Investigating the Benefits of Automated Negotiations in Enhancing People's Negotiation Skills*

Raz Lin[1], Yinon Oshrat[1] and Sarit Kraus[1,2]
[1] Department of Computer Science
Bar-Ilan University
Ramat-Gan, Israel 52900
{linraz,sarit}@cs.biu.ac.il

[2] Institute for Advanced Computer Studies
University of Maryland
College Park, MD 20742 USA

ABSTRACT

Negotiation surrounds our day-to-day lives. Research in the field of automated negotiations has suggested the design and use of automated negotiators, on one hand to allow facilitation of the negotiation process by human negotiators and, on the other hand to provide automated agents that can negotiate on behalf of humans. Many papers present innovative agents and evaluate their efficacy in negotiations with other automated agents or people. Others focus on building negotiation support systems with the purpose of helping negotiators reach an agreement. Yet, the question still remains whether these systems or agents have the potential of improving people's negotiation skills. In this paper we attempt to shed more light on this topic. By means of extensive simulations with human negotiators we examine and compare several training methods and their implications on the improvement of negotiation skills of human negotiators.

Categories and Subject Descriptors

H.1.2 [**Information Systems**]: User/Machine Systems

General Terms

Experimentation

Keywords

automated bilateral negotiation, opponent modeling

1. INTRODUCTION

The use of simulation and role-playing is common for training people in negotiations (e.g., the Interactive Computer-Assisted Negotiation Support system (ICANS) [14], the InterNeg Support Program for Intercultural REsearch (INSPIRE) [7] and virtual humans for training [6]). Surprisingly, extensive research of whether these simulations or role-playing indeed improve people's negotiation skills has not been conducted. Many of the negotiation support systems are only used as a mechanism to help negotiators reach an agreement, whereas the underlying assumption of other researchers is that role-playing improves people's negotiation skills [13].

The introduction of automated negotiator agents allows in-depth investigation into the question of whether role-playing and simulations indeed improve the negotiation skills of human negotiators. Automated negotiators have also been proposed for negotiation with humans in the literature [2, 5, 8, 9]; but again, they have not been evaluated for their efficacy in improving the negotiation skills of human negotiators. As we will demonstrate in the rest of the paper, automated agents capable of negotiating with people can provide a breakthrough in training human negotiators and enhancing their negotiation skills.

In this paper we provide results of extensive experiments involving human subjects. In the experiments we investigate several methods in which people are involved in negotiations and try to find out whether these methods allow them to improve their negotiation skills. We identified several possible ways of training people in negotiations. The first is actually undergoing the negotiation process. This can either involve negotiating with another human counterpart or with an automated negotiator. The second is by means of designing an automated negotiator.

In order to obtain an objective method for comparing the different training techniques, we propose a unique evaluation method, namely, using a *standardized automated negotiator*. To this end, we need an automated negotiator, which has been shown to be an efficient negotiator against human negotiators. Following a review of the literature on automated negotiations with human counterparts (e.g., *AutONA* agent [2], Colored-Trail agent [4], heuristic based agent [5]) we chose the *QOAgent* [9] as our standardized automated negotiator. The *QOAgent* has been shown to be an efficient automated negotiator, especially with respect to negotiating with people. Consequently, we believe it could serve as a good standardized negotiator for our experiments. In addition, its simulation environment is rich and supports bilateral multi-issue and multi-attribute negotiations, both with a human counterpart and automated agents.

*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.

Cite as: Investigating the Benefits of Automated Negotiations in Enhancing People's Negotiation Skills, Raz Lin, Yinon Oshrat, Sarit Kraus, *Proc. of 8th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2009)*, Decker, Sichman, Sierra and Castelfranchi (eds.), May, 10–15, 2009, Budapest, Hungary, pp. 345–352

Several groups of human negotiators participated in the simulations. The first group (we will refer to it hereafter as the control group) comprised human negotiators that were not given any training before they negotiated with our standardized automated negotiator. All the other groups performed one of the following training methods, before negotiating against the standardized automated agent:

- Classical role playing with another human counterpart;
- Negotiating with another automated negotiator;
- Designing an automated negotiator.

