An Agent Design for Repeated Negotiation and Information Revelation with People

(Extended Abstract)

Noam Peled
Gonda Brain Research Center
Bar Ilan University, Israel

Ya'akov (Kobi) Gal
Dept. of Information Systems Engineering
Ben-Gurion University, Israel

Sarit Kraus *
Dept. of Computer Science
Bar Ilan University, Israel

ABSTRACT

Many negotiations in the real world are characterized by incomplete information, and participants' success depends on their ability to reveal information in a way that facilitates agreement without compromising the individual gains of agents. This paper presents a novel agent design for repeated negotiation in incomplete information settings that learns to reveal information strategically during the negotiation process. The agent used classical machine learning techniques to predict how people make and respond to offers during the negotiation, how they reveal information and their response to potential revelation actions by the agent. The agent was evaluated empirically in an extensive empirical study spanning hundreds of human subjects. Results show that the agent was able (1) to make offers that were beneficial to people while not compromising its own benefit; (2) to incrementally reveal information to people in a way that increased its expected performance. The agent also had a positive effect on people's strategy, in that people playing the agent performed significantly higher than people playing other people. This work demonstrates the efficacy of combining machine learning with opponent modeling techniques towards the design of computer agents for negotiating with people in settings of incomplete information.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]

General Terms

Experimentation

Keywords

Human-robot/agent interaction, Negotiation

1. INTRODUCTION

In many negotiation settings, participants lack information about each other's resources and preferences, often hindering their ability to reach beneficial agreements. In such cases, participants can choose whether and how much information to reveal about their resources to others. This paper presents a novel agent design for repeated negotiation with people in settings where participants can choose to reveal information while engaging in a finite sequences of alternating negotiation rounds. Our study is conducted in an experimental framework called a "revelation game" in which people and agents repeatedly negotiate over scarce resources, there is incomplete information about their resources and preferences and they are given the opportunity to reveal this information in a controlled fashion during the negotiation. The proposed agent combines a prediction model of people's behavior in the game with a decision-theoretic approach to make optimal decisions. The parameters of this model were estimated from data consisting of human play. The agent was evaluated in an extensive empirical study that spanned hundreds of subjects. The results showed that the agent was able to outperform human players. In particular, it learned (1) to make offers that were significantly more beneficial to people than the offers made by other people while not compromising its own benefit, and increased the social welfare of both participants as compared to people; (2) to incrementally reveal information to people in a way that increased its expected performance. Moreover, the agent had a positive effect on people's strategy, in that people playing the agent performed significantly higher than people playing other people. Lastly, we show how to generalize the agent-design to different settings that varied rules and situational parameters of the game without the need to accumulate new data. Work in interest-based negotiation has studied different protocols that allows players to reveal their goals in negotiation in a controlled fashion [2, 7, 3]. Other works employed Bayesian techniques [4] or approximation heuristics [5] to estimate people's preferences in negotiation and integrated this model with a pre-defined concession strategy to make offers. Bench-Capon [1] provide an argumentation based mechanism for explaining human behavior in the ultimatum game. We extend these works in two ways, first in developing a strategic model of people's negotiation behavior and second in formalizing an optimal decision-making paradigm for agents using this model.

2. REPEATED REVELATION GAMES

The "repeated revelation game" is played on a board of colored squares. One square on the board is designated as the players' goal. The goal of the game is to reach the goal

*Sarit Kraus is also affiliated with the University of Maryland Institute for Advanced Computer Studies.

Appears in: *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013)*, Ito, Jonker, Gini, and Shehory (eds.), May, 6–10, 2013, Saint Paul, Minnesota, USA.

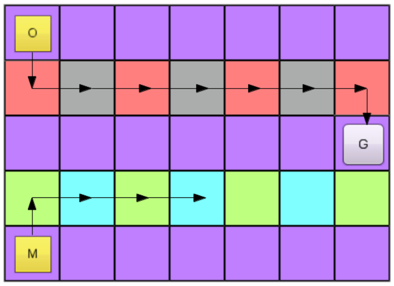

Figure 1: A snapshot showing the revelation game from the point of view of a person (the "M" player) playing against a computer agent (the "O" player).

