "me" and "O" players were missing three chips to get to the goal. The relevant path from the point of view of the "me" player is outlined.

Each subject played a single CT game. To standardize our experiments, in all CT games we ran, people were designated as first proposers, and the PURB agent was designated as the first responder. Each subject was randomly assigned one of the following dependency roles: a task co-dependent participant that was paired with a task co-dependent PURB agent; a task independent participant that was paired with a task dependent PURB agent; or, a task dependent participant that was paired with another task independent PURB agent.

5. EMPIRICAL METHODOLOGY

We focused on learning four key features that characterized human negotiation in our study. Specifically we considered the following learning models.

- Reliability Model - the extent to which a person was reliable in the negotiation.
- Acceptance Model - the likelihood of accepting a given proposal.
- Human Reached Goal Model - the likelihood for the person to get to the goal following accepting a given proposal.
- Agent Reached Goal Model - the likelihood for the agent to get to the goal.

5.1 Potential Features

We defined a general set of features to be used by the various learning models. To describe the features we lay out the following notation. Let n denote an arbitrary communication phase in the game. For any two participants i and j, let c_i denote the set of chips in possession of i at phase n in the game. Let $p^n = (p_i, p_j)$ denote a proposal at round n, where $p_i \subseteq c_i$ is the set of chips that i agreed to send to j, and let $p_i^* \subseteq c_i$ be the set of chips actually sent by i following the agreement. Similarly define c_j, p_j, and p_j^* from the point of view of player j. We can now describe the following features:

- The *current score* for i at round n measures the score in the game given its current set of chips c_i. This is defined as $s_i^n(c_i)$.
- the *resulting score* to player i at round n measures the score in the game that i would receive in the case that j sent all of its promised chips p_j during the transfer phase, and i sent all of its promised chips p_i. This is defined as $s_i^n(h_i)$, where $h_i = c_i \cup p_j \setminus p_i$.
- the *score-based-reliability* of player i at round n, denoted r_i^n is the extent to which player i fulfilled an agreement $p^n = (p_i, p_j)$ at round n. It computes reliability measures solely for negotiation rounds in which agreements were reached. This is defined as the ratio between the score to j given the chips that i actually transferred, and the score that j would receive if i fulfilled its agreement. Formally, we define
$$r_i^n = \frac{s_j(c^n \cup p_i^*)}{s_i^n(c^n \cup p_i)}$$

We now use the definitions above to describe the set of potential features for our models. Each feature is described from the point of view of a general player i. We also considered the symmetrical feature from the point of view of player j.

- The current round n.
- the current score for i at round n.
- the resulting score to i given proposal $p^n = (p_i, p_j)$ at round n.
- the score-base-reliability of i at round n.
- the weighted score-base-reliability wr_n of i. This is defined as follows: for $n > 1$ it is a weighted average of the score-base-reliability of i at round n and $n - 1$
$$(0.7 \cdot r_n) + (0.3 \cdot wr_{n-1})$$
For $n = 1$, the score-based-reliability of i is defined as an initial score-based-reliability of 1.
- the generosity of player i at round n. This feature clustered the chip sets offered by both players in round n into three classes measuring the difference in the number of chips proposed by i to i and j.
  - $|p_i| < |p_j|$ (player i offered less chips to j than player j offered to i),
  - $|p_i| = |p_j|$ (both players offered the same chips to each other),
  - $|p_j| < |p_i|$ (player i offered more chips to player j than did player j.)
- The dependency role of player i at round n, which was one of two classes: task independent, task dependent.
- Missing chips: this feature includes the total number of chips that player i needed to get to its goal given its position in the board at round n.

5.2 Learning Algorithms

Our study was based on the Weka framework (`http://www.cs.waikato.ac.nz/ml/weka/`), a repository of machine learning algorithms that is freely available on the web. We used the following learner models on all prediction tasks.

