

# Facilitating the Evaluation of Automated Negotiators using Peer Designed Agents\*

Raz Lin<sup>1</sup>, Sarit Kraus<sup>1,2</sup> and Yinon Oshrat<sup>1</sup>

<sup>1</sup> Department of Computer Science  
Bar-Ilan University  
Ramat-Gan, Israel 52900

<sup>2</sup> Institute for Advanced Computer Studies  
University of Maryland  
College Park, MD 20742 USA  
{linraz, sarit}@cs.biu.ac.il

Ya'akov (Kobi) Gal

Department of Information Systems Engineering  
Ben-Gurion University of the Negev  
Beer-Sheva, Israel 84105  
kobig@bgu.ac.il

## Abstract

Computer agents are increasingly deployed in settings in which they make decisions with people, such as electronic commerce, collaborative interfaces, and cognitive assistants. However, the scientific evaluation of computational strategies for human-computer decision-making is a costly process, involving time, effort and personnel. This paper investigates the use of Peer Designed Agents (PDA)—computer agents developed by human subjects—as a tool for facilitating the evaluation process of automatic negotiators that were developed by researchers. It compares the performance between automatic negotiators that interacted with PDAs to automatic negotiators that interacted with actual people in different domains. The experiments included more than 300 human subjects and 50 PDAs developed by students. Results showed that the automatic negotiators outperformed PDAs in the same situations in which they outperformed people, and that on average, they exhibited the same measure of generosity towards their negotiation partners. These patterns occurred for all types of domains, and for all types of automated negotiators, despite the fact that there were individual differences between the behavior of PDAs and people. The study thus provides an empirical proof that PDAs can alleviate the evaluation process of automatic negotiators, and facilitate their design.

## Introduction

Heterogeneous group activities in which people and computers interact are becoming increasingly prevalent. Examples of group activities in which computer systems participate include on-line 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.

Expert Designed Negotiators (EDNs) have been developed for a variety of purposes, whether as autonomous actors (Gal et al., 2004; Katz and Kraus, 2006), as proxies for individual people and organizations (Traum et al., 2008), or as a training tool (Fleming et al., 2009; Olsen, 1997; Lin and

Kraus, 2010). 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. For a recent survey of that describes studies that evaluate automatic agents that negotiate with people, see Lin and Kraus (2010).

Using people for the evaluation purposes is costly in terms of time, effort and money, making for a difficult task for researchers. The traditional view was that people cannot be replaced in the evaluation process, because of their diverse behavior that is affected by cognitive, social and cultural factors (Lax and Sebenius, 1992; Camerer, 2003).

This paper suggests a new paradigm for evaluating EDNs by comparing their performance to that of computer agents that were designed by other people, called Peer Designed Agents (PDA). We hypothesized that the aggregate performance of PDAs would reflect that of people, and EDNs that negotiate successfully with PDAs would also negotiate successfully with people.

This paradigm was inspired by the “strategy method” commonly used in behavioral economics, which is an experimental methodology that elicits from people a complete strategy specification. People state their actions for every possible situation that may arise in their interaction (Offerman, Potters, and Verbon, 2001; Selten, Mitzkewitz, and Uhlich, 1997; Selten et al., 2003). The assumption behind this method is that people are able to effectively encapsulate their own strategies if they are properly motivated, monetarily or otherwise.

The strategy method has also begun to be used within artificial intelligence research (Chalamish, Sarne, and Kraus, 2008; Rosenfeld and Kraus, 2009) and in electronic commerce applications such as eBay, where people can design automatic bidders to represent themselves in negotiation. As we will show in this paper, PDAs can represent the negotiation preferences of diverse types of people, and their use has the potential of alleviating some of the need for people in the evaluation of automated negotiators.

This study includes results from extensive experiments involving more than 300 human subjects and 50 PDAs. The experiments involved negotiations of people that interact with other people, people that interact with PDAs and people that interact with EDNs. The experiments were conducted in two different negotiation settings that simulate real-world

\*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.

Copyright © 2010, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

scenarios that require negotiators to reach agreement about the exchange of resources in order to complete their goals.

