GPNEG: General Purpose Negotiation Training Tool

Raz Lin and **Sarit Kraus**[*]
Department of Computer Science
Bar-Ilan University
Ramat-Gan, Israel 52900
{linraz, sarit}@cs.biu.ac.il

Jonathan Wilkenfeld [*]
Department of Government and Politics
University of Maryland
College Park, Maryland 20742
jwilkenf@gvpt.umd.edu

Abstract

Negotiation is an important mechanism for resolving conflicts between people. Many tasks in day-to-day life involve interactions among several people. Many of these interactions involve negotiating over a desired outcome. Negotiation in and of itself is not an easy task, and it becomes more complex under conditions of incomplete information. Therefore, training people in negotiation is beneficial, yet difficult. In this paper an innovative general purpose tool for negotiation training is proposed and demonstrated experimentally. The experiments demonstrate the ease of use and the general applicability of the proposed tool. Thus, our work contributes to the automation of negotiation in general, and to negotiation training in particular.

Introduction

Negotiation processes are becoming more and more common, yet in many situations they are conducted by ordinary people who have had no formal training in negotiations. Moreover, the increasing diverse settings in which negotiations occur means that the outcomes of the negotiation process can have unforeseen implications for the negotiators. Despite this, there are few generic automated tools that can offer people an efficient training experience. In addition, the issue of cross cultural negotiation has become increasingly important, and a training tool that addresses this aspect can make the difference between a successful negotiation and a failure.

In this paper we present GPNEG - a general purpose negotiation training tool. The development of GPNEG stems from the current lack of existing training tools. While there are tools in the market for negotiation training, as we describe in the related work section, these tools offer limited services and lack the rich data required for a robust training. GPNEG is an automated negotiation environment. Note that we do not intend it to replace humans in negotiation. GPNEG is intended to be used as an efficient decision support tool or as a training tool for negotiations with people. Thus,

it can be used to support training in real daily life negotiations, such as e-commerce, and it can also be used as the main tool in both conventional lecture and online courses, aimed at making the trainee a better negotiator.

Most negotiation tools today are domain-dependent and focus on a single issue negotiation (e.g., price). These tools do not provide an efficient training and learning experience for the trainee. Instead of providing the trainee a wide range of test cases, they constrain him to a predefined scenario, which is only a fragment of the variety of scenarios he/she might encounter in the real world.

In addition, our automated environment embodies an automated agent that plays against the trainee. This allows the trainee to use it *anytime* and test his/her capabilities and improvement. This agent was proven to play efficiently against other people (Lin *et al.* 2006). In contrast to existing automated agents, our automated agent does not assume that every side has complete information about the negotiation settings and preferences of the rival, or the rationality of both sides. Since most negotiations are done in situations where there is incomplete information about some settings of the negotiations (e.g., uncertainty regarding the preferences of the opponent), and since research has shown (Kraus, Hoz-Weiss, & Wilkenfeld 2007) that humans usually do not fall in the category of rational decision making, other tools fail to provide a realistic training environment.

GPNEG provides a generic environment for test-bedding different negotiation scenarios. Thus, it provides a domain-independent capability which can be used with almost every domain. In addition, it allows for either human-human negotiations or human-computer negotiation. This makes the tool available for use both in the academia and in training companies as the main tool for training people to negotiate more effectively.

The rest of the paper is organized as follows. We start by describing related work in the field of automated negotiation tools. Then we describe GPNEG and its features. We continue and present two case studies done using GPNEG. Finally, we provide a summary and discuss future work.

Related Work

Research on cross culture negotiation has demonstrated the importance of the context in which negotiation is being performed (Cohen 1997). Based on this our tool should be able

---

to model different opponent types and to deal with incomplete information.

The problem of modeling an automated agent for bilateral negotiation is not new for researchers in the fields of Multi-Agent Systems and Game Theory. However, most research makes simplifying assumptions that do not necessarily apply in genuine negotiations, such as assuming complete information (Faratin, Sierra, & Jennings 2002; Oprea 2002) or the rationality of the opponent (Faratin, Sierra, & Jennings 2002; Fatima & Wooldridge 2004; Fatima, Wooldridge, & Jennings 2005; Lin & Chou 2003). None of the above researchers has looked into the negotiation process in which there is both incomplete information and the opponent is bounded rational (for example, humans). While their approaches might be appropriate in their context, they cannot be applied to our settings.

Dealing only with the bounded rationality of the opponent several researchers suggested new notions of equilibria (e.g., the *trembling hand equilibrium* described in Rasmusen (Rasmusen 2001) (p. 139)) or other probability models. For example, Capra *et al.* (Capra *et al.* 1999) use what is called a *"standard logit model"*. In this model probabilities are assigned to the decisions. Those probabilities are proportional to exponential functions of expected payoffs. They use this model in order to enable the players to update their beliefs about other players. This model is equivalent to assuming that expected payoffs are subjected to deviations with an extreme value distribution. That is, the logit model assumes that the decisions are not perfect and may have some noise and it tries to deal with such situations. These errors can be interpreted either as unobserved random changes in preferences or as errors in responding to expected payoffs. Similar to Capra *et al.*, our agent also assigns probability to the believed type of the opponent. However, we try to avoid the need of adding a special mechanism that assumes that the actions of the opponent are characterized by noise.

Other researchers suggested shifting from quantitative decision theory to qualitative decision theory (Tennenholtz 1996). In using such a model we do not necessarily assume that the opponent will follow the equilibrium strategy or try to be a utility maximizer. Also, this model is better suited for cases in which the utility or preferences are unknown but can be expressed in ordinal scales or as preference relations (Dubois, Prade, & Sabbadin 2001). This approach seems appropriate in our settings, and using the maximin criteria, which is generally used in this context, enables our agent to follow a pessimistic approach regarding the probability that an offer will be accepted.