Thus, we were able to compare the results of a group of human negotiators, who had no training, and negotiated against the automated agent, with the results of the other groups of human negotiators, who also negotiated against the same automated agent after undergoing training.

Note that the focus of this paper is not on designing the best automated negotiator, but rather to demonstrate the benefits of using one both for training people in negotiation and as an objective evaluation measure. The results of our comprehensive experiments indeed show that using automated negotiator agents helps improve people's negotiation skills, while classical role playing between human negotiators generates no significant improvement in their negotiation skills.

This paper contributes to research on automated negotiations in several ways. First, it tackles the problem of evaluating the efficiency of simulations for improving people's negotiation skills. It proposes the use of a standardized automated negotiator and automated agents as a tool for training people in negotiation tasks and successfully improving their negotiation skills. Given the importance of negotiating efficiently and the extensive simulations currently implemented, this contribution cannot be overstated. Second, the simulation tool and the method that we describe can serve as a basis for further exploration of these issues and for the objective evaluation of other techniques used for training people in negotiations. Together, the negotiation environment will enable exploration of future research directions and thereafter our method can be used also to better understand behavioral and cognitive aspects of negotiations undertaken by human negotiators.

The rest of this paper is organized as follows. In Section 2 we review related work in the field of automated negotiators' design. We continue and describe our methods for improving people's negotiation skills in Section 4. In Section 5 we present the experiments we conducted, our methodology and our evaluation results. Finally, we conclude the paper with open questions and future directions for research.

2. RELATED WORK

At the University of Maryland, the International Communication and Negotiation Simulations (ICONS) project has been running for over twenty years now as a tool for teaching students on international relations and conflict resolutions, mainly based on bilateral negotiation simulations and role-playing [12]. The students act as diplomats and negotiate with each other on various international problems. However, it is interesting to note that the only method of evaluation used to date to determine whether the simulation indeed helps the students in learning negotiation dynamics has been a debriefing conducted by faculty members with the students following simulations.

Susskind and Corburn [13] also identified the problem of the existence of simulations aimed to teach negotiation via role-playing with automated agents, which lack evidence to back if it indeed supports the training process. However, the authors try to tackle this problem from the pedagogical aspect and attempt to understand the usefulness of simulations by questioning leading practitioners in the field about why and how they use simulations to teach negotiation. In addition, they assert that there cannot be a "one-size-fits-all approach to using simulations to teach negotiation". In this paper we attempt to present the evidence of methods that do allow this to occur.

Ross *et al.* [17], on the other hand, investigate how the negotiation skills of students in relevant courses can be improved. They suggest using an interactive video simulation, designed to teach negotiation skills via a dynamic negotiation with an automated agent. The simulation uses a single sales negotiation scenario. Using the simulations, Ross *et al.* tried to see whether the simulation helped students increase their learning of negotiation concepts. While they focus on measuring the reaction of the students to the simulation, they also showed that by interacting with it students better learned to recognize important decision points in the negotiation. While we also evaluate the improvement of the negotiation skills of the students using a single domain, our simulation environment is intended to be used in a wider range of domains and thus can be extensively used for future investigation of these issues. In addition, their evaluation method is more subjective and is based on paper-and-pencil pretest and posttest. Some of the questions are subjective in nature (e.g., "How valuable was using the simulator as a learning experience?" and "How valuable was the simulator for teaching bargaining skills?"), while the others consisted of a series of 20 multiple-choice problem situations. While they do show a significant improvement in the test scores, the underlying drawback of the evaluation method still remains. We, on the other hand, propose a more objective evaluation criteria in which the students' performance in the negotiation itself is compared using a standardized automated negotiator.