Results show that PDAs can be used as a mechanism for evaluating automated negotiators that reflects on the behavior of people, as well as to fine-tune and improve the strategy of the automated negotiators. In particular, EDNs were able to outperform people in our experiments in the same situations in which EDNs were able to outperform PDAs. These results suggest that PDAs can be used to better determine the proficiency of an automated negotiator when matched with people, as well as to compare between the performance of different EDNs and their effects on human behavior. Thus, this paper directly contributes to research on agent-design for human-computer negotiation by facilitating the evaluation process of automated negotiators.

## Related Work

A straightforward approach for evaluating or designing EDNs that should negotiate with people could have been using agents that use pre-defined strategies such as equilibrium- or heuristic-based agents. However, results from social sciences suggest that people do not follow equilibrium strategies (Erev and Roth, 1998; McKelvey and Palfrey, 1992; Camerer, 2003). Moreover, when playing with humans, the equilibrium strategy may not be optimal (Tversky and Kahneman, 1981). Recent works have adapted automated agents to incorporate heuristics that can capture deviations from equilibrium. For example, Kraus et al. (2008) observed that human players do not necessarily negotiate according to equilibrium strategies and their automated agent was able to achieve better performance after adding heuristics that allowed it to deviate from the equilibrium path. Hence, matching automated negotiators with other automated agents that follow equilibrium or other game theory paradigms cannot reflect on the proficiency of the negotiators when negotiating with people.

Grosz et al. (2004) 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, Sarne, and Kraus (2008) 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 (2009) 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. Our paper extends these works to a richer strategy space, and directly compares the extent to which people's performance against EDNs matches the performance of PDAs against EDNs.

The use of PDAs is directly related to the strategy method used in behavioral economics, in that subjects need to specify an action for every possible situation that may occur in the negotiation process (Offerman, Potters, and Verbon, 2001; Selten, Mitzkewitz, and Uhlich, 1997; Selten et al.,

2003). This is essentially a look-up table that matches an action to each possible history in the interaction. However, PDAs differ from the strategy method in that the strategy they prescribe may be the result of algorithms and computational processes.

Finally, we note that 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. PDAs have been studied in settings such as the Trading Agent Competition for Supply Chain Management (TAC SCM) (TAC Team, 2001) and Robocup (Asada et al., 1998) in which agents were designed to interact with other computer agents. This is the first work that studies the use of PDAs to evaluate EDNs for negotiating with people.

## Problem Description

We consider two different settings of bilateral negotiation in which participants, either automated negotiators or people, negotiate to reach an agreement on conflicting issues.

The first setting involved a multi-issue negotiation setting in which participants engage in repeated negotiation rounds in which they can propose an agreement which consists of values for a subset or all of the issues in the negotiation. The negotiation protocol proceeds as follows: At each time period each participant can propose a possible agreement for some subset of the issues (see Figure 1), and the other participant can either accept the offer, reject it or opt out. The protocol is an extension of the classic alternating offers protocol of Osborne and Rubinstein (1994, p. 118-121) to support partial agreements as well as an opting-out option. The setting was described to people using one of two possible scenarios: an employer and an employee negotiating over terms of a job contract, or as a diplomatic negotiation process between two countries.

The negotiation terminates when one of the following holds: (a) the negotiators reach agreement for all of the issues, (b) one of the participants opts out, thus forcing the termination of the negotiation with a default outcome, 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 is implemented. The utility for participants depends on the agreement and the time period in which it was reached.

We assume that there is a finite set of negotiator types and that negotiators do not know each other's 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.

In this setting we experimented with two EDNs called the *KBAgent* and the *QOAgent* which have been shown to negotiate proficiently with people (Oshrat, Lin, and Kraus, 2009; Lin et al., 2008). Both agents combine instance-based learning with non-classical decision-making methods to negotiating with people.