Several methods are proposed when dealing with incomplete information regarding the preferences of an opponent. For example, Bayes' theorem is the core component of the *Bayesian Nash equilibrium* ((Osborne & Rubinstein 1994), p. 24-29), and it is used to deduce the current state given a certain signal. One motivation for using this notion of equilibrium is that it allows one to compensate for incomplete information and enables a good adaptation in a negotiation with time-constraints. In finite horizon negotiations there are no past interactions to learn from and not enough time periods to build a complete model. Thus this model provides a good probabilistic tool to model the opponent, as opposed to using feed-forward neural networks (Oprea 2002) or genetic algorithms (Lin & Chou 2003), both of which require considerable time to facilitate adequate learning and are more sensitive to the domain in which they are run.

Zeng and Sycara (Zeng & Sycara 1998) also build on using Bayesian analysis as a learning mechanism in negotiations. Like them, we also use Bayes' theorem to update the believed type of the opponent. Thus, we allow the negotiator, at each time period, to act as if the opponent is of a certain type.

In the following section we present our general purpose negotiation tool.

GPNEG: A General Purpose Negotiation Training Tool

GPNEG is a general purpose negotiation training tool with a generic nature. It enables negotiation between two people or against one of two automated agents incorporated in the tool. We first describe the tools interface and its adaptive nature and then we describe the automated agents incorporated in the tool.

GPNEG Overview

The GPNEG tool provides a simulation environment which is adaptable such that any scenario and utility functions, expressed as a single issue or multi-issue attributes, can be used, with no additional changes in the configuration of the interface of the simulations or the automated agents. The automated agents can play either role in the negotiation, while the human counterpart accesses the negotiation interface via a web address. Note that there is also an option to load a scenario *without* attaching a utility to each issue, and thus add uncertainty to the negotiation process. The negotiation itself is conducted using a semi-formal language. Each player constructs an offer by choosing the different values constituting the offer. Then, the offer can be sent in plain English to the counterpart. In addition, the tool allows the participants to reach either a partial or full agreement. To make the negotiation richer, in addition to sending proposals to the opponent which upon acceptance are taken as commitments, the players can also send queries and promises. The difference between queries and promises to offers is that they are not binding, and even if accepted, both sides can backtrack from them. In all of these message types, though, a message can be sent regarding all the attributes of the negotiation or only some of them.

Table 1 lists the different parameters adjustable in the simulation environment. These parameters are set prior to loading the simulation, and are described below.

The scenario is inserted by defining the different issues and their attributes for the negotiation. For each issue and attribute an optional description can be given. The utility itself (an outcome or value attached for each attribute) is not mandatory. If no utility is given then it allows for uncertainty to rule the negotiation. If a utility is inserted, then the tool also allows the designer to add uncertainty regarding the exact preferences of each agent. This is done by adding

Table 1: Adaptable parameters in the negotiation environment.

Parameter	Optional (O) or Mandatory (M)
Scenario	M
Utility and agents types	O
Number of turns	M
Turn Length	M
Time effect	O
Opt-out value	O
Status-quo	O
Comments list	O
Threats list	O
Opponent type	M
Instruction	O

Figure 1: Main negotiation screen.

different agent types to the system. That is, different utilities can be loaded and related to different agent types (e.g., an agent that has a long-term orientation vs. an agent with a short-term orientation). The player can then be matched with one of these types.

Another parameter is the number of turns in the negotiation. The simulation tool allows each player to perform any number of interactions with the opponent player at any given time period (as opposed to the model of alternating offers (Osborne & Rubinstein 1994)). The number of turns for the negotiation can be set along with the length of each turn. The time effect is an optional parameter that assigns a time cost which influences the utility of each player as time passes (there can be different time costs for each player). The time effect can be either negative or positive. If no agreement is reached by the end of the final turn then a status quo agreement is implemented resulting in a status quo value to each player. Another option shipped with the tool is the option for each player to quit the negotiation at any given time if he/she decides that the negotiation is not proceeding in a favorable way. This results in the implementation of an opt-out outcome. To enable a rich interaction between the players there is also the option to send comments or threats to the other side throughout the negotiation. The possible values for the comments or threats can be loaded with the simulation environment as well. Finally, the simulation is loaded after setting the opponent type. This can be either two human players or a human player playing against one of the two automated agents shipped with the tool. The automated agents are described in the following section.

During each phase of the negotiation, the instructions and the attributes of the negotiation are accessible to the players. The players are also aware of the current turn and time left until the end of the turn and until the negotiation terminates. The history of past interactions is also easily accessible.