Similar to Ross *et al.* Druckman and Ebner [3] conducted an experiment in which they demonstrated that students practicing in simulations demonstrate a better understanding of negotiation concepts. They present the example of the benefits of learning by designing of simulations and provide observations that suggest that learning by designing provides more insights into the negotiation concepts than learning by negotiating. They conducted an experiment to investigate whether the design process increases short-term learning than actual negotiation. However, once again, the evaluation process was based on a questionnaire, comprising multiple-choice questions and open-ended questions. As we stated earlier, in this paper we provide more objective methods to evaluate the efficacy of automated agents in increasing people's negotiation skills.

Lastly, Kenny *et al.* [6] and Traum *et al.* [15] describe work on virtual humans that are used for interpersonal training for skills, such as: negotiation, leadership, interviewing and cultural training. The virtual agent also tries to model the opponent by reasoning about its mental state. Traum *et al.* tested their agents in several negotiation scenarios.

One such scenario is a simulation for soldiers that practice and conduct bilateral engagements with virtual humans, in situations in which culture plays an important role. In this case, the different actions can be selected from a menu which includes appropriate questions based on the history of the simulation thus far. They state that the subjects enjoy using the system for negotiations and that it also allows them to learn from their mistakes. Nonetheless, they did not perform extensive experiments to verify this claim.

3. PROBLEM DESCRIPTION

We consider the problem of people negotiating and how the interaction with automated negotiators can enhance the negotiation experience by improving the negotiation skills of human negotiators. We consider a bilateral negotiation in which two agents, either automated negotiators or people, negotiate to reach an agreement on conflicting issues. The negotiation can end either when (a) the negotiators reach a full agreement, (b) one of the agents opts out, thus forcing the termination of the negotiation with an opt-out outcome (OPT), or (c) a predefined deadline is reached, whereby, if a partial agreement is reached it is implemented or, if no agreement is reached, a status quo outcome (SQ) is implemented. Let I denote the set of issues in the negotiation, O_i the finite set of values for each $i \in I$ and O a finite set of values for all issues ($O_1 \times O_2 \times \ldots \times O_{|I|}$). We allow partial agreements, $\perp \in O_i$ for each $i \in I$. Therefore an offer is denoted as a vector $\vec{o} \in O$. Since no agreement is worse than any agreement, and a status quo is implemented if the deadline is reached, we assume that default values are assigned to each attribute. Thus, if both sides agree only on a subset of the issues and the deadline is reached, the unresolved issues are assigned their default value and thus a partial agreement can be implemented.

It is assumed that the agents can take actions during the negotiation process until it terminates. Let **Time** denote the set of time periods in the negotiation, that is $\textbf{Time} = \{0, 1, ..., dl\}$. Time also has an impact on the agents' utilities. Each agent is assigned a time cost which influences its utility as time passes. In each period $t \in \textbf{Time}$ of the negotiation, if the negotiation has not terminated earlier, each agent can propose a possible agreement, and the other agent can either accept the offer, reject it or opt out. Each agent can either propose an agreement which consists of all the issues in the negotiation, or a partial agreement. We use an extension of the model of alternating offers ([10], p. 118-121), in which each agent can perform up to $M > 0$ interactions with its counterpart in each time period.

The negotiation problem also involves incomplete information concerning the opponent's preferences. We assume that there is a finite set of agent types. These types are associated with different additive utility functions (e.g., one type might have a long term orientation regarding the final agreement, while the other type might have a more constrained orientation). Formally, we denote the possible types of agents $\textbf{Types} = \{1, \ldots, k\}$. Given $l \in \textbf{Types}$, $1 \leq l \leq k$, we refer to the utility of an agent of type l as u_l, and $u_l : \{(O \cup \{SQ\} \cup \{OPT\}) \times \textbf{Time}\} \rightarrow \mathbb{R}$. Each agent is given its exact utility function. The negotiators are aware of the set of possible types of the opponent. However, the exact utility function of the rival is private information.

We continue to describe the different uses of automated agents upon which the experiments were conducted to evaluate their efficacy in improving people's negotiation skills.