- K-Nearest Neighbors classifier (KNN): a method for classifying objects based on closest training examples in the feature space. K-NN is a type of instance-based learning, or lazy learning where the function is only approximated locally and all computation is deferred until classification. The k-nearest neighbor algorithm is amongst the simplest of all machine learning algorithms: an object is classified by a majority vote of its neighbors, with the object being assigned to the class most common amongst its K nearest neighbors (K is a positive integer, typically small).

- SMO-PUK: Sequential Minimal Optimization using the PUK (The Pearson VII function-based Universal Kernel) kernel function was used. Support Vector Machines (SVMs) application are used for solving classification and regression problems. SVM has high generalization performance and ability to model non-linear relationships with a suitable kernel function. This kernel function transforms the non-linear input space into a high dimensional feature space in which the solution of the problem can be represented as being a straight linear classification or regression problem. There are a lot of possible kernel functions that can be used to create such high dimensional feature space. The most commonly used kernel functions are the linear and polynomial inner-product functions and the Radial Basis Function (RBF). But, the applicability, suitability, performance and robustness of the PUK kernel function in comparison to the commonly applied kernels is robust and has an equal or even stronger mapping power as compared to the standard kernel functions leading to an equal or better generalization performance of SVMs.

- J48: J48 implements C4.5 algorithm, (a decision-tree algorithm) for generating a pruned or unpruned C4.5 decision tree. The decision trees generated by J48 can be used for classification. J48 builds decision trees from a set of labeled training data. It uses the fact that each attribute of the data can be used to make a decision by splitting the data into smaller subsets. To make the decision, the attribute with the highest normalized information gain is used. Then the algorithm recurs on the smaller subsets. The splitting procedure stops if all instances in a subset belong to the same class. Then a leaf node is created in the decision tree telling to choose that class.

Table 1 summarizes the number of instances used for each model, population. In Table 4, each item in the table is the success rate which is the correctly Classified Instance, based on running the algorithm (mentioned in parentheses) with ten-fold cross validation. In this cross-validation method, the data-set is randomly reordered and then split into 10 folds of equal size. In each iteration, one fold is used for testing and the other 9 folds are used for training the classifier. The test results are collected and averaged over all folds. This gives the cross-validation estimate of the accuracy.

5.3 Evaluation Criteria

We now formally define the learning models and the criteria we used to evaluate them.

Reliability Model Given that an agreement was reached over proposal $p^n = (p_i, p_j)$ at round n, the reliability model predicts the reliability of player j following the agreement. For classification purposes we discretized the reliability measure to four types: (1) whether player j did not send any of the promised chips($p_j^* = \emptyset$); (2) whether player j sent a set of chips p_j^* that were different than the set of chips it promised ($p_j^* \neq p_j$); (3) whether player j sent part of the promised chips ($p_j^* \subset p_j$); (4) whether player j sent all of the promised chips ($p_j^* = p_j$). In addition, we tested another option of 2 classes measure: (5) whether player j did not fulfill the agreement (classes (1) and (2)); (6) whether player j fulfilled the agreement in part or in full (classes (3) and (4)). We note that the definition of reliability in the evaluation criteria is different than the way reliability is measured as a feature in the learning models.

Proposal Acceptance Model Given a proposal $p^n = (p_i, p_j)$ at round n of the game, the Acceptance Model predicts whether the Human (player j) (1) accepts the proposal or (2) reject it. Each proposal within a game is an example or instance of the model.

Human Reached Goal Given a proposal $p^n = (p_i, p_j)$ at round n of the game, this model predicts if at the end of the game (1) the human will reach the goal or (2) not. Each offer within a game is an example or instance of the model.

Agent Reached goal Given a proposal $p^n = (p_i, p_j)$ at round n of the game, this model predicts if at the end of the game (1) the agent will reach the goal or (2) not. Each offer within a game is an example or instance of the model.

Table 1: Number of Instances

Model	Israel	Lebanon	U.S.	Lebanon & U.S.
Reliability	126	80	196	276
Acceptance	206	152	317	469
Human reached Goal	397	343	580	923
Agent reached Goal	397	343	580	923

6. RESULTS AND DISCUSSION

For the remainder of this paper we will use the abbreviations in Table 3 (right) to relate to each potential feature of the learning models:

Table 3 (left) lists the set of optimal features chosen for each learning tasks. The selection process for the features was carried out by hand on a held-out test-set that was not used to evaluate the learning modules.