Examples of the main screen, the offer generation and the receiving offer screens are given in Figures 1, 2 and 3, respectively. When receiving an offer the player can choose whether to accept or reject it, or to make a counter-offer. When rejecting an offer, a free text can be sent by the player to the opponent as to the reason for the rejection.
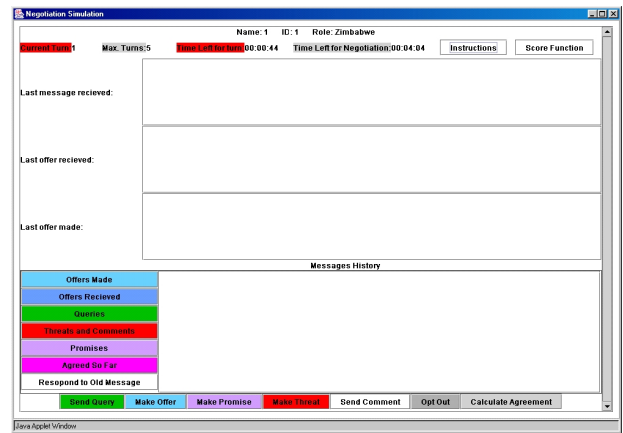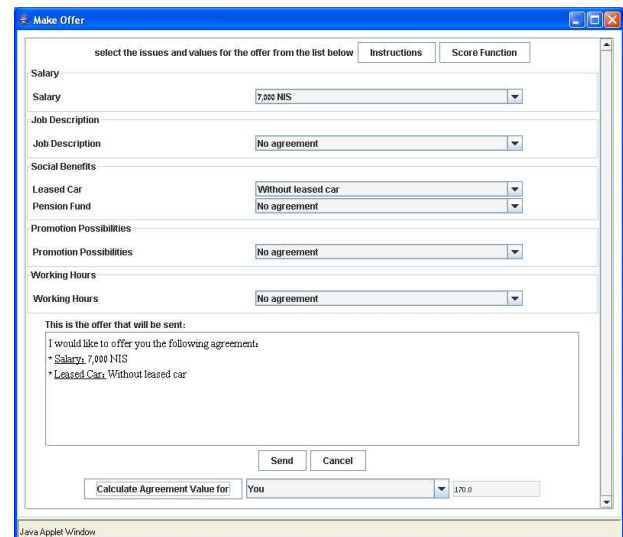

Figure 2: Generating offers screen.

If a utility is attached to each outcome, then the tool allows the player to use an 'outcome calculator'. This calculator enables the player to simulate different outcomes and to check their values for him/her and the opponent. This can be done for each of the agent types that were defined.

The Automated Agents

Two automated agents are delivered with the tool - the first is a simple equilibrium agent, shipped today with many automated tools. But the equilibrium agent yet has several disadvantages when it is matched against humans or if incomplete information is invovled. The second is our proposed automated agent, termed *QOAgent*, which plays efficiently under the assumptions of incomplete information and the bounded rationality of the opponent. The tool also allows for two automated agents to be matched against each other so the trainee can learn from this experience. We describe below the nature and characteristics of the *QOAgent*.

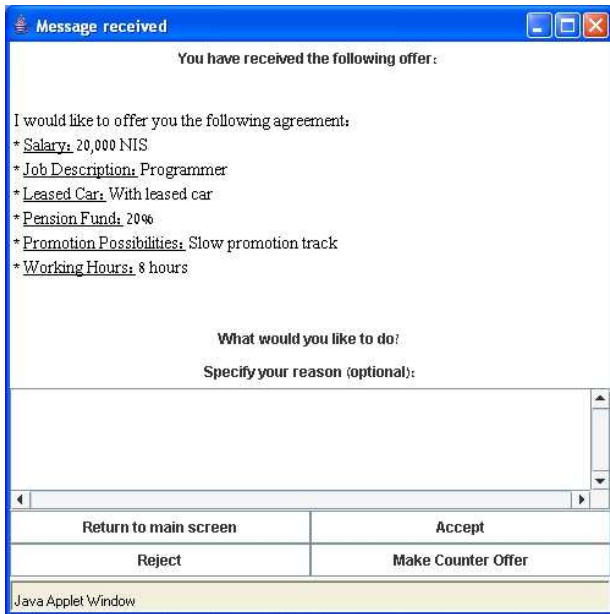The *QOagent* is built with two mechanisms: (a) a deci-

Figure 3: Receiving offers screen.

sion making mechanism, and (b) a mechanism for updating beliefs. While designing the agent we shifted from the traditional quantitative decision making (e.g., the *trembling hand equilibrium* described in Rasmusen (Rasmusen 2001) (p. 139) and the *Bayesian Nash equilibrium* (Osborne & Rubinstein 1994), p. 24-29) to a qualitative decision approach (Tennenholtz 1996). In using such a model we do not necessarily assume that the opponent will follow an equilibrium strategy or try to be a utility maximizer. We start by describing the decision making component, which is responsible both for generating offers and deciding upon acceptance of offers, and then we continue by describing the updating beliefs component.

The motivation behind the mechanism for generating offers is that the automated agent would like to propose an offer which yields him/her the highest utility value. However, due to conflicts of interests, there is a high probability that this agreement will be rejected by the opponent. To overcome this, our agent uses a qualitative decision strategy. Basically, the agent evaluates all possible offers based on its utility and the probability that the rival will accept them.

Since the opponent him/herself also tries to reason whether an offer will be accepted by our agent, we take this into account as well. That is, our agent tries to estimate, from the opponent's point of view, whether the opponent will accept the offer. Using the resultant value, our agent compares it with its own utility values. Similar to the qualitative decision theory, which uses the maximin value (Dubois, Prade, & Sabbadin 2001; Tennenholtz 1996), our agent selects the minimum value between those two values, under the pessimistic assumption that the probability that an offer is accepted is based on the agent that favors the offer the least. After calculating the minimum value between all the offers, our agent selects the offer with the maximum value among all the minima, in order to also try and maximize its own utility. Thus, our qualitative offer generation mechanism selects, intuitively, the best offer among the offers that the agent believes that might be accepted.

Seemingly, our method of generating offers is a non-classical method. However, not only were we able to show its efficacy by empirical experiments, in which it was used in negotiations with bounded rational agents, but also showed that it also conforms to some properties from classical negotiation theory, which are mainly used by mediators. Detailed description of the agent and its qualitative approach can be found in (Lin *et al.* 2006).