4. ENHANCING PEOPLE'S NEGOTIATION SKILLS

Using a simulation environment in which training of negotiation skills can be done, we propose several ways to apply automated negotiators in order to train human negotiators. Moreover, we systematically compare these methods to evaluate their efficiency in training people. We first describe the simulation environment and then we continue to present the different methods used to enhance the negotiation skills that we tested.

4.1 The Simulation Environment

The simulation environment we used 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 simulations' interface. The automated agents can play either role in the negotiation, while the human counterpart accesses the negotiation interface via a web address. 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. 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.

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. Another parameter is the number of turns in the negotiation. The simulation environment allows each player to perform any number of interactions with the opponent player at any given time period (that is, it extends the model of alternating offers [10]). The number of turns in 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 for each player. Each player can also 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. Finally, the simulation is loaded after setting the opponent type. This can be either two human players or a human player playing against an automated agent.

During each phase of the negotiation, the instructions and 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. When receiving an offer the player can choose whether to accept or reject it, or make a counter-offer.

4.2 Methods for Enhancing Negotiation Skills

We used the *QOAgent*, an automated negotiator which has been previously shown to be an efficient negotiator against human counterparts [9], to evaluate the methods we applied. The *QOAgent* incorporates two mechanisms. The first involves a learning mechanism based on the Bayesian updating rule to try and compensate for the incomplete information and learn the opponent's model. The second mechanism is a decision making one which uses a non-classical model of offers' valuation, rather than the traditional quantitative decision making model. In essence, the decision making valuation component takes into account the agent's utility function, as well as the believed type of the opponent. This data is used both for deciding whether to accept or reject an offer and for generating an offer. The *QOAgent* behavior is not deterministic. First, it is based on the modeling of the opponent which is dynamically generated based on the dynamics of the negotiation. Second, it uses randomization in the decision making mechanism when deciding which offer to accept.

In order to evaluate the different training methods that we propose, all people negotiated against the same agent (the *QOAgent*) after undergoing their training session. In addition, a control group was used to compare the different results. This group consisted of people that had not undergone any training before negotiating against the *QOAgent*.

The first training method was the classical role playing of two people, that is, negotiating with another person. While role playing might be the simplest training method and used in classes, in the general case it is hard to find human negotiators with whom one can train. Extensive training with people requires a great deal of scheduling to match the different human negotiators.

The second approach that we evaluate is role playing with an automated negotiator. In this approach, the human negotiator is matched with the *KBAgent*. This is a different automated agent that was designed in order to improve the performance of the *QOAgent* by using a general opponent modeling technique. The *KBAgent* uses a database of past negotiation sessions between specific agent types to allow it be more efficient in subsequent negotiations with agents of that specific type. Based on the existing database, the agent performs an offline learning, founded on the kernel based density estimation ([16], Chapter 2). This type of learning allows the agent to attach acceptance probabilities to each possible agreement and then use these probabilities in its decision making component, either when proposing a new offer or when determining its concession rate. After negotiating with the *KBAgent*, the human negotiators are matched with the original automated agent. We then compared their results when matched with the standardized automated negotiator to that of the control group. In this training method we used a specific automated agent which has advantages over the classical role playing model of two human negotiators. This is due to the fact that automated negotiators are always available and no scheduling is needed. Moreover, there are many types of automated negotiators (e.g., [2, 5, 15]) and the human negotiator can train each time with a different type and observe the different attitudes towards the negotiation and explore different negotiation strategies. In addition, there are several automated agents that negotiate efficiently with people in several domains and thus the human negotiator can train himself/herself in various domains in an easy manner.

The third approach that we examine is the improvement of negotiation skills due to the actual design of an automated negotiator by the human subjects. In this case, the human negotiators were given a task to implement an efficient automated agent for the Job Candidate domain, which is described in Section 5.1. The implementation was done in the same simulation environment as the *QOAgent*. The students were provided skeleton classes to help them implement their agents. This also allowed them to focus on the strategy and the behavior of the agent, and eliminate the need to implement the communication protocol or the negotiation protocol. In addition, it provided them with a simulation environment in which they could test their agents and their strategies. On the one hand, this approach is more time consuming than the previous approaches. On the other hand, it provides the human negotiator a great deal of understanding of negotiation techniques, strategies and different attitudes towards negotiations. Yet, it requires the human negotiator to have knowledge in computer science in order to code the strategy of the agent's behavior.