Table 1 lists the number of instances for data used by the learning models. For all learning tasks, the lowest number of instances were in Lebanon. This is because the CT

Table 2: Behavior Measures for People (in percentages)

Model / Population	Israel	Lebanon	U.S.	Lebanon & U.S.
Reliability	55	92	65	
Acceptance	60	39	75	
Human reached Goal	80.92	95	78	86.1
Agent reached Goal	83.05	87.7	87	87.6

Table 3: Features used for Learning Modules

Model /Population	Israel	Lebanon	U.S.	Lebanon & U.S.
Reliability	CS, RS, WPR	CS, RS, WPR	CS, RS, WPR	CS, RS, WPR
Acceptance	CS, RS, OG	CS, RS, OG	CS, RS, PR, PR OG RN	CS, RS, OG
Human reached Goal	CS, PR R, MC	CS,	CS, PR	CS, PR
Agent reached Goal	CS, PR R	CS, PR R	CS, PR R	CS, PR R

Feature Key	Description (FBP=For Both Players)
CS	Current Score(FBP)
RS	Resulting Score(FBP)
PR	Previous Reliability(FBP)
WPR	Weighted Previous Reliability(FBP)
OG	Offer Generosity
R	player's Role(FBP)
MC	Missed Chips(FBP)
RN	Round Number

games played in Lebanon were on average shorter than in the other countries because players were able to reach agreements quicker.

We first analyze the prediction accuracy of each of the learning models, as shown in Table 4. With respect to the reliability model, for all countries, the features that were optimal for prediction were the current and resulting scores for both players, and weighted reliability. The best results in all countries were achieved using the PUK learner. The highest accuracy rate (97.5%) was achieved in Lebanon, the lowest accuracy rate (84.127%) was achieved in Israel, and the accuracy rate in the U.S was between these two (85.714%). To understand these success rates, we note that the reliability of people in each of the countries echo these results: As shown in Table 2, people in Lebanon exhibited higher reliability than people in Israel (92% versus 65%), and people in the U.S. exhibited higher reliability than people in Israel (65% versus 55%).

We concluded that because the reliability of people in Lebanon was significantly higher than the reliability in Israel and the U.S., it was easier to learn the extent to which people in Lebanon were reliable than the U.S. or Israel. In contrast to Lebanon, the success rate of the reliability for the U.S. population and Israel are similar. Their success rate was nearly the same (84.127% and 85.7143% correspondingly). The main difference in their behavior, as shown in Table 2 was that in the U.S., people were more likely not to deliver any of their chips (and entirely renege on their commitments) than in Israel (61% versus 41%).

With respect to the acceptance model, the highest success rate of predicting the likelihood of acceptance was in Lebanon (75%). Interestingly, people in Lebanon were less likely to agree to offers than in the U.S. or Israel. However, when they accept or reject an offer, they were more likely to fulfill the agreement. To predict acceptance, the optimal features included the scores and generosity features for players. The success rate to predict acceptance in the U.S. was a bit less, 73.5016%. In addition, in the U.S, removing any parameters from the feature set caused the success rate to be reduced. For example, the round number parameter was important. The success rate to predict acceptance in Israel was 70%. To predict acceptance, the optimal features included the scores and offer generosity parameters.

The Israel population accept 56% from the offers. From this, 35% were mistakenly classified, means, classified as FALSE, though their real class was TRUE. As we can see in this acceptance model's success rate table, when learning the model on the combined populations, Lebanon and U.S. the success rate is significantly reduced, only 65.8849%, as compared to learning them separately. This shows that learning from combined population is not good enough as learning each population specifically. To explain this, we note that there was a difference in the likelihood of acceptance in both countries. In the U.S., people were more likely to accept offers than in Lebanon, and we expect that the difference in this behavior made it more difficult to learn from the combined populations.

