A Cultural Sensitive Agent for Human-Computer Negotiation

Galit Haim
Computer Science Department
Bar Ilan University
Ramat Gan, Israel
haimga@macs.biu.ac.il

ABSTRACT

People's cultural background has been shown to affect the way they reach agreements in negotiation and how they fulfil these agreements. This paper presents a novel agent design for negotiating with people from different cultures. Our setting involved an alternating-offer protocol that allowed parties to choose the extent to which they kept each of their agreements during the negotiation. A challenge to designing agents for such setting is to predict how people reciprocate their actions over time despite the scarcity of prior data of their behavior across different cultures. Our methodology addresses this challenge by combining a decision theoretic model with classical machine learning techniques to predict how people respond to offers, and the extent to which they fulfill agreements. The agent based its initial strategy on a general model of the population in each culture, and adapted its behavior to its particular partner over time. This agent was evaluated empirically by playing with 157 people in three countries—Lebanon, the U.S., and Israel—in which people are known to vary widely in their negotiation behavior. The agent was able to outperform people in all countries under conditions that varied how parties depended on each other at the onset of the negotiation. This is the first work to show that a computer agent can learn to outperform people when negotiating in three countries representing different cultures.

1. INTRODUCTION

The dissemination of technology across geographical and ethnic borders is opening up opportunities for computer agents to negotiate with people of diverse cultures and backgrounds. For example, electronic commerce (e.g., ebay), crowd-sourcing (e.g., Amazon Turk) and deal-of-the-day applications (e.g., Groupon) already involve computer agents that make decisions together with people from different countries. People's cultural background has been shown to be a key determinant of the way they make and keep their agreements with others [3]. It is thus important for agent designers to model how people from various cultures respond to different kinds of decision-making behavior employed by others. To succeed in such settings computer agents need to adapt to the culture and particular behavior of the individual they interact with.

Appears in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Conitzer, Winikoff, Padgham, and van der Hoek (eds.), June, 4–8, 2012, Valencia, Spain.

This paper presents a novel agent-design for settings in which participants repeatedly negotiate over the exchange of scarce resources, and agreements are not binding. Such settings characterize the real-world applications shown above, in that participants make commitments to purchasing items or carrying out tasks, they can choose whether and how to fulfill these commitments, and these decisions affect their future interactions with the other participants. For example, a seller that delivers an item very late to a buyer, or does not deliver an item at all, may be negatively reciprocated by the buyer in a future transaction.

2. IMPLEMENTATION USING COLORED TRAILS

Our empirical setting consisted of a game that interleaved negotiation to reach agreements and decisions of whether and how much to fulfil the agreement. The game was configured using the Colored Trails (CT) game [1] and played on a 7x5 board of colored squares. 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.

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, using an alternating offers protocol. 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.

An advantage of using CT 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 that time.

3. LEARNING PEOPLE'S BEHAVIOR

We constructed probabilistic models of people's behavior from data collected in the game as follows: We defined a set of features representing aspects of the game as well as players' reliability measures. We trained classifiers for predicting people's behavior using the subset of features that performed well on a held-out set of data instances, and maximized the

	Average	
PAL	People	
Leb.	132.6	94.86
U.S.	**192.6**	75.77
Isr.	152.75	97.85

Table 1: Performance comparison for each condition and country

likelihood of the training set. These classifiers were incorporated into the influence diagram described in the last section and used by PAL to adapt to people's negotiation behavior in each country.

To meet these challenges we used three sources of data to train our classifiers. First, we used the 222 game instances consisting of people playing the hand-designed agent used by Gal et al. [2]. In addition, in the U.S. and in Israel, we *were* also able to collect 112 game instances of people playing other people. Lastly, in Lebanon, we collected 64 additional games in which people played a variant of the agent used by Gal et al. that was programmed to be significantly less reliable when fulfilling its agreement. In this way, we were able to collect data of people's reactions to more diverse negotiation behavior in the game.