To evaluate each training method, the people had to negotiate against the *QOAgent*, which served as the standardized automated negotiator. The results of the people when negotiating against the *QOAgent* after being trained were compared to the control group of people, who negotiated against the *QOAgent* without undergoing training in advance.

5. EXPERIMENTS

The experiments were conducted using the simulation environment and a given multi-attribute multi-issue domain. We begin by describing the domain which was used in all the experiments and then continue to describe the experimental methodology and results.

5.1 The Negotiation Domain

For the negotiation domain we chose a Job Candidate domain, which is related to the subjects' experience, and thus they could better identify with it. In this domain, which was first described in [9], 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. Below are the issues under negotiation:

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. Thus, a total of 3 possible values are allowed for this issue.

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 his/her work place and his/her prestige. The possible values are (a) QA, (b) programmer, (c) team manager, or (d) project manager. Thus, a total of 4 possible values are allowed for this issue.

3. **Social benefits**. The social benefits are an addition to the salary and thus impose an extra expense on 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) providing 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. Thus, a total of 12 possible values ($3 \times 4 = 12$) are allowed for this issue.

4. **Promotion possibilities**. This issue describes the commitment by the employer regarding the 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. Thus, a total of 3 possible values are allowed for this issue.

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. Thus, a total of 3 possible values are allowed for this issue.

In this scenario, a total of 1,296 possible agreements exist ($3 \times 4 \times 12 \times 3 \times 3 = 1296$).

Each turn in the scenario equates to two minutes of the negotiation, and the negotiation is limited to 28 minutes. If the parties 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 agents 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 slim and it is impossible to negotiate anymore. Time also has an impact on the negotiation. As time advances the candidate's utility decreases, as the employer's good impression of the job candidate decreases. The employer's utility also decreases as the candidate becomes less motivated to work for the company.

The utility values range 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.

As there is also incomplete information, we assume that there are three possible types of agents for each role. These types are associated with different additive utility functions. The different types are characterized as ones with short-term orientation regarding the final agreement, long-term orientation and a compromising orientation. Detailed score functions for the domain can be found in Appendix A.

5.2 Experimental Methodology

We ran an extensive set of simulations, consisting of a total of 148 human negotiators. The human negotiators were mostly computer science undergraduate and graduate students, while a few were former students who are currently working in the Hi-Tech industry. Table 1 summarizes the number of different human subjects we had per each method we evaluated. Each subject served only one specific role in the negotiations (either the employer role or the job candidate one).

While other research has attempted to evaluate the efficiency of the training method, researchers have mostly based their findings on questionnaires or allowed human negotiators to re-negotiate, and thus the results are bias. While well-designed questionnaires may be constructed to allow providing objective and useful insights, most papers we found (e.g., [3, 17]) rely on questionnaires, which are subjective. For example, subjects were asked how they evaluated their negotiation experiment, whether they believe they are better trained now and the sort. We, on the other hand, evaluate the efficiency of the negotiation method by using an objective measure. This is done by using a standardized negotiator and comparing the results of the people that negotiate with it after their training method to a control group of people who negotiated against it without undergoing training.

Approach/Role	Employer	Job Candidate
Control Group	18	16
Training via Human Negotiation	18	18
Training via Automated Negotiator	20	20
Training via Agent Design	19	19

Table 1: Number of subjects in each evaluation method.

Each simulation was divided into two parts: (i) training method, and (ii) negotiating against the standardized agent. 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. While the subjects knew that they will negotiate twice, they did not know in advance against whom they played (whether it is a human negotiator or an automated one).

We continue with the description of the experimental results.