We also note that the games in the U.S. were generally longer than in other countries. The round number in Lebanon games was up to six, whereas in U.S, the games last up to 16 rounds, which is significantly longer.

The following explains results with respect to the likelihood for reaching the goal. Human Reached Goal model:

- For the Israeli model, many features in the parameters list were needed to get the high success rate of 94.71%. In comparison to Lebanon, where only one parameter (CS) was needed to get high success rate of 93.8776%. On the other hand, for both the U.S and the combined populations, only two parameters of the model were taken: CS and PR. In these cases, adding any other parameter from the basic model, lead the success rate to be reduced.

- In Lebanon, when they accept an offer, they really fulfill the agreement, and therefore, they quickly reached the goal.

- In U.S., the best success rate, 81.0345% was with two parameters in the model, CS and PR. The false positive was 18 instances from 399, i.e 4.5%, and the false negative was 92 instances from 181, i.e, 50%.

- When learning both Lebanon and U.S. population, the parameters needed were CS and PR, and the success rate was 85.2654%, which is much less when learning

Table 4: Optimal Success Rate (in percentages)

Model / Population	Israel	Lebanon	U.S.	Lebanon and U.S.
Reliability	84.127 (PUK)	97.5 (PUK)	85.7143 (PUK)	87.3188 (PUK)
Acceptance	70.3883 (KNN:k=19)	75(J48)	73.5016 (KNN; k=9)	65.8849 (KNN; k=9)
Human reached Goal	94.7103 (J48)	93.8776 (J48)	81.0345 (J48)	85.2654 (J48)
Agent reached Goal	87.1537 (J48)	87.172 (J48)	86.8966 (J48)	86.9989 (J48)

Lebanon separately (93.8776%), and a bit better than learning U.S. separately (81.0345%). Again, this shows that learning each population separately gives better result than learning population together.

Agent Reached Goal model:

- In case of the PURB agent reached goal, it can be seen that in all kind of populations, and even in combined populations, the agent reached the goal in 87%. In addition, for each population learning, this success rate was generated with the 3 parameters, CS, PR and R.

As shown in Table Models Average likelihood, People in Lebanon reached the goal significantly more often than the PURB agent(95% of the time versus 87% of the time), while in the U.S. the PURB agent reached the goal significantly more often than people (95% of the time versus 87% of the time).

To conclude this section, we showed that there were significant differences in the negotiation behavior of people across the three countries. We have shown that these differences affect the performance of learning models in two ways: the types of features that they use to achieve best performance, and the extent to which they can predict behavior in each country. As we have noted in the beginning of this section, the ethnicity demographics of the subjects in each country were different, while other demographic components, such as age and gender, were equal. Therefore we concluded that the differences in behavior can be accounted to the cultural discrepancies between subjects.

7. CONCLUSIONS

This paper investigated the use of machine learning models to predict decisions relating to people's negotiation behavior across 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 were 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.

Our results showed that models that learned to negotiate from using data from particular countries were able to as well as, or better than models that learned from combining the data for all countries. We also showed that there were considerable differences in people's negotiation behavior across the different countries, and that these differences affected the prediction accuracy of the models. We are currently adapting these learning models towards building a cohesive agent that will be able to negotiate successfully with people across cultures.

8. ACKNOWLEDGEMENTS

This research is based upon work supported in part by the U.S. Army Research Laboratory and the U.S. Army Research Office under grant number W911NF-08-1-0144 and under NSF grant 0705587.
Thanks very much to Yael Ejgenberg for her help in programming the PURB agent, and to Swapna Reddy, Leandra King and Konstantin Pozin for their valuable assistance in administering experiments.

References

AAMAS, 2006.

A. Byde, M. Yearworth, K.Y. Chen, C. Bartolini, and N. Vulkan. Autona: A system for automated multiple 1-1 negotiation. In *Proceedings of the 4th ACM conference on Electronic commerce*, pages 198–199. ACM New York, NY, USA, 2003.

C.K.W. De Dreu and P.A.M. Van Lange. The impact of social value orientations on negotiator cognition and behavior. *Personality and Social Psychology Bulletin*, 21: 1178–1188, 1995.

