

# Human-Agent Decision-making: Combining Theory and Practice

Sarit Kraus

Department of Computer Science  
Bar-Ilan University  
Ramat-Gan Israel  
sarit@cs.biu.ac.il

Extensive work has been conducted both in game theory and logic to model strategic interaction. An important question is whether we can use these theories to design agents for interacting with people? On the one hand, they provide a formal design specification for agent strategies. On the other hand, people do not necessarily adhere to playing in accordance with these strategies, and their behavior is affected by a multitude of social and psychological factors. In this paper we will consider the question of whether strategies implied by theories of strategic behavior can be used by automated agents that interact proficiently with people. We will focus on automated agents that we built that need to interact with people in two negotiation settings: bargaining and deliberation. For bargaining we will study game-theory based equilibrium agents and for argumentation we will discuss logic-based argumentation theory. We will also consider security games and persuasion games and will discuss the benefits of using equilibrium based agents.

## 1 Introduction

Agents that interact proficiently with people may be useful for training [26], supporting [23, 12, 11, 1] and even replacing people in many applications [10, 22]. We are considering the agent-human interactions as being a *strategic* activity [14]. That is, we assume that when the automated agent engages in the interaction, it should act as best it can to realize its preferences. *Game theory* is the mathematical theory of strategic decision-making [28] and thus it seems that game theory might be an appropriate analytical tool for understanding how a strategic agent can and should act, and might also be useful in both the design of automated agents and protocols for the interactions. However, game theory assumes that all players will act as best they can to realize their preferences. Unfortunately, humans tend to make mistakes, and they are affected by cognitive, social and cultural factors [8, 25, 4]. In particular, people's observed behavior does not correspond to game theory-based equilibrium strategies [13, 29].

Another approach for the development of automated agents is the (non-classical)-logic approach. The agent is given a logical representation of its environment and its desired goals, and it reasons logically in order to generate its activities. When interacting with people, the environment consists also of the human model. Yet, modeling people's behavior is a big challenge. We have incomplete information about the person's preferences, and we have to cope with the uncertainties inherent in human decision-making and behavior. Human behavior is diverse, and cannot be satisfactorily captured by a simple abstract model. In particular, human decision-making tends to be very noisy: a person may make different strategic decisions in similar situations.

In this paper we survey briefly a few of the agents that we built over the years that interact proficiently with people. In most of the cases, deploying only a game-theory approach or logical-based approach was not beneficial. Heuristics and machine-learning techniques were augmented into the formal

Table 1: Multi-issue negotiations

Settings	Agent Name	Agent Properties	Scenarios	Significance vs people
Bilateral, single-issue, full information, complex actions, agreements not enforceable	EQH	SPE with manually designed heuristics	fishing dispute	One role
Bilateral, uncertainty, multi-issue	QO-agent	Qualitative decision-making Non-deterministic behavior	job interview tobacco	One role
	KBAgent	Machine learning, qualitative decision-making non-deterministic behavior	job interview tobacco	Both roles
	NegoChat	KBagent algorithms, AAT Anchoring, NLP module	job interview	Both roles

models to lead to agents that interact proficiently with people. We will discuss three negotiation settings: multi-issue negotiations, games where the players interleave negotiations with resource exchange while attempting to satisfy their goals and argumentation settings. Finally, we will discuss security games.

## 2 Multi-issue negotiations

Over the years we designed and implemented several automated agents for multi-issue negotiations. In multi-issue negotiations the players need to reach an agreement on several issues. Each issue is associated with a set of possible values and the players need to agree on a specific value for each issue. The negotiations can end with the negotiators signing an agreement or with one of the sides opting out of the negotiations. In addition, if the crisis does not end within a pre-specified deadline then the status quo is implemented. Each outcome of the negotiations is associated with a utility score for both players. A summary of our agents is presented in Table 1.