5.3 Experimental Results

Table 2 summarizes the control group's results, that is, the average utility scores of human negotiators without prior training when playing against the *QOAgent*. Throughout this section, we also evaluate the significance of the results. The significant test was performed by applying the *t-test* on the results. The *t-test* is a statistical hypothesis test in which the test statistics has a *t-distribution* if the null hypothesis is true. This test requires a normal distribution of the measurements ([1], Chapter 3). Thus, it is used in our analysis in order to compare the utility values of the different simulation methods, which have continuous values. To analyze the significant difference in the end turn we use the *Wilcoxon signed-rank test*, which is a non-parametric alternative to the paired t-test for the case of two related samples or repeated measurements on a single sample. This test does not require any assumptions regarding the distribution of the measurements ([11], Chapter 5).

Tables 3 and 4 summarize the average utility scores achieved by the human negotiators and the average end turn of the negotiation during all the simulations, respectively. The tables also present the statistical significance of the results compared to the control group. All the results are of the negotiations between people and the *QOAgent* after having

Role	Average	Standard Deviation
Employer	431.78	80.83
Job Candidate	320.5	112.71

Table 2: Average utility scores and standard deviation of the control group against the standardized negotiator.

Method	Role	Average	Std.	p-value
Control Group	Employer	431.78	80.83	
	Job Can.	320.5	112.71	
Training via	Employer	448.56	66.08	0.25
Human Negotiation	Job Can.	383.83	112.73	0.05
Training via	Employer	468.6	38.94	0.04
Automated Negotiator	Job Can.	433	102.84	0.002
Training via	Employer	466.84	46.26	0.06
Agent Design	Job Can.	391.53	76.75	0.02

Table 3: Comparison of the average utility scores and standard deviation of human negotiators using different training methods and the control group.

undergone training.

We can see that the classical training method of role playing between humans allows the human negotiators to achieve higher utility values (448.56 and 383.83 for the employer and job candidate roles, respectively) compared to the control group. However, this is only significant for one of the roles (the job candidate role, with p-value < 0.05). This is also the case if we compare the average turn in which the negotiation ended. Most negotiations ended more quickly after the training method (4.67 and 3.33 in both roles) as compared to the negotiations of the control group (5.06 and 5.37).

Comparing the next training method, in which role playing was done with another automated agent, we find that the average utility obtained by people using this training method is significantly higher for both roles compared to the control group (468.6 compared to 431.78 with a p-value < 0.04 for the employer role and 433 compared to 320.5 with a p-value < 0.002 for the job candidate role). The average end turn of the negotiation is also lower in this case for both roles (2.65 and 3.45 for the employer and job candidate, respectively) and even significantly lower in the case of the employer role (p-value < 0.001).

In the final training method that we evaluated, training via design of an automated negotiator, the results are also better than the control group, in terms of average utility values. The average utility values of the people in the training group were 466.84 (p-value < 0.06) and 391.53 (p-value < 0.02) for the employer and job candidate roles, respectively, compared to 431.78 and 320.5 for the control group. In the case of the employer role, the negotiation terminated faster (3.89 compared to 5.05 for the control group, p-value < 0.06), while in the job candidate role there was no significant difference between the end turns of the negotiation compared to the control group.

Method	Role	Average	Std.	p-value
Control Group	Employer	5.06	2.1	
	Job Can.	5.37	4.31	
Training via	Employer	4.67	3.05	0.47
Human Negotiation	Job Can.	3.33	2.89	0.14
Training via	Employer	2.65	2.16	<0.001
Automated Negotiator	Job Can.	3.45	2.72	0.18
Training via	Employer	3.89	2.08	0.06
Agent Design	Job Can.	5.58	2.87	0.4

Table 4: Comparison of the average end turn and standard deviation of human negotiators using different training methods and the control group.

Role	Average	Std.	Average	Std.	p-value
	Training via Human Negotiation		**Training via Agent Design**		
Employer	448.56	66.08	466.84	46.26	0.17
Job Can.	383.83	112.73	391.52	76.75	0.40

Table 5: Comparison of the average utility scores and standard deviation of training via human negotiations versus training via agent design.