S. G. Ficici and A. Pfeffer. Simultaneously modeling humans' preferences and their beliefs about others' preferences. In *Proc. 7th International Joint Conference on Autonomous Agents and Multi-agent Systems (AAMAS)*, 2008.

Y. Gal and A. Pfeffer. Modeling reciprocity in human bilateral negotiation. In *Proc. 22nd National Conference on Artificial Intelligence (AAAI)*, 2007.

Y. Gal, B. Grosz, S. Shieber, A. Pfeffer, and A. Allain. The effects of task contexts on the decision-making of people and computers. In *Proce. of the Sixth International Interdisciplinary Conference on Modeling and Using Context, Roskilde University, Denmark*, 2007.

Y. Gal, S. DâĂŹsouza, P. Pasquier, I. Rahwan, and S. Abdallah. The effects of goal revelation on computer-mediated negotiation. In *Proceedings of the Annual meeting of the Cognitive Science Society (CogSci), Amsterdam, The Netherlands*, 2009.

M. Gelfand and N. Dyer. A cultural perspective on negotiation: Progress, pitfalls, and prospects. *Applied Psychology*, 49(1):62–99, 2001.

M.J. Gelfand, M. Higgins, L.H. Nishii, J.L. Raver, A. Dominguez, F. Murakami, S. Yamaguchi, and M. Toyama. Culture and egocentric perceptions of fairness in conflict and negotiation. *Journal of Applied Psychology*, 87(5):833–845, 2002.

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 Autonomous Agents and Multi-agent Systems (AAMAS)*, 2004.

C.M. Jonker, V. Robu, and J. Treur. An agent architecture for multi-attribute negotiation using incomplete preference information. *Autonomous Agents and Multi-Agent Systems*, 15(2):221–252, 2007.

E. Kamar, E. Horvitz, and C. Meek. Mobile Opportunistic Commerce: Mechanisms, Architecture, and Application. *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, page 1087, 2008.

R. Katz and S. Kraus. 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)* aam (2006).

P. Kenny, A. Hartholt, J. Gratch, W. Swartout, D. Traum, S. Marsella, D. Piepol, and CA Marina Del Rey. Building interactive virtual humans for training environments. In *Proceedings of I/ITSEC*. NTSA, 2007.

S. Kraus and D. J. Lehmann. Designing and building a negotiating automated agent. *Computational Intelligence*, 11:132–171, 1995.

S. Kraus, P. Hoz-Weiss, J. Wilkenfeld, D.R. Andersen, and A. Pate. Resolving crises through automated bilateral negotiations. *Artificial Intelligence*, 172(1):1–18, 2008.

R. Lin and S. Kraus. Can automated agents proficiently negotiate with humans? *Commun. CACM*, 53(1):78–88, 2010.

R. Lin, S. Kraus, J. Wilkenfeld, and J. Barry. Negotiating with bounded rational agents in environments with incomplete information using an automated agent. *AIJ*, 172(6-7):823–851, 2008.

R. Lin, Y. Oshrat, and S. Kraus. Investigating the benefits of automated negotiations in enhancing people's negotiation skills. In *Proc of AAMAS-09*, pages 345–352, 2009.

R. Murphy. HumanâĂŞrobot interaction in rescue robotics. *IEEE Transactions On Systems, Man, and CyberneticsâĂŤPart C: Applications and Reviews*, 34(2), May 2004.

Y. Oshrat, R. Lin, and S. Kraus. Facing the challenge of human-agent negotiations via effective general opponent modeling. In *AAMAS*, pages 377–384, 2009.

D. Rajarshi, J. E. Hanson, J. O. Kephart, and G. Tesauro. Agent-human interactions in the continuous double auction. In *Proc. 17th International Joint Conference on Artificial Intelligence (IJCAI)*, 2001.

N. Schurr, P. Patil, F. Pighin, and M. Tambe. Using multiagent teams to improve the training of incident commanders. In *Proc. 5th International Joint Conference on Autonomous Agents and Multi-agent Systems (AAMAS)* aam (2006).