Another important issue is when to accept offers sent by the opponent. The agent needs to decide what to do when it receives an offer from its opponent, offer_{opp}, at time $t-1$. If we refer to the automated agent as agent 1 and the opponent player as agent 2, if $u_1(\text{offer}_{opp}) \geq u_1(QO(t))$ then our agent accepts the offer. Otherwise, our agent should not immediately rule out accepting the offer it has just received. Instead, it should take into consideration the probability that its counter-offer will be accepted or rejected by the opponent. This is done by comparing the believed utility of the opponent from the original offer as compared with the opponent's utility from our offer. If the difference is lower than a given threshold T, that is $\mid u_2(QO(t)) - u_2(\text{offer}_{opp}) \mid \leq T$, then there is a high probability that the opponent will be indifferent between its original offer and our counter-offer, so our agent will reject the offer and propose a counter-offer (taking a risk that the offer will be rejected), since the counter-offer has a better utility value for our agent. If the difference is greater than the threshold, i.e., there is a higher probability that the opponent will not accept our counter-offer, our agent will accept the opponent's offer with a given probability, which is attached to each outcome. To this end we define the rank number, which is associated with each offer and a given utility function u, denoted $rank(\text{offer})$. The rank number of an offer is calculated by ordering all offers on an ordinal scale between 1 and $|O|$ according to their utility values, and dividing the offer's ordering number by $|O|$. That is, the agent will accept an offer with a probability $rank(\text{offer}_{opp})$ and reject and make a counter-offer $QO(t)$ with probability $1 - rank(\text{offer}_{opp})$. The intuition behind this is to enable the agent also to accept agreements based on their relative values, on an ordinal scale of $[0..1]$, and not based on their absolute values.

For the update beliefs component we use the Bayesian updating rule, which is based on Bayes' theorem (Osborne & Rubinstein 1994). We basically assert that there is a set of different agent types. The bounded rational agent should be matched to one such type. In each time period, the agent consults the component in order to update its belief regarding the opponent's type. For every offer that is received or response to offers from the opponent the agent updates its beliefs and acts as if the opponent is of a given type (the one with the highest probability).

In the next section we describe the experiments which we ran in order to test the adaptive nature of the simulation tool as well as the efficacy of our proposed automated agent which is shipped with the tool.

Cross Cultural Negotiations Using GPNEG

As we stated above, the issue of cross-cultural negotiation is important. Negotiation over the same issues against parties from different countries, for example, can result with distinct agreements, and it is vital that the negotiator will be aware of these variations. It has been shown that in different countries the attitude regarding the negotiation and the actions during it are quite different. A Chinese negotiator will appear to concede more often while in the UK it is common to use pressure tactics to impose a deal on the other side. The same tactic, however, against a negotiator from Greece will most likely backfire.

Our tool supports the modeling of different negotiation styles for the opponents. This is simply done by modeling different preferences for the opponent. Using it, the negotiator can be matched against any of these types at any given time.

In addition, it is common for people to adopt different negotiation styles in the way they pursue the closing of the deal. For example, one might want to address all the issues in a single agreement, while the others might prefer reaching an agreement incrementally, each time committing to one issue. This process is also facilitated by GPNEG, as it enables each side to propose a full agreement or only a partial agreement. It also allows a party to improve on a previously accepted agreement by negotiating over previously agreed upon issues. However, no backtracking from an agreement is possible, only opting-out.

Case Study

To test the efficiency of the proposed agent, we have conducted experiments in two distinct domains. These experiments show that the agent is capable of negotiating in various domains. That is, once the utility functions are adapted to fit the parameters of a specific scenario, the agent is in a position to interact effectively in the negotiation process. In the following subsections we describe the two domains in which our tool was tested.

In the first domain one player gains as time advances, while the other loses, the status-quo value for one of the players is much higher than for the opponent, and there is an option to reach partial agreements. In the second domain, both players lose as time advances, and the status-quo value for both players is quite similar. In both domains we modeled three possible player types, and thus a set of six different utility functions was created for each domain. These sets describe the different types or approaches towards the negotiation process and the other party. For example, type (a) has a long term orientation regarding the final agreement, type (b) has a short term orientation, and type (c) has a compromise orientation.

Each negotiator was assigned a utility function at the beginning of the negotiation but had incomplete information regarding the opponent's utility. That is, the different possible types of the opponent were public knowledge, but the exact type of the opponent was unknown. The negotiation lasts at most 14 time periods, each with a duration of two minutes. If an agreement is not reached by the deadline then the negotiation terminates with a status quo outcome. Each party can also opt out of the negotiation if it decides that the negotiation is not proceeding in a favorable way.

The first domain was based on an international crisis, while the second domain dealt with negotiation over the hiring terms after a successful job interview. We describe the two domains in the following subsections.

The domains reflect the applicability and gererality of our tool for many other domains. In addition, our results indicate that the automated agent, implemented in the tool, reached more agreements and played more effectively than its human counterparts. Moreover, in most of the cases, it played significantly better than the human counterparts. This reinforced our contention that this is an effective training tool for negotiation purposes. Next we describe the two scenarios.

Domain 1: The World Health Organization's Framework Convention on Tobacco Control In this scenario England and Zimbabwe are negotiating in order to reach an agreement growing out of the World Health Organization's Framework Convention on Tobacco Control, the world's first public health treaty. The principal goal of the convention is "to protect present and future generations from the devastating health, social, environmental and economic consequences of tobacco consumption and exposure to tobacco smoke."

The leaders of both countries are about to meet at a long scheduled summit. They must reach agreement on the following issues:

1. **The total amount to be deposited into the Global Tobacco Fund to aid countries seeking to rid themselves of economic dependence on tobacco production**. This issue has an impact on the budget of England and on the effectiveness of near-term and long-range economic benefits to Zimbabwe. The possible values are (a) $10 billion, (b) $50 billion, (c) $100 billion, or (d) no agreement

2. **Impact on other aid programs**. This issue affects the net cost to England and the overall benefit to Zimbabwe. If other aid programs are reduced, the economic difficulties for Zimbabwe will increase. The possible values are (a) no reduction, (b) reduction equal to half of the Global Tobacco Fund, (c) reduction equal to the size of the Global Tobacco Fund, or (d) no agreement.