### 2.1 EQH agent

The first agent, EQH, that we developed was for crisis scenarios and the setting was quite complex [24]. In addition to the message exchange in a semi-structured language, players could take actions during the negotiations and agreements were not enforceable. In particular, opting out in a crisis is a stochastic action and thus the agents are uncertain about the result. In addition to the main issue of the negotiation or opting out, there are various other parameters of an agent's action. These parameters influence the utility of the negotiators from the crisis. Time plays an important role in the crisis [42]. The specific scenario we used for the experimental study was a fishing dispute between Canada and Spain. We formalized the crisis scenario as a game and identified a subgame-perfect equilibrium. We ran preliminary experiments when the automated agent followed its subgame perfect equilibrium strategy. However, the human negotiators who negotiated with it became frustrated and the negotiation often ended with no agreement. The frustration of the human negotiators was mainly due the lack of flexibility of the agent. Since the proposed and accepted agreements of the subgame perfect negotiation did not

change over the negotiation time, the agent did not compromise.

To address this limitation of the equilibrium-based agent, we incorporated several heuristics to the EQH agent. We allowed the owner of the agent to determine the way the agent will deviate from the equilibrium strategies by determining parameters that influence the agent's behavior which are instantiated before the beginning of negotiations. In order to provide the agent with some flexibility when playing against people, we allowed the agent to consent to agreements that have a lower utility than it would have obtained according to the relevant subgame perfect equilibrium strategy agreement. Therefore we added the margin parameter that determines the largest number of points lower than the desired utility value to which the agent will agree.

An additional parameter is the number of negotiation units by which the agent will increase or decrease its first offer from the agreement specified in its equilibrium strategy. Human negotiators usually begin negotiations with an offer higher (or lower, depending on the negotiator's role) than the value they would eventually like to reach at the end of negotiations. This leaves bargaining space and our agent uses this type of strategy. Another parameter indicates whether the agent will send the first message in the negotiation or will wait for its opponent to make the first offer. The default value of this parameter, following some literature recommendations [16], was that the agent will send the first offer, since we wanted a trigger to initiate negotiations with the other agent.

Another heuristic concerns opting out. Given our assumptions, while rational agents will not opt out, people may opt out. If the agent's expected utility from opting out is higher than its expected utility from its opponent opting out, it will try to predict whether its opponent is going to opt out. If so, it will opt out first. The heuristic for the prediction of whether an opponent will opt out is based on the messages sent by the opponent. For example, when a threatening message is received, or when a comment message indicating that the negotiations are heading in a dangerous direction is received, the estimation that the opponent may opt out increases.

We ran extensive experiments for evaluating the equilibrium agent with the heuristics (EQH agent) [24]. We compared the results of the humans to those of the agents and concluded that the EQH agent received a higher utility score playing both roles, but the results were only statistically significant when the agent played just one of the roles. Furthermore, when an agent participates in a negotiation, the sum of the utilities are significantly higher than when two humans play since the agent always proposes Pareto-optimal offers while people reach agreements that are not.

While the EQH agent was based on the subgame perfect equilibrium strategies, it required the introduction of many heuristics, and its success compared with people was only in one role. The main open question is whether it is possible to provide formal methodology that will lead to an agent that is similar to the EQH without the need to manually design the EQH heuristics. Furthermore, we are aiming for an agent that can achieve a significantly higher utility score than people in both roles. Toward this challenges, we next tried to use a qualitative approach, to introduce incomplete information into the environment and to improve the agent's results in both roles.

## 2.2 QOagent and KBagent

The QOagent was designed to interact with people in environments of bilateral negotiations with incomplete information when the agreements consist of multiple issues [27]. With respect to incomplete information, each negotiator keeps his preferences private, though the preferences might be inferred from the actions of each side (e.g., offers made or responses to offers proposed). Incomplete information is expressed as uncertainty regarding the utility preferences of the opponent, and it is assumed that there is a finite set of different negotiator 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). Lastly, the negotiation is conducted once with each opponent. The experiments were run on two distinct domains. In the first domain, England and Zimbabwe negotiate in order to reach an agreement evolving from the World Health Organization's Framework Convention on Tobacco Control, the world's first public health treaty. In the second domain a negotiation takes place after a successful job interview between an employer and a job candidate.

