Automated Agents that Proficiently Negotiate with People: Can We Keep People out of the Evaluation Loop?*

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

Research on automated negotiators has flourished in recent years. Among the important issues considered is how automated negotiators can be designed that are capable of negotiating proficiently with people. Nonetheless, conducting experiments with people is timely and costly making the evaluation of these automated negotiators a very difficult process. Moreover, it makes it difficult to devise and revise the agent's strategies constantly, as each time it requires to gather an additional set of people for the experiments. In this paper we investigate the use of Peer Designed Agents (PDAs, computer agents developed by human subjects) as a method for evaluating automated negotiators. We have examined the negotiation results and its dynamics in extensive simulations with more than 300 human negotiators and more than 50 PDAs in two distinct negotiation settings. Results show that computer agents perform better than PDAs in the same negotiation contexts in which they perform better than people, and that on average, they exhibit the same measure of generosity towards their negotiation partners. Thus, we found that using the method of peer designed negotiators embodies the promise of relieving some of the need for people when evaluating automated negotiators.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multiagent systems*

General Terms

Bilateral Negotiations

Keywords

peer designed agents, evaluation techniques, automated negotiation

1. INTRODUCTION

Heterogenous group activities in which people and computers interact are becoming increasingly prevalent. Examples of group activities in which computer systems participate include online auctions, assistive care, and military systems. In particular, the design of automated negotiators that can proficiently negotiate with people is receiving growing attention in AI [1, 8, 9, 10, 11, 14, 19, 25].

Expert Designed Negotiators (EDNs) have been developed for a variety of purposes, whether as autonomous actors [5, 9], as proxies for individual people and organizations [25], or as a training tool [4, 17]. The evaluation process of automated negotiators is a critical aspect of their design process. Traditionally, this task is done by measuring the performance of EDNs when interacting with actual people.

However, using people for experimentation purposes is timely and costly, making the evaluation process a very difficult task for researches. This becomes even harder when the negotiation is taken place in open systems. Designing agents that model the human behavior during negotiations is also a very difficult task, especially due to the diverse behavior of people which makes it difficult to capture it by a monolithic model. People tend to make mistakes, and they are affected by cognitive, social and cultural factors, etc. [12]. Thus, it is commonly assumed that people cannot be substituted in the evaluation process of EDNs designed to negotiate with people. While the design issues of EDNs is important, and can be found in the literature [13], we found little, if any, literature focusing on the important issue of the evaluation of EDNs designed to negotiate with people. Thus, in this paper we only focus on the evaluation process of these EDNs.

The question which now arises is, even though people and agents behave differently, whether one can use agents to evaluate EDNs and reflect from the behavior of the EDNs with other agents to their behavior with people. Following this intuition, we can look at the strategy method [16, 21, 22] which is an experimental methodology which requires people to elicit their actions. The assumption behind this method is that people are able to effectively encapsulate their own strategies if they are properly motivated, monetarily or otherwise. This approach is well accepted in experimental economics and has also begun to be used within artificial intelligence research [2, 20]. The application of this methodology within the study of automated negotiator agents implies that peer designed agents (PDAs) can be developed that represent the negotiation preferences of diverse people. As we will demonstrate in the rest of this paper, the use of PDAs

*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.

has the potential of alleviating some of the need for people in the evaluation of automated negotiators.

In this paper we provide results of extensive experiments involving more than 300 human subjects and 50 PDAs. The results shed light as to the prospect of using PDAs to better determine the proficiency of an automated negotiator when matched with people, as well as to compare the behavior of different EDNs. The experiments involved negotiations of people with people, people with PDAs and people with EDNs to be evaluated. The experiments were conducted in two distinct negotiation environments. Results show that PDAs can alleviate the evaluation process of automatic negotiators, and facilitate their design in a variety of settings.