3. **Trade issues**. Both countries can use trade policy to extract concessions or provide incentives to the other party. They can use restrictive trade barriers such as tariffs (taxes on imports from the other country) or they can liberalize their trade policy by increasing imports from the other party. There are both benefits and costs to these policies: tariffs may increase revenue in the short run but lead to higher prices for consumers and possible retaliation by affected countries over the long run. Increasing imports can cause problems for domestic industries. But it can also lead to lower consumer costs and improved welfare. Thus, the possible values are divided between Zimbabwe's (a) reducing tariffs or (b) increasing tariffs on imports, and England's (a) reducing or (b) increasing imports. Both can also choose not to agree on this.

4. **Creation of a forum to explore comparable arrangements for other long-term health issues**. This issue relates to the precedent that may be set by the Global Tobacco Fund. If the fund is established, Zimbabwe will be highly motivated to apply the same approach to other global health agreements. This would be very costly to England. The possible values are (a) creation of a fund, (b) creation of a committee to discuss the creation of a fund, (c) creation of a committee to develop an agenda for future discussions, or (d) no agreement

Thus, a total of 576 possible agreements exist. While on the first two issues there are contradicting preferences for England and Zimbabwe, for the last two issues there are options which might be jointly preferred by both sides.

Each turn in the scenario is equivalent to a week of the summit, while the summit is limited to 14 weeks. If no agreement is reached within the specified time limit, the Framework Convention will be seen as an empty document, devoid of any political significance. This will be a blow to England, which has invested political capital to reach an agreement, in the hope of gaining support for other, perhaps more important, international agreements in the future. It will also, however, save it money in the near term. For Zimbabwe, failure to reach an agreement will create a major financial hardship and deprive them of a precedent that can be used for future negotiations. Consequently, England is better able to accept a failure than Zimbabwe is. This outcome is modeled for both players as the status-quo outcome.

Opting out of the negotiation is also an option. Opting out by England means trade sanctions imposed by England on Zimbabwe (including a ban on the import of tobacco from Zimbabwe), while if Zimbabwe opts out then it will boycott all British imports. However, if England opts out it also saves the funds that would have been spent on the Tobacco Fund, and if Zimbabwe opts out it loses the opportunity for financial gain and for assistance in reducing the health problems that arise from tobacco use. Consequently, England will likely be more willing to opt out if the negotiations are not going its way, and Zimbabwe will be more willing to continue negotiations until agreement is reached.

Time also has an impact on the negotiation. Creation of the fund is more urgent for Zimbabwe than for England. Consequently, Zimbabwe has an incentive to reach an agreement earlier rather than later, thus as time advances it loses utility. On the other hand, England gains as time advances, as it postpones the time at which it must transfer money to the fund.

Taking into account the different types of players, we can say, for example, that a player that represents Zimbabwe and has a short term orientation, focuses on short term redistribution of resources, insists on the largest possible current assistance and help with long-term health problems, as well as trade concessions. On the other hand, a player which represents England with the same short term orientation, for example, aims to minimize current cost, limit impact on trade, and maintain economic and political position in the near term.

Domain 2: The Job Candidate   In this scenario, a negotiation takes place after a successful job interview between an employer and a job candidate. In the negotiation both the employer and the job candidate wish to formalize the hiring terms and conditions of the applicant. Both sides must reach agreement on the following issues:

1. **Salary**. This issue dictates the total net salary the applicant will receive per month. The possible values are (a) $7,000, (b) $12,000, or (c) $20,000.

2. **Job description**. This issue describes the job description and responsibilities given to the job applicant. The job description has an effect on the advancement of the candidate in the work place and his/her prestige. The possible values are (a) QA, (b) programmer, (c) team manager, or (d) project manager.

3. **Social benefits**. The social benefits are an addition to the salary and thus impose an extra expense to the employer, yet they can be viewed as an incentive for the applicant. The social benefits are divided into two categories: company car and the percentage of the salary allocated, by the employer, to the candidate's pension funds. The possible values for a company car are (a) giving a leased company car, (b) no leased car, or (c) no agreement. The possible value for the percentage of the salary deposited in pension funds are (a) 0%, (b) 10%, (c) 20%, or (d) no agreement.

4. **Promotion possibilities**. This issue describes the commitment by the employer regarding the fast track for promotion for the job candidate. The possible values are (a) fast promotion track (2 years), (b) slow promotion track (4 years), or (c) no agreement.

5. **Working hours**. This issue describes the number of working hours required by the employee per day (not including over-time). This is an integral part of the contract. The possible values are (a) 8 hours, (b) 9 hours, or (c) 10 hours.

In this scenario, a total of 1,296 possible agreements exist.

Each turn in the scenario equates to two minutes of the negotiation, and the negotiation is limited to 28 minutes. If the sides do not reach an agreement by the end of the allocated time, the job interview ends with the candidate being hired with a standard contract, which cannot be renegotiated during the first year. This outcome is modeled for both players as the status-quo outcome.

Each side can also opt-out of the negotiation if it feels that the prospects of reaching an agreement with the opponent are getting slim and it is impossible to negotiation any more. Opting out by the employer entails the postponement of the project the candidate was interviewing for, with the possible prospect of its cancellation and a considerable amount of expense.

Opting-out by the job candidate will make it very difficult for him/her to find another job, as the employer will spread his/her negative impression of the candidate to other CEOs of large companies.

Time also has an impact on the negotiation. As time advances the candidate loses utility, as the good impression of

the employer has of the job candidate decreases. The employer also loses utility as the candidate becomes less motivated to work for the company.