We first formalized the scenario as a Bayesian game and computed the Bayesian Nash equilibrium. Though we did not run simulations of the Bayesian Nash equilibrium agent against human negotiators, we ran two human negotiations. We found out that the opponent's utility score from the offers suggested by the equilibrium agent are much lower than the final utility values of the human negotiations. By also analyzing the simulation process of the human negotiations, we deduced that without incorporating any heuristics into the equilibrium agent, the human players would not have accepted the offers proposed by it which will lead to low utility scores for the equilibrium agent, similar to the low score of the equilibrium agent in the fishing dispute.

Therefore, we developed the QOAGENT. For the decision-making process, the approach used by the QOAGENT tries to take the utility of both sides into consideration. While the QOAGENT's model applies utility functions, it is based on a non-classical decision-making method, rather than focusing on maximizing the expected utility: the maximin function and a qualitative valuation of offers. Using these methods, the QOAGENT generates offers and decides whether to accept or reject proposals it has received. As for incomplete information, the QOAGENT tackles this problem using a simple Bayesian update mechanism. After each action, this mechanism tries to infer which negotiator type best suits the opponent.

The effectiveness of this method was demonstrated through extensive empirical experiments by [27].

The results of the experiments showed that the automated agent achieved higher utility scores than the human counterpart. This can be explained by the nature of our agent both in reference to accepting offers and generating offers. Using the decision-making mechanism we allow the agent to propose agreements that are good for it, but also reasonable for its opponent. In addition, the automated agent makes straightforward calculations. It evaluates the offer based on its attributes, and not based on its content. In addition, it also places more weight on the fact that it loses or gains as time advances. This is not the case, however, when analyzing the logs of the people. It seems that people put more weight on the content of the offer than on its value. This was more evident in the Job Candidate domain with which the human subjects could more easily identify. Yet, this does not explain why, in both domains, similar to the EQH agent experiments, these results are significant only for one of the sides. In the England-Zimbabwe domain, the results are significant when the agent played the role of England, while in the Job Candidate domain these results are significant when it played the role of the job candidate.

In order to improve the QOAGENT, we extended it by using a generic opponent modeling mechanism, which allows the agent to model its counterpart's population and adapt its behavior to that population [32]. The extended agent, called KBAGENT, is an automated negotiator that negotiates with each person only once, and uses past negotiation sessions of others as a knowledge base for generic opponent modeling. The database containing the a relatively small number of past negotiation sessions is used to extract the likelihood of acceptance of proposals and which proposals may be offered by the opposite side. The performance of KBAGENT in terms of its counter-offer generation and generic opponent modeling was tested against people in the Tobacco and the Job interview domains.

The results of these tests indicate that the KBAGENT negotiates proficiently with people and even achieves higher utility score values than the QOAGENT. Moreover, the KBAGENT achieves significantly better agreements, in terms of utility score, than the human counterparts in *both* roles. These results

Your role: **Employer.**
Remaining time: 27:40
Your scores
Partner's scores

This is a simulation of **negotiation** between an Employer and a Job Candidate, about the Candidate's job conditions.

- You play the role of the **Employer**.
- You communicate with your opponent using **chat** only.
- The negotiation relates to several **issues**. For each issue, there are several possible **values**.
- Your score depends on the **values** in the final agreement, and on **time**. Click "Your scores" for details.
- Use the menus below as a draft. Your partner cannot see them - you communicate with **chat** only.
- This draft is **not binding** until both of you sign it.

Issue	Employer draft
Salary:	12,000 NIS
Job Description:	-- select --
Leased Car:	No agreement
Pension Fund:	-- select --
Promotion Possibilities:	-- select --
Working Hours:	9 hours
<b>Your score:</b>	294

[Sign the agreement](#)

Please write a complete, grammatically correct English sentence.

History of actions			
#	Proposer	Action	Offer
3	Employer (You)	Message	I want to pay 12,000
4	Candidate	Message	I agree to { Salary : 12,000 NIS }. Now let's talk about the other issues. I offer you will be the project manager
5	Employer (You)	Message	I need QA
6	Candidate	Message	I cannot agree to { Job Description : QA }. I require Project Manager
7	Candidate	Message	my offer is project manager

Click here to quit the negotiation immediately: [Opt Out](#)

Figure 1: The negotiation system's interface for NegoChat.

indicate that integrating general opponent modeling into qualitative decision-making is beneficial for automated negotiations.