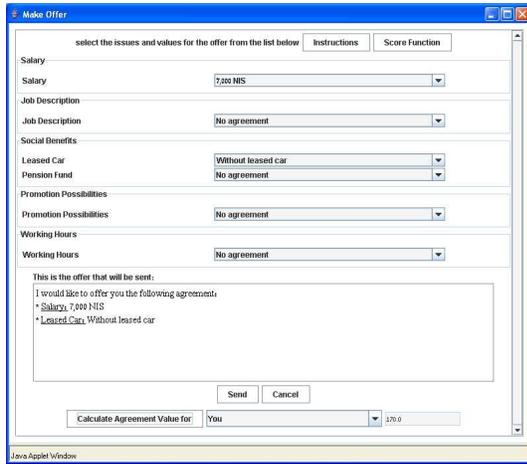
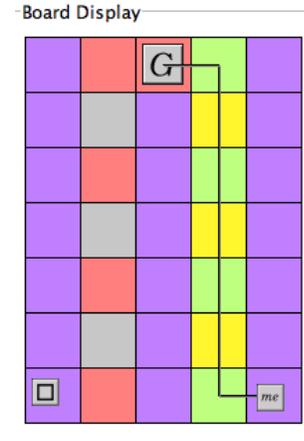


Figure 1: Bilateral Negotiation: Generating offers screen.

The second negotiation setting involved playing the Colored Trails (CT) game (Grosz et al., 2004) which is a general negotiation test-bed that provides an analogy to task-settings in the real-world.<sup>1</sup> CT is played on a 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). Each participant obtains a score at the end of the game that is computed as follows: 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 transfer phase and a movement phase. In the negotiation phase the players can request or promise to send chips. Then, in the transfer phase, the players choose the extent to which they meet their commitments by deciding the amount of chips to send. 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 termi-

<sup>1</sup>Colored Trails is Free Software and can be downloaded at <http://www.eecs.harvard.edu/ai/ct>



(a) Board Panel

Player Chips Display

	1	0	13	10	0
	5	0	10	0	9

(b) Chip Display Panel (showing the chips in the possession of both participants)

Figure 2: Snapshots of Colored Trails GUI

nates when either side reaches the goal square or if any of the player have not moved in three consecutive turns.

The EDN that was used, called the Personality and Utility Rule Based (PURB) agent combined a social utility function that represented the behavioral traits of other participants, as well as a rule-based mechanism that used the utility function to make decisions in the negotiation process.

## Empirical Methodology

The following experiments were performed to answer two fundamental questions: (a) whether the behavior of EDNs negotiating with PDAs reflects their behavior negotiating with people, and (b) whether PDAs can be used in lieu of people to compare the behavior of different EDNs. We begin by describing the settings which were used in the different experiments and then continue to describe the experimental methodology and results.

## The Negotiation Settings

Two different settings were used (for complete descriptions refer to Lin et al. (2008)). The first setting was a Job Candidate domain. In this domain, a negotiation takes place between an employer and a job candidate. In the negotiation both the employer and the job candidate needed to formalize the hiring terms and conditions of the applicant. In this scenario, five 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. In this domain, 4 different attributes are under negotiation, resulting with a total of 576 possible agreements.

In both domains status quo agreement was implemented if an agreement was not reached by a certain deadline. 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 we used a  $7 \times 5$  board. There were two types of board games that were used: (a) asymmetric and (b) symmetric. In each game the same  $7 \times 5$  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.

We ran an extensive set of simulations, consisting of more than 300 human negotiators and more than 50 PDAs. The human negotiators were current or former computer science undergraduate and graduate students. 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 either the dependent player or the independent player). Prior to the experiments, the subjects were given identical instructions regarding the experiment and the domain. The subjects were instructed to achieve the best possible agreement. They were told that performance in the game depended on their own score but not on the score of the other participant.

The PDAs were automated negotiators designed by different people than the ones that played the EDNs. The students were given a task to implement an automated agent for a given negotiation setting. Students did not take a prior course in decision- or game-theory. The implementation was done in the same setting as the people who played the EDNs. The students were provided skeleton classes and APIs to facilitate the implementation of 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 the agent strategies that they developed.

## Results and Discussion

The 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.