This paper contributes to research on automated negotiations by tackling the important issue of the evaluation process of automated negotiators. Some of the drawbacks of designing automated negotiators that are capable of negotiating with people stems from the fact that their evaluation involves people and the cumbersome task of orchestrating the experiments with people. Perhaps simplifying this evaluation process could, in turn, make the design of these agents more accessible and cause more agents to be designed for proficiently negotiating with people. Thus the importance of our novel research, which suggests using PDAs as an unbiased mechanism for evaluating automated negotiators, which can reflect on the behavior of people, as well as allowing fine-tuning and improving the strategy of the automated negotiators.

The rest of this paper is organized as follows. In Section 2 we review related work in the field of automated negotiators' evaluation. We continue and describe the problem domain in Section 3. In Section 4 we present the experiments we conducted, our methodology and our evaluation results. Finally, we provide a summary and discuss the results.

2. RELATED WORK

Important differences exist between designing an automated agent that can successfully negotiate with a human counterpart and designing an automated agent to negotiate with other automated agents. As this paper does not focus on the design of such automated agents, we will not survey related work on this topic. A more detailed review on the design of automated agents capable of negotiating with people can be found, for example, in [13].

A straightforward approach for evaluating or designing EDNs that should negotiate with people could have been using equilibrium agents. However, results from social sciences suggest that people do not follow equilibrium strategies [3, 15]. Moreover, when playing with humans, the theoretical equilibrium strategy is not necessarily the optimal strategy [26]. In this respect, equilibrium based automated agents that play with people must incorporate heuristics to allow for "unknown" deviations in the behavior of the other party. For example, Kraus *et al.* [10] report on the design of a perfect equilibrium agent and experiments conducted with people. They observed that the human players do not necessarily follow equilibrium strategies and thus their automated agent performed poorly. Only after adding heuristics and argumentation tailored to their specific settings were they able to achieve effective negotiation by their agent. Hence, it is apparent that matching automated negotiators with other automated agents following equilibrium or other game theory paradigms cannot reflect on the proficiency of the negotiators when negotiating with people. Even when not following equilibrium strategy, automated negotiators might not be suitable for evaluation purposes. For example, Hindriks and Tykhonov [7] showed the limitation of automated negotiators and the deterioration in their performance when deployed in domains other than the ones the designers had initially created them for.

Thus, in order to replace people in the evaluation loop, to date, one cannot rely on specific automated agents. Instead we examine the use of peer designed agent as a type of strategy method. Similarly to developing agents, using the strategy method requires subjects to specify their choices for all information sets of the game and not only the ones that occur during the course of a play of a game [16, 21, 22]. Despite the similarity, asking subjects to design and program agents is different from the strategy method. Developing agents requires subjects to implement much more complex strategies like using heuristics and learning algorithms potentially to make decisions in situations not originally considered. The use of PDAs has been extensively studied within the context of the Trading Agent Competition for Supply Chain Management (TAC SCM) [23]. In TAC one needs to design a trading agent that participates in auctions for certain good. The use of PDA's within this domain demonstrates the benefits of a large set of PDAs for evaluation purposes of EDNs. Yet, in this context, both the PDAs and the EDNs were used for interacting with other computer agents, and not for interaction with people.

Grosz *et al.* [6] experimented with people designing agents for a game called Colored Trails. They observed that when people design agents, they do not always follow equilibrium strategies. Moreover, in their analysis they showed that people demonstrated more helpfulness, which led to higher scores, than their designed agents. Chalamish *et al.* [2] report on large-scale experiments in which people programmed agents which were shown to successfully capture their strategy in a set of simple games. They conclude that peer designed agents can be used instead of people in some cases. In another settings, Rosenfeld and Kraus [20] report on experiments done with PDAs designed for optimization problems. Based on the experiments they conclude that theories of bounded rationality can be used to better simulate people's behavior. However, the settings of Chalamish *et al.* [2] were relatively simple, while our settings have richer strategy space and are much more complicated. Hence, it is not straightforward that PDAs can be used to simulate people's behavior and thus replace them for evaluation purposes.