Experiment Results

We evaluated the performance of the agent against human subjects, all of whom were computer science undergraduates at Bar-Ilan University in Israel. The experiment involved 88 simulations with human subjects, divided into 44 pairs, such that 44 simulations were run for each domain. Each simulation was divided into two parts: (i) negotiating against another human subject, and (ii) negotiating against the automated agent. The subjects did not know in advance against whom they played. Also, in order not to bias the results as a consequence of the subjects getting familiar with the domain and the simulation, for exactly half of the subjects the first part of the simulation consisted of negotiating with a human opponent, while the other half negotiated first with the automated agent. The outcome of each negotiation is either reaching a full agreement, opting out, or reaching the deadline without an agreement. Prior to the experiments, the subjects were given oral instructions regarding the experiment and the domain. The subjects were instructed to play based on their score functions and to achieve the best possible agreement for them.

The main goal of the experiments was to verify that the automated agent is capable of achieving better agreements than a human playing the same role, and to facilitate an earlier end to the negotiation as compared to negotiations without the agent. A secondary goal was to check on whether indeed the agent facilitated an increase in the social welfare of the outcome, that is, improved the utility scores for both parties, as compared to negotiations without an automated agent.

Results for the England-Zimbabwe Domain Table 2 summarizes the average utility values of all the negotiations, the average ranking of the agreements reached, and the average of the sums of utility values and ranking of the agreements in all the experiments in the England-Zimbabwe domain. H_{Zim} and H_{Eng} denote the utility value gained by people playing the role of Zimbabwe or England, respectively, and Q_{Zim} and Q_{Eng} denote the utility value gained by the QO agent playing either role.

The utility values ranged from -575 to 895 for the England role and from -680 to 830 for the Zimbabwe role. The Status-Quo value in the beginning of the negotiation was 150 for England and -610 for Zimbabwe. England had a fixed gain of 12 points per time period, while Zimbabwe had a fixed loss of -16 points.

First, we examined the final utility values of all the negotiations for each player, and the sums of the final utility values. When the automated agent played the role of England the average utility value achieved by the automated agent was 565.1, while the average for the human playing the role of England was 331.8. The results show that our agent achieved significantly higher utility values as opposed to a human agent playing the same role (using 2-sample t-test: $t(22) = 3.10, p < 0.004$). On the other hand, when

Table 2: Final negotiations utility values, ranking values, sums of utility values and sums of ranking values. England-Zimbabwe Domain

Parameter	Avg	Stdev
$\mathbf{Q_{Eng}}$ vs. H_{Zim}	565.1	283.30
$rank(\mathbf{Q_{Eng}})$ vs. $rank(H_{Zim})$	0.82	0.17
$\mathbf{H_{Eng}}$ vs. H_{Zim}	331.8	210.6
$rank(\mathbf{H_{Eng}})$ vs. $rank(H_{Zim})$	0.58	0.19
$\mathbf{Q_{Zim}}$ vs. H_{Eng}	18.45	223.1
$rank(\mathbf{Q_{Zim}})$ vs. $rank(H_{Eng})$	0.64	0.13
$\mathbf{H_{Zim}}$ vs. H_{Eng}	-92.6	247.90
$rank(\mathbf{H_{Zim}})$ vs. $rank(H_{Eng})$	0.60	0.15
$\mathbf{H_{Zim}}$ vs. $\mathbf{Q_{Eng}}$	-322.55	265.94
$rank(\mathbf{H_{Zim}})$ vs. $rank(Q_{Eng})$	0.41	0.16
$\mathbf{H_{Eng}}$ vs. $\mathbf{Q_{Zim}}$	311.50	204.79
$rank(\mathbf{H_{Eng}})$ vs. $rank(Q_{Zim})$	0.57	0.18
Sum - $\mathbf{H_{Eng}}$ vs. $\mathbf{Q_{Zim}}$	330	222.8
Sum - $rank(\mathbf{H_{Eng}})$ vs. $rank(\mathbf{Q_{Zim}})$	1.21	0.07
Sum - $\mathbf{H_{Zim}}$ vs. $\mathbf{Q_{Eng}}$	242.5	409.4
Sum - $rank(\mathbf{H_{Zim}})$ vs. $rank(\mathbf{Q_{Eng}})$	1.25	0.03
Sum - $\mathbf{H_{Eng}}$ vs. $\mathbf{H_{Zim}}$	239.2	298.8
Sum - $rank(\mathbf{H_{Eng}})$ vs. $rank(\mathbf{H_{Zim}})$	1.17	0.07

the agent played the role of Zimbabwe, there is no significant difference between the utility values of the agent and the human player, though the average utility value for the automated agent was higher (18.45) than that of the humans (-92.6). One explanation for the higher values achieved by the QO agent is that the QO agent is more eager to accept agreements than humans. When playing the Zimbabwe side, which has a negative time cost, accepting agreements sooner rather than later allowed the agent to gain higher utility values than the human playing the same side.

This is also supported when we consider the results of the ranking of the agreements. When the automated agent played the role of Zimbabwe, the average ranking it achieved was similar to the ranking the human players achieved playing the same role (0.64 and 0.60). On the other hand, when the automated agent played the role of England it achieved significantly higher ranking values than the human playing the same role, with an average of 0.82 as compared to only 0.58 (using 2-sample Wilcoxon test, $p < 0.002$).

Comparing the sum of utility values of both negotiators, based on the role the agent played, we show that this sum is higher when the negotiations involved the agent. When the automated agent played the role of Zimbabwe, the sum of utility values was 330 as opposed to only 239.2 when two humans were involved. When the automated agent played the role of England, the sum of utility values was 242.5. In comparing the sum of the rankings, we note that when the automated agent was involved the sum of rankings was higher than when only humans were involved (an average of 1.21 and 1.25 when the automated agent played the role of Zimbabwe and England respectively, and an average of 1.17 when the human players played against each other).

