An Automated Negotiator for an International Crisis

Penina Hoz-Weiss,[1] Sarit Kraus,[1,3] Jonathan Wilkenfeld,[2,3] and Tara E. Santmire[2]

1Dept. of Computer Science
Bar-Ilan University, Ramat Gan 52900

2Dept. of Government and Politics 3 Institute for advanced Computer Studies
Univ. of Maryland, College Park MD 20742

sarit@cs.biu.ac.il

Introduction

This demo presents an automated agent that negotiates efficiently with human players in a simulated bilateral international crisis. The agent negotiates in a situation characterized by time constraints, deadlines, full information, and the possibility of opting out of negotiation. The specific scenario that we focus on in this demo concerns a crisis between Spain and Canada over access to a fishery in the North Atlantic. Canada blames Spain for over-fishing near its territorial waters and thereby damaging the flatfish stock. The countries have agreed to meet and negotiate over the fishery dispute. If an agreement is not reached by the beginning of the next fishing season, a status quo outcome will be implemented. The status quo outcome is not equally advantageous to both parties.

We developed an automated agent that can play the role of either side in such negotiations. The negotiation is conducted using a semi-formal negotiation language. The language consists of seven types of messages, including detailed offers and counteroffers, threats and promises. The human players are provided with a decision support system that helps them to analyze the scenario and to compare the utility points associated with various outcomes. They are also provided with a language editor to facilitate the composition of messages during the negotiation. The model used in constructing the automated agent is based on a formal analysis of the fishing dispute scenario using game theoretic methods and heuristics for argumentation.

The Fishing Dispute Scenario

Canada and Spain have agreed to meet in an attempt to negotiate an agreement regarding the fishery dispute. Each party must consider five possible ways of ending the crisis: (1) An agreement on Total Allowable Catch (TAC) for the season. The TAC can be between 1 ton and 54 tons. (2) An agreement on limiting the length of the fishing season. (3) Canada enforces conservation measures with military force against Spain. This can result in either success, partial success or failure. (4) Spain enforces its right to fish throughout the fishery with military force against Canada. This can result in either success, partial success or failure. (5) Status quo. The following are world state parameters that are also negotiable and affect the players' utilities: (a) Canada subsidizes the removal of Spain's ships (0, 5, 10, 15, 20 ships). (b) Spain reduces the amount of pollution caused by the fishing fleet (0%, 15%, 25%, 50%). (c) Canada imposes trade sanctions on Spain. (d) Spain imposes trade sanctions on Canada. The negotiation takes time and is divided into time periods. If the negotiation does not end by the beginning of the fishing season, then the status quo will be implemented. During the negotiation, each of the parties has the capability to make requests, threats, offers, conditional offers and counteroffers, as well as to comment on the negotiation. See (Hoz-Weiss 2001, Hoz-Weiss et. al) for the values and examples of the utility functions of the players in the fishing dispute.

 This research is supported by NSF, grant IRI-9820657.

Agent Design

The automated agent is a program written in FCC that handles a negotiation process. The agent is programmed to play in simulations of the fishing dispute. It can play either side in the process of negotiation. During the simulation the agent receives messages sent by humans, analyzes them and responds. It also initiates a discussion on one or more parameters of the agreement.

The agent is based on a simultaneous negotiation model. At the beginning of the crisis it computes by backward induction the subgame perfect equilibrium (Osborne & Rubinstein 1990, Kraus 2001,Hoz-Weiss 2001). It stores the offers that it should make in each time period according to the equilibrium strategy in an array, referred to as the strategy array.

We demonstrate the way the agent computes the equilibrium strategy in the case where the deadline for the negotiations is 16 time periods and the world state parameters are not taken into consideration. If Canada makes the last offer at period 15, it will offer TAC = 1, since the expected utility for Spain from opting out is 477.1 and the expected utility for Spain if Canada opts out is 379.2. If TAC=1 is accepted Canada will attain 630 and Spain 560. This is better for Spain than opting out and the best possible outcome for Canada. If Spain makes the last offer, it will offer TAC = 39 since Spain would attain 940 and Canada 440, and the expected utility for Canada from opting out is 438.8. Thus, TAC=39 is better for Canada than opting out and the best possible outcome for Spain. When going backward to *t*= 14 each country that will make an offer has to offer an agreement that will provide the other party a utility which is higher than the expected utility from opting out and higher than the expected utility in period 15. Both Canada and Spain, if they make an offer at *t* = 14, will offer TAC = 21. Continuing the backward induction, at the first time period the agreement that will be offered by a country will be TAC = 34 which gives Spain a utility of 750 and Canada a utility of 535. The other country will accept the agreement.