3. PROBLEM DESCRIPTION

We consider the problem of evaluating the proficiency of EDNs designed to negotiate with people. We consider a general environment of bilateral negotiation in which two agents, either automated negotiators or people, negotiate to reach an agreement on conflicting issues. We consider two distinct bilateral negotiation environments. The first involves a day-to-day scenario in which the parties negotiate to reach an agreement on conflicting goals, while the second involves playing a game. We describe both environments below.

The first negotiation environment involved a multi-attribute multi-issues negotiation settings (see Figure 1). In this environment, 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. 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 [18, p. 118-121], in which each agent can perform up to $M > 0$ interactions with its counterpart in each time period.

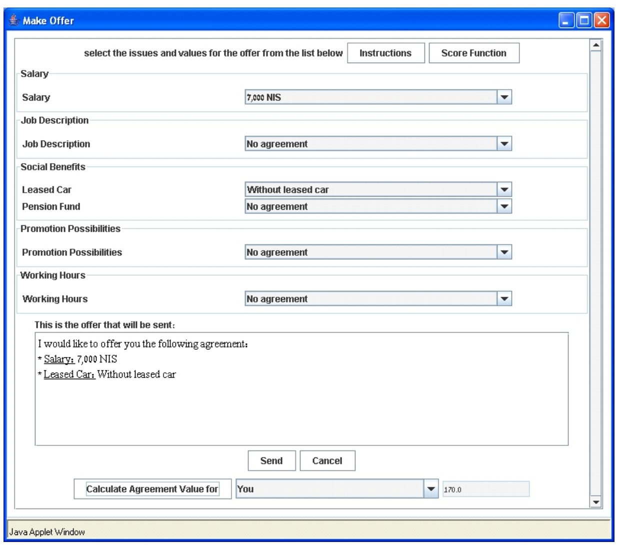

Figure 1: Bilateral Negotiation: Generating offers screen.

In order to make the settings more realistic, it also involved 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). 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 developed a simulation environment which is adaptable such that any scenario and utility function, expressed as multi-issue attributes, can be used, with no additional changes in the configuration of the interface of the simulations or the automated agent. The agents (PDAs or EDNs) 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 agent constructs an offer by choosing the different values constituting the offers. Then, the offer is constructed and sent in plain English to its counterpart.

In this environment we experimented with two state-of-the-art automated negotiators, *KBAgent* and *QOAgent* which were shown by Oshrat *et al.* [19] and Lin *et al.* [14] to negotiate proficiently with people. Both agents are domain independent and apply non-classical decision making method, rather than focusing on maximizing the expected utility. They also apply different learning mechanism to determine the type of their counterpart. Both agents were shown to reach more agreements and played more effectively than their human counterparts, when the effectiveness is measured by the score of the individual utility. Since they were shown to be proficient negotiators with people they can serve as our baseline for evaluating the PDAs as a strategy method for replacing people in the evaluation loop.

The second negotiation environment involved playing the Colored Trails (CT) game [6] which is a game played on a nxm board of colored squares. Players are issued colored chips and are required to move from their initial square to a designated goal square. To move to an adjacent square, a player must turn in a chip of the same color as the square. Players must negotiate with each other to obtain chips needed to reach the goal square (see Figure 2). 100 points are given for reaching the goal square and 10 points bonus are given for each chip left for each agent at the end of the game. If the player did not reach the goal, 15 points penalty are given for each square from its final position to the goal square. Note that in this game, the performance of the agent does not depend on the outcome of the other player. Agreements are not enforceable, allowing players to promise chips but not transferring them. In addition, each player can see the entire game board.

The simulation environment we used in this setting is adaptable such that different variations of the game can be set. The size of the board, number and color of total chips and chips given to each player can be changed. The automated agents can play both sides in the game, while the human counterpart accesses the game via a web address. The game itself is split into turns, where each turn is divided to a negotiation phase, a commitment phase and a movement phase. In the negotiation phase the players can request or promise to send chips. Then, in the commitment phase, the players can send the actual chips or withdraw from the agreement. This might result in one agent sending the promised chips in return to be given other chips, while the other agent fails to deliver. In the movement phase, the players can choose to move to adjacent squares, given they have the required colored chips. The game terminates when either side reaches the goal square or if no player has moved in three consecutive turns.