**Evaluating EDNs when Matched with PDAs versus when Matched with People** We begin by examining the final outcomes in each experimental setting. Table 1 summarizes the average performance achieved by each side in each experiment for the Job candidate and England-Zimbabwe domains, while Table 2 summarizes the average performance in the CT game. 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, we see whether the comparison between the play of PDAs versus EDNs with PDAs versus PDAs is similar to the comparison between people versus people and people versus EDN negotiations (as had been done, for example, by Lin et al. (2008); Oshrat, Lin, and Kraus (2009)).

When the *KBAgent* negotiated 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* negotiated 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). This was consistent with the *QOAgent*'s achievements when matched with people (lines (4),(6) in Table 1), with one exception: in the employer role, the *QOAgent* outperformed PDA, but did not outperform people in the same role.

In the CT game, results were consistent with the pattern described above: When the *PURB* agent played in the symmetric settings and in the asymmetric game as the independent role, the performance achieved by the EDN were higher than the average utilities of the PDAs. When it played the dependent role in the asymmetric game, its performance was lower than the average utility of the PDAs (lines (1),(2) in Table 2). The same pattern is apparent 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).

Interestingly, the performance of two EDNs (the *KBAgent* and the *QOAgent*) when matched with PDAs indicated 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 were 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 confirmed the hypothesis that in conditions where an EDN outperform PDAs,

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

Table 1: Performance in the bilateral negotiation.

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

Table 2: Performance in the CT game.

the EDN will also outperform people. In addition, we showed that our results could also be used to compare between different EDNs and reflect on their proficiency when matched with people.

We also compared the extent to which the EDNs and PDAs exhibited similar behavioral traits during negotiation, such as generosity and selfishness. We say that an 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. We found that both EDNs exhibited similar behavioral traits 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).

**Evaluating the Performance and Behavior of People versus PDAs** On an individual basis, it has been shown that PDAs negotiate differently than people (e.g., Grosz et al. (2004); Rosenfeld and Kraus (2009)), yet our experiments indicated that in the aggregate, PDAs can reflect on the results when the EDNs are matched with people. We examined the pattern of behavior demonstrated by people and 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 performance 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 behavior of EDNs. The results demonstrate the similarity between people and PDAs when matched with EDNs. For example, in Figure 3 we can observe that PDAs achieve 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.

## 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. Our results revealed 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.

The three EDNs we used in the study used different decision-making paradigms (e.g., learning vs. decision-theory) and were evaluated on different domains, (e.g., complete vs. incomplete information) using hundreds of subjects. At least one of these domains have already been shown in the past to reflect task settings in the real world (Gal et al., 2007). The paper shows that on average their performance is representative of people’s behavior in the following qualitative (not quantitative) sense: An EDN that outperforms PDAs in score will also outperform people in score. We do not claim that the extent to which an EDN outperforms PDAs can be used to predict its performance against people.

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 assess 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 ne-

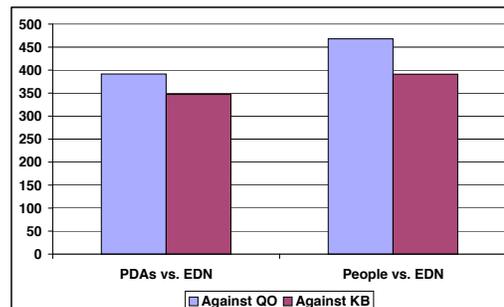


Figure 3: Comparing overall performance between people and PDAs when matched with EDNs.

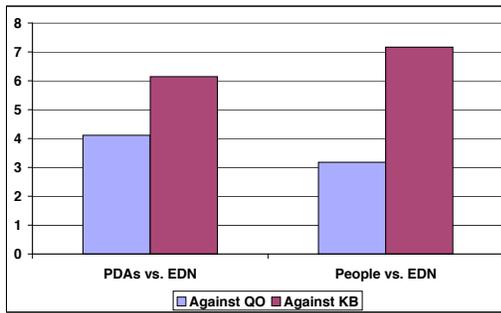


Figure 4: Comparing end-turn performance of EDNs when matched with people and PDAs.

gotiators. 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.

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, make the design of these agents more accessible and cause more agents to be designed for proficiently negotiating with people.