If the agent was playing against a rational opponent, who has the ability to identify the subgame perfect equilibrium, this would be sufficient. In particular when one version of the agent plays against another version, the agreement is reached in the first time period in the array. However, humans do not necessarily follow equilibrium strategies, and when the automated agent follows its equilibrium strategy the human negotiators become frustrated and often the negotiation ends with no agreement. Therefore, the formal theory is not enough and we added heuristics and argumentation to complete the formal model and make the agent an effective negotiator with humans.

There are two main activities that the agent performs during the negotiations: (a) Responding to incoming messages, and sending counteroffers that serve the interests of the automated agent: The specific message depends on the incoming messages, the strategy array, the world state and the agent's parameters that are specified below. The agent maintains the state of the world during the negotiation. (b) Sending messages regarding issues that have not yet been discussed: Every three minutes the agent checks which parameter was not negotiated recently and then it sends a message regarding that parameter. The specific heuristics can be found in (Hoz-Weiss 2001).

As part of the heuristics we used there is a set of parameters that influences the agent's behavior. These parameters are instantiated before the beginning of the negotiations: (i) A parameter that indicates whether the agent sends the first message in the negotiation or waits for the opponent to make the first offer. (ii) A parameter that determines if the agent will use the full offer message or will use partial offers to negotiate each issue separately. (iii) The number of negotiation units (tons of fish in the fishing dispute) the agent will increase or decrease its offer by. (iv) The agent agrees to an agreement that yields a utility that is lower than the desired utility by at most the number of points specified by this parameter.

The agent is sensitive to the risk level of its human opponent and will change its view of the human's utility function accordingly. The agent begins with the assumption that its opponent is risk neutral. We use a heuristic method to decide whether to change the estimation of the risk attitude of the opponent. When the agent decides that its opponent is risk prone, it changes the opponent's utility function. This leads the agent to a recalculation of the strategy array.

Experiments

In order to evaluate the agent's performance in negotiation situations and to compare it to the performance of humans, we conducted simulations with Computer Science students at Bar Ilan University. The students were introduced to the Generalized Decision Support System (GDSS) for the Fishing Dispute simulation, by which they could evaluate different outcomes in terms of utility point values. A total of 45 simulations were run: 15 simulations were human against human, and 30 simulations were human against the agent. In 14 simulations the agent played Spain and in 16 simulations the agent played Canada (Hoz-Weiss 2001).

Comparing the results of the humans to those of the agents for those simulations that ended with an agreement, the agent plays Spain's role significantly better than the human does (agent's average utility: 845; humans' average utility: 723; $t=-5.957$, $p<0.01$) and the role of Canada just as well as a human (agent's average utility: 607; humans' average utility: 612). When looking at the results that include all the outcomes, again, the agent playing Spain played significantly better than the human playing Spain ($t=-2.51$, $p<0.05$). The results for Canada did not show a significant difference between the agent and human players.

In addition, the average sum of utility points in simulations where agreements were reached with only humans was 1336 and the average sum of the simulations where an agent was involved was 1439. We conclude that when an agent participates in a negotiation the sum of the utilities was significantly higher than when two humans played ($t=-4.916$, $p<0.01$).

References

P. Hoz-Weiss. An Automated Negotiator for a Fishing Dispute, Master thesis, Bar-Ilan Univ., Ramat-Gan 2001.

P. Hoz-Weiss, S. Kraus, J. Wilkenfeld and T. E. Santmire. An Automated Agent for Bilateral Negotiations with Humans, in M. Huget editor, Communication in MAS Springer-Verlag. (To appear).

S. Kraus, Strategic Negotiation in Multiagent Environments. MIT Press, 2001.

Osborne, M. J. and A. Rubinstein 1990. Bargaining and Markets, Academic Press.