An EDN was used in this environment. The agent, which is called *PURB*, is *P*ersonality traits, *U*tility and *R*ules *B*ased,

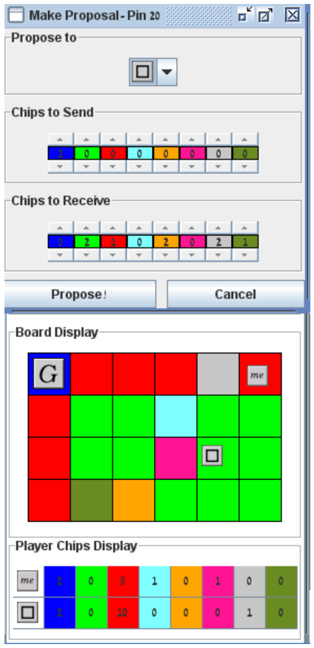

Figure 2: The Colored-Trail game screenshot.

and extends the agent reported by Talman *et al.* [24] to allow it to play proficiently with people. Its behavior is characterized as high reliability and medium generosity. The agent can change its personality level for cooperation and reliability and also model these traits of its opponent. A utility function is used to evaluate each possible action and proposal and randomization is also used to choose between different choices.

4. EXPERIMENTS

The experiments were conducted using the simulation environments mentioned above. The experiments were designed to answer two fundamental questions: (a) whether the behavior of EDNs against PDAs can reflect on their behavior against people?, and (b) whether PDAs can be used to compare the behavior of different EDNs instead of people? We begin by describing the environments which were used in the different experiments and then continue to describe the experimental methodology and results.

4.1 The Negotiation Environments

For the bilateral negotiation environment two domains were used, which were described by Lin *et al.* [14]. The first domain is a Job Candidate domain, which is related to the subjects' experience, and thus they could better identify with it. In this domain, 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. In this scenario, 5 different attributes are negotiable with a total of 1,296 possible agreements that exist. The second domain involved reaching an agreement between England and Zimbabwe evolving from 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". In this domain, 4 different attributes are under negotiation, resulting with a total of 576 possible agreements.

In both domains if an agreement is not reached by the end of the allocated time a status quo agreement is implemented. In addition, time also has an impact and the sides might lose or gain as time progresses. In the Job candidate domain both sides lose as time progresses, while in the England-Zimbabwe domain, England gains while Zimbabwe loses as time progresses. Also, each side can choose to opt out of the negotiation at any time. As there is also incomplete information in each domain, 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 and a compromising orientation.

For the CT game environment, two types of games were used: (a) asymmetric and (b) symmetric. In each game the same 7x5 board was used. Each game differed by the ability of each player to reach the goal square with or without the assistance of the other player. The asymmetric game was characterized by one of the player having 15 chips and being *dependant* of the other player and needed to exchange chips in order to reach the goal square, while the other player had 24 chips and was *independent* of its counterpart, and thus could reach the goal square without doing any exchange. The symmetric game, on the other hand, was characterized by the two players having 24 chips and being dependent and needing each other's chips to reach the goal square.

4.2 Experimental Methodology

We ran an extensive set of simulations, consisting of more than 300 human negotiators and more than 50 PDAs. 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. Each subject served only one specific role in the negotiations (e.g., in the bilateral negotiations either the employer role or the job candidate one, and in the CT game environment either the dependent player or the independent player). 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 score function was private information and unknown for the other side.

The PDAs were automated negotiators designed by people. The students were given a task to implement an efficient automated agent for a given negotiation environment (different students designed PDAs for the different negotiation environments). The implementation was done in the same simulation environment as the negotiation itself. The students were provided skeleton classes, having all the necessary server-communication functionality, 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. The students were able to first negotiate or play the CT game before submitting their PDAs.

4.3 Experimental Results

The main goal of the experiments was to analyze whether the strategy method of PDAs can be used to replace people in the evaluation process of EDNs designed to negotiate proficiently with people. In addition, we wanted to find whether this method can also be used to evaluate and compare different automated negotiators and obtain from it which will be a more proficient negotiator with people.