Table 3: Final negotiations utility values, ranking values, sums of utility values and sums of ranking values. Job Candidate Domain

Parameter	Avg	Stdev
$\mathbf{Q_{Can}}$ vs. H_{Emp}	409	93.95
$rank(\mathbf{Q_{Can}})$ vs. $rank(H_{Emp})$	0.75	0.19
$\mathbf{H_{Can}}$ vs. H_{Emp}	309.7	140.2
$rank(\mathbf{H_{Can}})$ vs. $rank(H_{Emp})$	0.56	0.29
$\mathbf{Q_{Emp}}$ vs. H_{Can}	437.3	121.7
$rank(\mathbf{Q_{Emp}})$ vs. $rank(H_{Can})$	0.77	0.19
$\mathbf{H_{Emp}}$ vs. H_{Can}	410.6	114.0
$rank(\mathbf{H_{Emp}})$ vs. $rank(H_{Can})$	0.75	0.20
$\mathbf{H_{Can}}$ vs. Q_{Emp}	342.45	114.40
$rank(\mathbf{H_{Can}})$ vs. $rank(Q_{Emp})$	0.58	0.24
$\mathbf{H_{Emp}}$ vs. Q_{Can}	448.82	82.41
$rank(\mathbf{H_{Emp}})$ vs. $rank(Q_{Can})$	0.74	0.21
Sum - $\mathbf{H_{Emp}}$ vs. $\mathbf{Q_{Can}}$	852.8	132.2
Sum - $rank(\mathbf{H_{Emp}})$ vs. $rank(\mathbf{Q_{Can}})$	1.49	0.23
Sum - $\mathbf{H_{Can}}$ vs. Q_{Emp}	779.7	199.0
Sum - $rank(\mathbf{H_{Can}})$ vs. $rank(\mathbf{Q_{Emp}})$	1.35	0.24
Sum - $\mathbf{H_{Emp}}$ vs. $\mathbf{H_{Can}}$	720.3	212.5
Sum - $rank(\mathbf{H_{Emp}})$ vs. $rank(\mathbf{H_{Can}})$	1.30	0.27

However, this is only significant when the automated agent played the role of England (using 2-sample Wilcoxon test, $p < 0.001$).

Another important aspect of the negotiation is the outcome - whether a full agreement was reached or whether the negotiation ended with no agreement (either status-quo or opting out) or with a partial agreement. While only 64% of the negotiations involving only people ended with a full agreement, more than 72% of the negotiations involving the automated agent ended with a full agreement. Using *Fisher's Exact test* we determined that there is a correlation between the kind of agent the opponent is (be it an agent or a human) and the form of the final agreement (full, partial or none). The results show that there is a significantly higher probability of reaching a full agreement when playing against the agent ($p < 0.006$).

Results for the Job Candidate Domain Table 3 summarizes the average utility values of all the negotiations, the average ranking of the agreements reached, and the average of the sums of utility values and ranking of the agreements in all the experiments in the Job Candidate domain. H_{Can} and H_{Emp} denote the utility value gained by people playing the role of the job candidate or the employer, respectively, and Q_{Can} and Q_{Emp} denote the utility value gained by the QO agent playing either role.

The utility values ranged from 170 to 620 for the employer role and from 60 to 635 for the job candidate role. The Status-Quo value in the beginning of the negotiation was 240 for the employer and -160 for the job candidate. Both players had a fixed loss per time period - the employer of -6 points and the job candidate of -8 points per period.

We started by examining the final utility values of all the negotiations for each player, and the sums of the final util-

ity values. The results show that when the automated agent played the role of the job candidate, it achieved significantly higher utility values with an average of 409.0, as opposed to a human agent playing the same role, with an average of 309.7 (using 2-sample *t-test*: $t(22) = 2.76, p < 0.008$). On the other hand, when playing the role of the employer, there is no significant difference between the utility values of the agent and the human player, though the average utility value for the automated agent was higher than for the human (437.3 vs. 410.6).

This is also supported when considering the results of the ranking of the agreements. When the automated agent played the role of the employer, the average ranking it achieved was higher than the ranking the human players achieved playing the same role (0.77 vs. 0.75). Moreover, when the automated agent played the role of the job candidate it achieved significantly higher ranking values than the human playing the same role, with an average of 0.75 as compared to 0.56 (using 2-sample Wilcoxon test, $p < 0.03$).

Comparing the sum of utility values of both negotiators, based on the role the agent played, we show that this sum is also significantly higher when the negotiations involved the agent, when it played the role of the job candidate with an average of 852.8 (using 2-sample *t-test*: $t(22) = 2.48, p < 0.002$). When the automated agent played the role of the employer, the sum of utility values was higher, though not significantly (779.7 vs. 720.3). When comparing the sum of the rankings, we note that when the automated agent was involved the sum of rankings was higher than when only people were involved (an average of 1.49 and 1.35 when the automated agent played the role of the job candidate and the employer respectively, and an average of 1.30 when people played against each other). However, this is only significant when the automated agent played the role of the job candidate (using 2-sample Wilcoxon test, $p < 0.041$).

Another important aspect of the negotiation is the outcome - whether a full agreement was reached or whether the negotiation ended with no agreement (either status-quo or opting out) or with a partial agreement. While only 72% of the negotiations involving only people ended with a full agreement, 86% of the negotiations involving the automated agent ended with a full agreement. Using *Fisher's Exact test* we determined that there is a correlation between the kind of the opponent agent (be it the agent or the human) and the form of the final agreement (full, partial or none). The results show that there is a significantly higher probability of reaching a full agreement when playing against the agent ($p < 0.006$).