We trained multi-layered neural network classifiers to implement the various models described above using the WEKA framework.[1] We selected the features for each learning task based on their performance (measured by mean-square classification error) on a held-out set of instances as well as measuring the likelihood of the models on the training set.

The variance in people's reliability measures in Lebanon was lower than their variance in the U.S. and in Israel. This explains why it was easier to learn models of these behaviors in Lebanon. Interestingly, the best performance for predicting the likelihood that players got to the goal was obtained in the U.S. This is explained by the fact that the games in the U.S. were significantly longer than the games played in Lebanon and in Israel, making for more opportunities for players to get to goal.

4. EMPIRICAL METHODOLOGY AND RESULTS

To make decisions, PAL used influence diagram , together with the machine learning models and the training data described in Section 3. To evaluate PAL's performance when playing against new people in the game, we recruited 157 subjects from the three countries. All results reported to be significant have been tested for significance in the $p < 0.05$ range using statistical ANOVA tests.

4.1 Comparison of Performance

Table 1 reports performance (in average score per game) for each of the countries and for each dependency condition. As shown by the Table, PAL was able to outperform people in all dependency conditions and in all countries: On average, PAL achieved 192.6 points in the U.S. (right-hand column in boldface), compared to 75.77 points for people; 132.6 points in Lebanon, compared to 94.86 points for people; and 152.75 points in Israel, compared to 97.85 points for people. As shown in Figure 1, PAL was also able to reach the goal significantly more often than people in all dependency conditions and in all countries. The best performance for PAL and the worst performance for people occurred in the U.S: As Table 1 shows, PAL's average performance in the U.S. (192.6 points) was significantly higher than its performance in Lebanon (152.75 points) and Israel (132.6 points points), while people's average performance in the U.S. (75.77 points) was significantly lower than in Lebanon (94.86 points) and Israel (97.85 points).

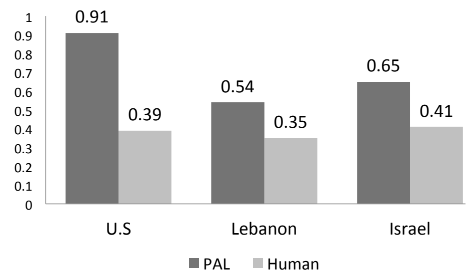

Figure 1: Getting to the Goal (in percentage of games)

5. CONCLUSIONS AND FUTURE WORK

This paper proposed a novel agent design for human-computer negotiation in different cultures. It focused on settings where participants engage in repeated rounds of negotiation and agreements are not binding. To succeed in such settings agents need to reason about the effects of their negotiation behavior over time, and to adapt to people's reaction to their behavior in different cultures. The proposed agent design combined a decision theoretic approach with classical machine learning techniques to model people's behavior. This agent was evaluated empirically by playing with 157 people in three countries—Lebanon, the U.S., and Israel. The results show that the agent was able to outperform people in all countries and when varying how parties depended on each other in the negotiations. The agent based its initial strategy on a general model of the population in each culture, and adapted its behavior to its particular partner over time. We are currently pursuing two future directions. First, we are investigating the use of Markov Chain Monte Carlo sampling techniques for more efficient inference in the game. Second, we are extending our setting to include groups of more than two players and a more advanced negotiation protocol for dividing up tasks among group members.

6. REFERENCES

[1] Y. Gal, B. Grosz, S. Kraus, A. Pfeffer, and S. Shieber. Agent decision-making in open mixed networks. *Artificial Intelligence*, 174(18):1460–1480, 2010.

[2] Y. Gal, S. Kraus, M. J. Gelfand, H. Khashan, and E. Salmon. Negotiating with people across cultures using an adaptive agent. *ACM Transactions on Intelligent Systems and Technology*, 3(1), 2012.

[3] M. J. Gelfand and S. Christakopoulou. Culture and negotiator cognition: Judgment accuracy and negotiation processes in individualistic and collectivistic cultures. *OrganizationŒăal Behavior and Human Decision Processes*, 79(3):248–269, 1999.

[1] http://www.cs.waikato.ac.nz/ml/weka/.