4.3.1 Evaluating EDNs when Matched with PDAs versus when Matched with People

We start by examining the final outcomes in each experimental setting. Table 1 summarizes the final utilities achieved by each side in each experiment for the Job candidate and England-Zimbabwe domains, while Table 2 summarizes the final utilities in the CT game environment. All results are statistically significant within the $p < 0.05$ range. The first question we address relates to the evaluation of the performance of EDNs and people and whether one can take people out of this evaluation loop. More precisely, if we wish to compare between the behavior of people and EDNs, can we gather results of PDAs-EDNs negotiation and then PDAs-PDAs negotiations instead of gathering results of people-people negotiations and then people-EDN negotiations (as had been done, for example, in [14, 19])?

When the *KBAgent* was matched with PDAs it was able to achieve higher utility values than the average of the PDAs matched against themselves (lines (1),(2) in Table 1). This is also consistent with the *KBAgent*'s achievements when matched with people (lines (3),(4) in Table 1). Similarly, when the *QOAgent* negotiates with PDAs it was able to achieve higher utility values than the average of the PDAs when matched against themselves (lines (2),(5) in Table 1). Again, this was consistent with the *QOAgent*'s achievements when matched with people (lines (4),(6) in Table 1). A similar phenomenon was observed in the CT game. When the *PURB* agent played in the symmetric settings and in the asymmetric game as the independent role, the final utilities achieved by the EDN were higher than the average utilities of the PDAs. When it played the dependent role in the asymmetric game, its final utility was lower than the average utility of the PDAs (lines (1),(2) in Table 2). The same relation is revealed when comparing the *PURB*'s utility when playing with people and the average utilities of people playing with one another (lines (3),(4) in Table 2).

It is also interesting to see whether the performance of two EDNs (the *KBAgent* and the *QOAgent*) when matched with PDAs can be used as a prediction of whom will perform better when matched with people. The *KBAgent* was shown to perform better when matched with people than the *QOAgent* (lines (3),(6) in Table 1). In three out of the four sides in the two domains, this is also reflected when they are matched with the PDAs, with the *KBAgent* achieving higher utility scores than the *QOAgent* (lines (1),(5) in Table 1). To summarize, the experimental results indeed show that using PDAs can help in the evaluation process of EDNs designed to negotiate proficiently with people. The results support the hypothesis that the final utility values can serve as a good indication for evaluating the proficiency of the automated negotiator. Moreover, they can also be used to compare between different EDNs and reflect on their proficiency when matched with people.

		Job Can. Domain		Eng-Zim Domain	
		u_{employer}	$u_{\text{job can}}$	u_{eng}	u_{zim}
(1)	***KBAgent* vs. PDAs**	437.7	415.8	720.0	-14.5
(2)	**PDAs vs. PDAs**	368.2	355.1	251.8	-83.7
(3)	***KBAgent* vs. People**	472.8	482.7	620.5	181.8
(4)	**People vs. People**	423.8	328.9	314.4	-160.0
(5)	***QOAgent* vs. PDAs**	466.1	396.8	663.4	-36.5
(6)	***QOAgent* vs. People**	417.4	397.8	384.9	35.3

Table 1: Final utility results in the bilateral negotiation environment.

		Asymmetric game		Symmetric game
		$u_{\text{independent}}$	$u_{\text{dependent}}$	$u_{\text{dependent}}$
(1)	***PURB* vs. PDAs**	180.53	35.00	131.36
(2)	**PDAs vs. PDAs**	178.38	45.25	111.48
(3)	***PURB* vs. People**	187.08	81.94	157.83
(4)	**People vs. People**	181.45	97.26	130.67

Table 2: Final utility results in the CT game environment.