We continued to test whether some of the training methods that we evaluated were better than others. To this end we compared the results of the people in each group when matched with the standardized automated negotiator after undergoing training. That is, we compared (a) the groups of people training via human negotiations and those trained via agent design, (b) training via automated negotiator versus training via agent design, and (c) training via automated negotiators compared to training via human negotiations. Tables 5, 6 and 7 summarize the comparison between the average utility scores of these training methods. While some training methods enabled the negotiators to achieve higher utility values than others, the results were not significant in any role or training method. In addition, if we compare the average end turn in the different training methods, we can see that training via automated negotiator allows reaching an agreement significantly faster than training via agent design (p-value < 0.01 and p-value < 0.006 for the Employer and Job Candidate roles, respectively). It was also significantly faster than training via human negotiation when playing the role of the Employer.

5.4 Discussion of Results

The experimental results indeed show that using automated agents can enhance the negotiation experience and thereby improve people's negotiation skills. However, im-

Role	Average	Std.	Average	Std.	p-value
	Training via Automated Negotiator		**Training via Agent Design**		
Employer	468.6	38.94	466.84	46.26	0.45
Job Can.	433	102.84	391.52	76.75	0.08

Table 6: Comparison of the average utility scores and standard deviation of training via automated negotiator versus training via agent design.

Role	Average	Std.	Average	Std.	p-value
	Training via Automated Negotiator		**Training via Human Negotiation**		
Employer	468.6	38.94	448.56	66.08	0.14
Job Can.	433	102.84	383.83	112.73	0.08

Table 7: Comparison of the average utility scores and standard deviation of training via automated negotiator versus training via human negotiations.

provement of people's negotiation skills is less evident with the classical role playing training method. The latter allowed people to achieve higher utility values and to finish the negotiation faster, however, the results were only significant in reference to one of the roles.

We are encouraged by the fact that the use of automated agents enables significant improvement of people's negotiation skills, which was shown throughout the experiments. Whether achieved by training with automated negotiators or by designing strategies for automated negotiators, people successfully improved their average utility scores after undergoing training.

Surprisingly, no significant differences were found between the different training methods. Yet, in all methods higher utility values were achieved when the automated agents were involved as opposed to the classical role simulation with two people. Since we did not focus in this paper on the design of an automated negotiator we did not experiment with other automated agents other than the *QOAgent* and the *KBAgent*. It might be the case that another automated negotiator could achieve higher utility values and cause the difference between the methods to be significant.

In addition, higher utility values were achieved in the training via automated negotiators compared to the design of an automated agent. Though the results are not significant, it is quite obvious that training via automated negotiators is a much more simpler task and less time consuming then training via agent design.

6. CONCLUSIONS

In this paper we presented an extensive systematic experimentation to answer the question whether simulation role-playing indeed improves the negotiation skills of humans. To do so, we provide an objective measure to evaluate the improvement of negotiation skills based on a standardized automated agent.

Our results reveal the potential embodied in automated agents as a key for training methods for people in negotiations, which turned out to be better than the classical method of role playing of two people. This fact motivates us to continue to introduce automated agents capable of negotiating efficiently with people, and to further evaluate their efficacy not only in the negotiation process, but also as a training facility.

Future work warrants careful investigation due to lack of significant differences between the various automated training techniques. We will evaluate other training methods and compare them to the standardized automated negotiator. In addition, we would like to support our findings by validating the training methods on a different domain. In other words, people will be trained in one domain and the improvement in their negotiation skills will be investigated in a different domain.

7. REFERENCES

[1] S. R. Brown and L. E. Melamed. *Experimental Design and Analysis.* Sage Publications, Inc., CA, USA, 1990.

[2] A. Byde, M. Yearworth, K.-Y. Chen, and C. Bartolini. AutONA: A system for automated multiple 1-1 negotiation. In *Proceedings of CEC'03*, pages 59–67, 2003.

[3] D. Druckman and N. Ebner. Onstage or behind the scenes? relative learning benefits of simulation role-play and design. *Simulation and Gaming*, to appear.