Discussion: Results against People Using GPNEG we were able to show the efficacy of the automated agent embodied in the system. We can now use the tool also to analyze the simulation logs and provide some insights to the behavior of both negotiators during the experiments.

The results show that the automated agent achieved higher utility values than the human counterpart. This can be explained by the nature of our agent both for accepting offers and generating offers. Using the qualitative offer mechanism we allow the agent to propose agreements that are good for

it, but also reasonable for its opponent. This is opposed to people, who generally tend to reject offers that do not seem to be good for them based predominantly on the content of the offer and less on the utility value of the offer. Another explanation for the high scores of the agent is its rationality, compared to the bounded rationality of people, in terms of assessing the offers. The automated agent makes straightforward calculations. It evaluates the offer based on its attributes, and not based on its *content*. In addition, it also puts more weight to the fact that it loses or gains as time advances. This is not the case, however, when analyzing the logs of the people. It seems that people put more weight to the content of the offer than to its value. This was more evident in the Job Candidate domain which the human subjects could more easily identify with.

Yet, this does not explain why, in both domains, these results are significant only for one of the sides. In the England-Zimbabwe domain, the results are significant when the agent played the role of England, while in the Job Candidate domain these results are significant when it played the role of the job candidate. It is interesting to note that our results, which show that the automated agents play significantly better when playing one of the sides, while playing just as well when playing the other side, are not unique. Kraus *et al.* (Kraus, Hoz-Weiss, & Wilkenfeld 2007) also experimented with an automated agent playing against humans in a fishing dispute domain, and they presented similar results.

Another possible explanation for this phenomenon can be found by examining the logs of the negotiations and the values of the agreements. In both domains we can see that the England side and the job candidate sides are the more dominant sides and have more leverage than the other side. For example, for England this is represented by the fact that it gains as time advances so it can put more pressure on the other side to accept agreements. For the job candidate side, this is more of a psychological interpretation. It seems that the job candidate side has less to lose in the negotiation. While both the employer and the job candidate lose as time passes, the status quo agreement ensures the hiring of the job candidate.

Conclusions and Future Work

Due to the unique and innovative nature of our proposed tool, we expect that the marketing of this type of tool will make waves both in the academia and in the market. Thus, it will allow for future collaboration and long-term relationships both in the academia and the commercial market.

Our future work involves improving the negotiation tool as well as adding important aspects to assist in the negotiation experiments. First we plan to include a observing/commenting agent that will observe the trainee in the negotiation and will comment on his/her actions. To provide a good commentary component we will incorporate an efficient analysis tool that will investigate the actions taken by the subject along with the history of the current negotiation and past negotiations. Based on this the program can suggest or direct the subject to the optimal behavior. We will also allow for a comparison between the people's acts and actions another automated agent would have taken in the same situation. Thus, we will be able to provide the subject a wide variety of information and comparison of his/her actions for future uses.

In addition, we will work on incorporating an argumentation component in the automated agent that will attempt to change the opponent's preferences.

Finally, we propose to develop an adaptive training manager that will allow for dynamically presenting the user with different negotiation scenarios as well as facing different variation of the *QOAgent*. This is an important feature in the future GPNEG. Analyzing the behavior of the user and comparing it to other behaviors or to the automated agents behavior will for allow better understanding of his/her weaknesses in the negotiation process. This in turn will allow for the adaptation of both the scenario and the agent matched with the user to best tackle these issues and make the experience more beneficial and fruitful.

References

Capra, C. M.; Goeree, J. K.; Gomez, R.; and Holt, C. A. 1999. Anomalous behavior in a traveler's dilemma? *American Economic Review* 89(3):678–690.

Cohen, R. 1997. *Negotiating Across Cultures: Communication Obstacles in International Diplomacy*. Washington, D.C: United States Institute of Peace Press, fourth edition.

Dubois, D.; Prade, H.; and Sabbadin, R. 2001. Decision-theoretic foundations of qualitative possibility theory. *European Journal of Operational Research* 128:459–478.

Faratin, P.; Sierra, C.; and Jennings, N. R. 2002. Using similarity criteria to make issue trade-offs in automated negotiations. *AIJ* 142(2):205–237.

Fatima, S., and Wooldridge, M. 2004. An agenda based framework for multi-issue negotiation. *AIJ* 152:1–45.

Fatima, S.; Wooldridge, M.; and Jennings, N. R. 2005. Bargaining with incomplete information. *AMAI* 44(3):207–232.

Kraus, S.; Hoz-Weiss, P.; and Wilkenfeld, J. 2007. Resolving crises through automated bilateral negotiations. *AIJ*.

Lin, R. J., and Chou, S. T. 2003. Bilateral multi-issue negotiations in a dynamic environment. In *Proc. of the AAMAS Workshop on Agent Mediated Electronic Commerce (AMEC V)*.

Lin, R.; Kraus, S.; Wilkenfeld, J.; and Barry, J. 2006. An automated agent for bilateral negotiation with bounded rational agents with incomplete information. In *Proceedings of the 17th European Conference on Artificial Intelligence (ECAI)*, 270–274.

Oprea, M. 2002. An adaptive negotiation model for agent-based electronic commerce. *Studies in Informatics and Control* 11(3):271–279.

Osborne, M. J., and Rubinstein, A. 1994. *A Course In Game Theory*. Cambridge MA: MIT Press.

Rasmusen, E. 2001. *Games and Information: An Inroduction to Game Theory*. Blackwell Publishers.

Tennenholtz, M. 1996. On stable social laws and qualitative equilibrium for risk-averse agents. In *KR*, 553–561.

Zeng, D., and Sycara, K. 1998. Bayesian learning in negotiation. *International Journal of Human Computer Studies* 48(1):125–141.