square. To move to an adjacent square requires surrendering a chip in the color of that square. Each player starts the game with a set of 16 chips. The allocation of the chips was chosen such that no player can reach the goal using only his chips, but there are some chip exchanges that let both players reach the goal. Players have full view of the board, but cannot observe the other player's chips. An example of a CT revelation game is shown in Figure 1. The gray icon "G" represent the goal, and the "M" and "O" icons represent the players starting positions. Each round in our CT game progresses in three phases with associated time limits. In the first "revelation" phase, both players can choose to reveal a subset of their chips.The revelation decision is truthful, that is, players cannot reveal chips that are not in their possession. In the "proposal phase", one of the players can offer to exchange a (possibly empty) subset of its chips with a (possibly empty) subset of the chips of the other player. Following an accepted proposal, the chips are transferred automatically. If the responder rejects the proposal (or no offer was received following a three minute deadline), it will be able to make a counter-proposal. In the "movement phase", the players can move towards the goal using the chips they have. In the next round the players' roles are switched: the first proposer in the previous round becomes the responder for the first proposal. The game ends after both players reach the goal, or after 5 rounds. At the end of the game, both players are moved towards the goal according to their chips, and their score is computed as follows: 60 points bonus for reaching the goal; 5 points for each chip left in a player's possession and 10 points deducted for any square in the path between the players' final position and the goal-square. This path is computed by the Manhattan distance. Note that players' motivation in the game is to maximize their score, not to beat the other participants.

The agent designed for the study, termed MERN (Maximal Expectation-based Revelation and Negotiation agent) developed uses a decision-theoretic approach to negotiate in revelation games. It is based on a model of how humans make decisions in the game. MERN makes decisions in the game by using Expectimax search. Due to the large action space and the exponential increase in the size of the tree, spanning the entire game is not feasible. A key challenge to designing strategies for MERN is how to assign values to intermediate states in the game. To address this challenge MERN uses a heuristic value function to assign utilities to intermediate rounds of the game. The value function is an estimate of the score that MERN will receive at the end of the game. MERN uses a model of human behavior that defines equivalence classes over proposals and revelations in the game and uses machine learning to estimate the parameters of the model. We recruited 410 subjects (186 played with each other, 185 with MERN and 39 with an expert designed agent) using Amazon Mechanical Turk [6]. Each participant played only one game.

The results revealed two opposing patterns in MERN's negotiation strategy. The first pattern included cooperative behavior: MERN learned to make proposals that were significantly more helpful to people (83.52 ± 31.17, $n = 44$), than proposals made by people to other people (37.52 ± 43.49, $n = 204$). The second pattern was competitive: The proposals made by MERN to other people were significantly more competitive (24.89 ± 37.44) than proposals made by people to other people (10.47 ± 57.11). Playing this "hard-headed" strategy affected people's behavior, in that people's average acceptance rate (per proposal) when interacting with MERN (19%) was significantly lower than when interacting with other people (36%). However, for those games in which MERN rejected peoples' proposals in rounds 1-4, peoples' average acceptance rate per game for the competitive offers made in those rounds was very high (88%). Thus, MERN was able to learn to make competitive offers to people that were eventually likely to be accepted. This strategy also affected the efficiency of the offers made by MERN in the game in that 78% of MERN's proposals were pareto optimal, while only 17% of peoples' proposals.

3. ACKNOWLEDGMENTS

S. Kraus and N. Peled were supported in part by ERC grant #267523, the Google Inter-university center for Electronic Markets and Auctions, ARO grants W911NF0910206 and W911NF1110344 and MURI grant number W911NF-08-1-0144. Y. Gal was supported in part by Marie Curie grant #268362.

4. REFERENCES

[1] T. Bench-Capon, K. Atkinson, and P. McBurney. Altruism and agents: an argumentation based approach to designing agent decision mechanisms.

[2] S. G. Ficici and A. Pfeffer. Simultaneously modeling humans' preferences and their beliefs about others' preferences. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 1*, pages 323–330, 2008.

[3] Y. Gal, S. D'souza, P. Pasquier, I. Rahwan, and S. Abdallah. The effects of goal revelation on computer-mediated negotiation. In *CogSci*, 2009.

[4] K. Hindriks and D. Tykhonov. Opponent modelling in automated multi-issue negotiation using bayesian learning. In *Proc. of AAMAS*, pages 331–338, 2008.

[5] 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.

[6] A. Kittur, E. H. Chi, and B. Suh. Crowdsourcing user studies with Mechanical Turk. page 453. ACM Press, 2008.

[7] P. Pasquier, R. Hollands, F. Dignum, I. Rahwan, and L. Sonenberg. An empirical study of interest-based negotiation. In *Proc. of ICEC*, 2007.