[4] S. Ficici and A. Pfeffer. Modeling how humans reason about others with partial information. In *Proceedings of AAMAS'08*, pages 315–322, 2008.

[5] C. M. Jonker, V. Robu, and J. Treur. An agent architecture for multi-attribute negotiation using incomplete preference information. *JAAMAS*, 15(2):221–252, 2007.

[6] P. Kenny, A. Hartholt, J. Gratch, W. Swartout, D. Traum, S. Marsella, and D. Piepol. Building interactive virtual humans for training environments. In *Proceedings of Interservice/Industry Training, Simulation and Education Conference (I/ITSEC)*, 2007.

[7] G. E. Kersten and S. J. Noronha. Www-based negotiation support: design, implementation, and use. *Decision Support Systems*, 25(2):135–154, 1999.

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

[9] 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.

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

[11] S. Siegel. *Non-Parametric Statistics for the Behavioral Sciences.* McGraw-Hill, NY, USA, 1956.

[12] B. A. Starkey and E. L. Blake. Simulation in international relations education. *Simulation and*

Gaming, 32(4):537–551, 2001.

[13] L. E. Susskind and J. Corburn. Using simulations to teach negotiation: Pedagogical theory and practice. Working Paper 99-1, Program on Negotiation at Harvard Law School, 1999.

[14] E. M. Thiessen, D. P. Loucks, and J. R. Stedinger. Computer-assisted negotiations of water resources conflicts. *GDN*, 7(2):109–129, 1998.

[15] D. Traum, S. Marsella, J. Gratch, J. Lee, and A. Hartholt. Multi-party, multi-issue, multi-strategy negotiation for multi-modal virtual agents. In *Proceedings of the 8th International Conference on Intelligent Virtual Agents*, 2008.

[16] M. Wand and M. Jones. *Kernel Smoothing*. Chapman & Hall, London, 1995.

[17] J. William H. Ross, W. Pollman, D. Perry, J. Welty, and K. Jones. Interactive video negotiator training: A preliminary evaluation of the mcgill negotiation simulator. *Simulation and Gaming*, 32(4):451–468, 2001.

APPENDIX

A. THE JOB CANDIDATE DOMAIN (I) SHORT-TERM, (II) LONG-TERM AND (III) COMPROMISE ORIENTATION SCORE FUNCTIONS

OUTCOMES	Job Candidate Outcome Weight / Importance			Employer Outcome Weight / Importance		
	(i)	(ii)	(iii)	(i)	(ii)	(iii)
Salary	20%	30%	15%	20%	15%	10%
7,000 NIS	3	2	3	8	7	7
12,000 NIS	6	6	5	6	6	6
20,000 NIS	8	9	6	3	3	4
Job Description	15%	25%	20%	20%	30%	20%
QA	2	-2	2	4	2	3
Programmer	4	3	4	6	6	6
Team Manager	5	6	6	4	3	4
Project Manager	6	8	8	2	1	3
Leased Car	20%	5%	10%	10%	10%	10%
Without leased car	-5	-5	-2	3	4	5
With leased car	5	5	2	-2	2	4
No agreement	0	0	0	0	0	0
Pension Fund	10%	5%	10%	10%	10%	10%
0% pension fund	-2	-2	-2	3	6	6
10% pension fund	3	4	3	4	4	4
20% pension fund	5	6	5	3	3	3
No agreement	0	0	0	0	0	0
Promotion Possibilities	5%	25%	35%	10%	20%	20%
Slow promotion track	4	1	-2	3	8	6
Fast promotion track	5	5	5	3	5	4
No agreement	0	0	0	0	0	0
Working Hours	30%	10%	10%	30%	15%	30%
10 hours	3	3	4	8	8	9
9 hours	5	4	5	6	6	6
8 hours	7	5	6	3	4	3
Time effect	**-8**	**-8**	**-8**	**-6**	**-6**	**-6**
Status Quo	**160**	**135**	**70**	**240**	**306**	**306**
Opting out	**150**	**75**	**80**	**210**	**150**	**215**