We also observed similarity in the measure of generosity the EDNs exhibited towards their negotiation partners. We say that the agent is *generous* if it proposes an offer which utility is higher than the previous offer it proposed for the other side, regardless of its value for the proposing agents. Similarly, we say that an agent is *selfish* if it proposes an offer which increases its own utility while decreasing the other side's utility, as compared to the previous proposal. When comparing the generosity and selfishness rates of the *KBAgent* and the *QOAgent* we found a correlation in the rates when they were matched with PDAs, as compared to when they were matched with people (generosity (selfishness) rates of 88% (5.9%) for the *KBAgent* as compared to 85% (11.3%) of the *QOAgent* when matched with people and 88% (5.56%) as compared to 71% (22.83%) when matched with PDAs).

4.3.2 Evaluating the Performance and Behavior of People versus PDAs

PDAs behave differently than people (e.g., [6, 20]), yet our experiments indicated that playing against PDAs can reflect on the results when the EDNs are matched with people. To bolster our confidence from these results we examined the pattern of behavior demonstrated people and with PDAs when matched with the EDNs. Due to space limitation we only present the results on one of the domains and negotiation's sides, the results are similar in the other domains and sides. Figure 3 compares the final utilities achieved by

PDAs and people when matched with the EDNs in the job candidate domain when playing the role of the employer, while Figure 4 compares the times in which the negotiations terminated. Note, that we compare between the behavior of people and PDAs and not the EDNs behavior. The results demonstrate the similarity between people and PDAs when matched with EDNs. For example, in Figure 3 we can observe that PDAs achieve somewhat higher utilities when matched with the *QOAgent* as compared to the *KBAgent*. The same trend is then observed when people are matched with both agents.

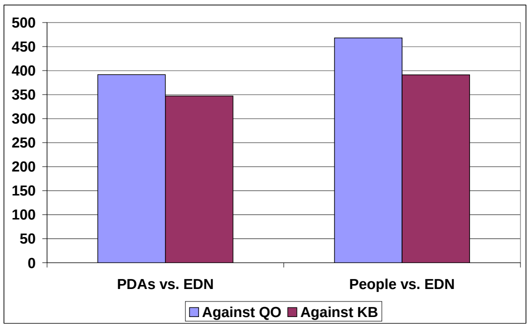

Figure 3: Comparing the final utility results of people and PDAs when matched with the EDNs.

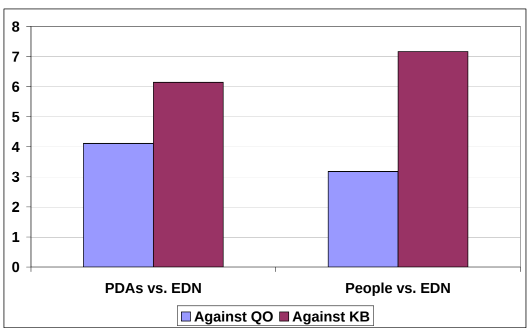

Figure 4: Comparing end turns of EDNs' negotiations when matched with people and PDAs.

5. CONCLUSIONS

The importance of designing proficient automated negotiators to negotiate with people and evaluating them cannot be overstated. Yet, evaluating agents against people is a tiresome task, due to the cost and time required. In this paper we presented an extensive systematic experimentation to answer the question whether people can be kept out of the evaluation loop when evaluating automated negotiators designed specifically to negotiate proficiently with people. To do so, we evaluated several negotiation behavioral parameters in an extensive set of experiments with people and with peer designed agents. In the bottom line, our results reveal that playing well against peer designed agents can reflect on the proficiency of the automated negotiator when matched with people. Moreover, we showed that while PDAs results with different negotiation outcomes than people, there is a common behavioral pattern when they are both matched with EDNs.

There are fundamental benefits of using PDAs instead of people. First, PDAs are accessible 24/7 and can be used whenever needed. In addition, PDAs are not biased and thus can be used several times to asses the EDN's behavior. Thus, they allow the agent designer to revised and change her agent with the ability to evaluate each design and compare it to previous designs. Lastly, it allows different EDNs to be matched on the same set of PDAs and obtain an objective evaluation of the results.

While people cannot be kept completely out of the evaluation loop, we demonstrated the promise embodied in peer designed agents for evaluation purposes of automated negotiators. Thus, evaluating on peer designed agents could and should serve as a first extensive attempt to validate the agent's proficiency and strategy design before continuing on to evaluation with people.