## References

- Asada, M.; Stone, P.; Kitano, H.; and Drogoul, A. 1998. The RoboCup physical agent challenge: Goals and protocols for phase I. *Lecture Notes in Computer Science* 1395.
- Camerer, C. F. 2003. In *Behavioral Game Theory. Experiments in Strategic Interaction*. Princeton University Press. chapter 2.
- Chalamish, M.; Sarne, D.; and Kraus, S. 2008. Programming agents as a means of capturing self-strategy. In *AA-MAS*, 1161–1168.
- Erev, I., and Roth, A. 1998. Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibrium. *American Economic Review* 88(4):848–881.
- Fleming, M.; Olsen, D.; Stathes, H.; Boteler, L.; Grossberg, P.; Pfeifer, J.; Schiro, S.; Banning, J.; and Skochelak, S. 2009. 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.
- Gal, Y.; Pfeffer, A.; Marzo, F.; and Grosz, B. J. 2004. Learning social preferences in games. In *AAAI*, 226–231.
- Gal, Y.; Grosz, B.; Shieber, S.; Pfeffer, A.; and Allain, A. 2007. The effects of task contexts on the decision-making of people and computers. In *CONTEXT*.
- Grosz, B.; Kraus, S.; Talman, S.; and Stossel, B. 2004. The influence of social dependencies on decision-making: Initial investigations with a new game. In *AAMAS*, 782–789.
- Katz, R., and Kraus, S. 2006. Efficient agents for cliff edge environments with a large set of decision options. In *AA-MAS*, 697–704.
- Kraus, S.; Hoz-Weiss, P.; Wilkenfeld, J.; Andersen, D. R.; and Pate, A. 2008. Resolving crises through automated bilateral negotiations. *AIJ* 172(1):1–18.
- Lax, D. A., and Sebenius, J. K. 1992. Thinking coalitionally: party arithmetic, process opportunism, and strategic sequencing. In Young, H. P., ed., *Negotiation Analysis*. The University of Michigan Press. 153–193.
- Lin, R., and Kraus, S. 2010. Can automated agents proficiently negotiate with humans? *CACM* 53(1):78–88.
- Lin, R.; Kraus, S.; Wilkenfeld, J.; and Barry, J. 2008. Negotiating with bounded rational agents in environments with incomplete information using an automated agent. *AIJ* 172(6-7):823–851.
- McKelvey, R. D., and Palfrey, T. R. 1992. An experimental study of the centipede game. *Econometrica* 60(4):803–836.
- Offerman, T.; Potters, J.; and Verbon, H. A. A. 2001. Cooperation in an overlapping generations experiment. *Games and Economic Behavior* 36(2):264–275.
- Olsen, D. E. 1997. Interview and interrogation training using a computer-simulated subject. In *IITSEC*.
- Osborne, M. J., and Rubinstein, A. 1994. *A Course In Game Theory*. Cambridge MA: MIT Press.
- Oshrat, Y.; Lin, R.; and Kraus, S. 2009. Facing the challenge of human-agent negotiations via effective general opponent modeling. In *AAMAS*, 377–384.
- Rosenfeld, A., and Kraus, S. 2009. Modeling agents through bounded rationality theories. In *IJCAI*, 264–271.
- Selten, R.; Abbink, K.; Buchta, J.; and Sadrieh, A. 2003. How to play (3x3)-games: A strategy method experiment. *Games and Economic Behavior* 45(1):19–37.
- Selten, R.; Mitzkewitz, M.; and Uhlich, G. R. 1997. Duopoly strategies programmed by experienced players. *Econometrica* 65(3):517–556.
- TAC Team. 2001. A trading agent competition. *IEEE Internet Computing* 5(2):43–51.
- Traum, D.; Marsella, S.; Gratch, J.; Lee, J.; and Hartholt, A. 2008. Multi-party, multi-issue, multi-strategy negotiation for multi-modal virtual agents. In *IVA*, 117 – 130.
- Tversky, A., and Kahneman, D. 1981. The framing of decisions and the psychology of choice. *Science* 211:453–458.