As noted, people behave differently than PDAs. Future work warrants careful investigation on the differences between the behavior of PDAs and people. This investigation might allow for a better understanding of people and the better design of automated agents specifically designed to negotiate with people.

6. ACKNOWLEDGMENTS

We thank Ya'akov (Kobi) Gal for his helpful remarks and suggestions.

7. REFERENCES

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

[2] M. Chalamish, D. Sarne, and S. Kraus. Programming agents as a means of capturing self-strategy. In *AAMAS*, pages 1161–1168, 2008.

[3] I. Erev and A. Roth. Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibrium. *American Economic Review*, 88(4):848–881, 1998.

[4] M. Fleming, D. Olsen, H. Stathes, L. Boteler, P. Grossberg, J. Pfeifer, S. Schiro, J. Banning, and S. Skochelak. Virtual reality skills training for health care professionals in alcohol screening and brief intervention. *Journal of the American Board of Family Medicine*, 22(4):387–398, 2009.

[5] Y. Gal, A. Pfeffer, F. Marzo, and B. J. Grosz. Learning social preferences in games. In *AAAI*, pages 226–231, 2004.

[6] B. Grosz, S. Kraus, S. Talman, and B. Stossel. The influence of social dependencies on decision-making: Initial investigations with a new game. In *AAMAS*, pages 782–789, 2004.

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

[8] 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.

[9] R. Katz and S. Kraus. Efficient agents for cliff edge environments with a large set of decision options. In

AAMAS, pages 697–704, 2006.

[10] 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.

[11] S. Kraus and D. Lehmann. Designing and building a negotiating automated agent. *Computational Intelligence*, 11(1):132–171, 1995.

[12] D. A. Lax and J. K. Sebenius. Thinking coalitionally: party arithmetic, process opportunism, and strategic sequencing. In H. P. Young, editor, *Negotiation Analysis*, pages 153–193. The University of Michigan Press, 1992.

[13] R. Lin and S. Kraus. Can automated agents proficiently negotiate with humans? *CACM*, 53(1):78–88, January 2010.

[14] 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.

[15] R. D. McKelvey and T. R. Palfrey. An experimental study of the centipede game. *Econometrica*, 60(4):803–836, 1992.

[16] T. Offerman, J. Potters, and H. A. A. Verbon. Cooperation in an overlapping generations experiment. *Games and Economic Behavior*, 36(2):264–275, 2001.

[17] D. E. Olsen. Interview and interrogation training using a computer-simulated subject. In *Interservice/Industry Training, Simulation and Education Conference*, 1997.

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

[19] Y. Oshrat, R. Lin, and S. Kraus. Facing the challenge of human-agent negotiations via effective general opponent modeling. In *AAMAS*, pages 377–384, 2009.

[20] A. Rosenfeld and S. Kraus. Modeling agents through bounded rationality theories. In *IJCAI*, pages 264–271, 2009.

[21] R. Selten, K. Abbink, J. Buchta, and A. Sadrieh. How to play (3x3)-games: A strategy method experiment. *Games and Economic Behavior*, 45(1):19–37, 2003.

[22] R. Selten, M. Mitzkewitz, and G. R. Uhlich. Duopoly strategies programmed by experienced players. *Econometrica*, 65(3):517–556, 1997.

[23] TAC Team. A trading agent competition. *IEEE Internet Computing*, 5(2):43–51, 2001.

[24] S. Talman, M. Hadad, Y. Gal, and S. Kraus. Adapting to agents' personalities in negotiation. In *AAMAS*, pages 383–389, 2005.

[25] D. Traum, S. Marsella, J. Gratch, J. Lee, and A. Hartholt. Multi-party, multi-issue, multi-strategy negotiation for multi-modal virtual agents. In *IVA*, pages 117 – 130, 2008.

[26] A. Tversky and D. Kahneman. The framing of decisions and the psychology of choice. *Science*, 211:453–458, 1981.