### 2.3 NegoChat Agent

All the agents we discussed so far negotiated with the human counterpart either using a structured language or using a menu-driven interaction. They lack the natural language processing support required to enable real world types of interactions. To address this challenge we first developed an NLP module that translates the free text of the human player to the agent's formal language. We modified the KBagent by adding this module without changing the KBagent strategy and ran an experiment in which the modified KBagent played with people in a chat-like environment (see Figure 1 for the negotiation system's interface for chat-based negotiations). We found that simply modifying the KBagent to include an NLP module is insufficient to create a good agent for such settings and the revised agent achieved relatively low utility scores. The main observation was that people in chat-based negotiations make and accept partial agreements and follow issue-by-issue negotiations while the KBagent proposes full offers and has difficulties reaching partial agreements. To address this limitation, we developed NegoChat, which extended the KBagent focusing on strategies that allow for partial agreements and issue-by-issue interactions. NegoChat's algorithm is based on bounded rationality, specifically anchoring and Aspiration

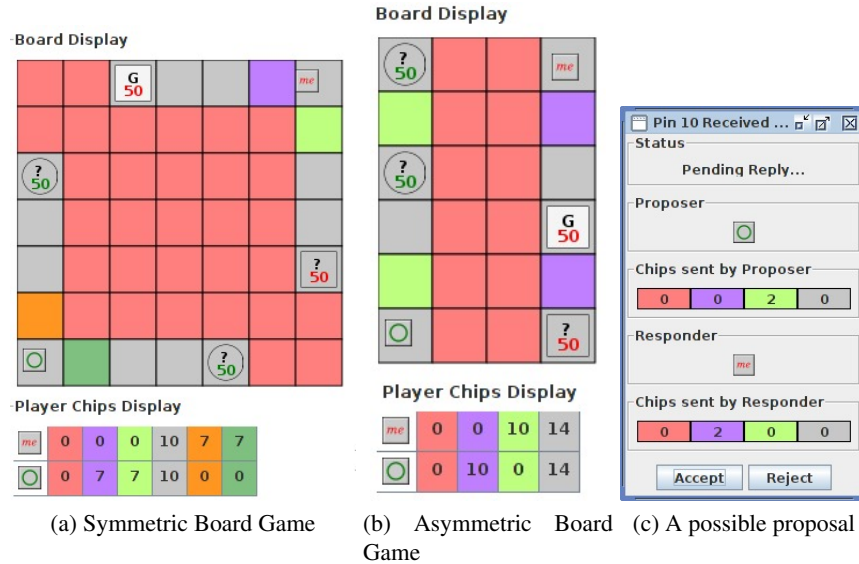


Figure 2: Two CT revelation games

Adaptation Theory (AAT). The AAT was used for deciding on the order in which the issues will be discussed. The agent begins each negotiation interaction by proposing a full offer based on the KBagent's strategy, which serves as its anchor. Assuming this offer is not accepted, NegoChat then proceeds to negotiate via partial agreements, proposing the next issue for negotiation based on people's typical urgency (according to AAT).

We evaluated the NegoChat agent in extensive experiments negotiating with people in the job interview domain. We compared its performance to the performance of the KBagent that also negotiated with (different) people using the same NLP module. The NegoChat agent achieved significantly better agreements (i.e., higher utility score) in less time. However, people playing against KBagent, on average, did better. This implies that some of NegoChat's success is evidently at the cost of the person's score and consequently the social welfare score of this agent is not significantly better than that of KBagent. As our goal is to maximize the agent's utility score this should not be seen as a fault. However, future generations of automated agents may decide to implement different strategies to maximize social welfare.

### 3 Negotiations and Actions Interleaving

In most situations, negotiation is not done in isolation but is associated with agreement implementation and other activities. We developed agents that can interact with people in such settings. These studies were carried out in a configurable system called Colored Trails (CT)<sup>1</sup>. It is a game played by two or more participants on a board of colored squares. CT is an abstract, conceptually simple but highly versatile game in which players negotiate and exchange resources to enable them to achieve their individual or group goals. It provides a realistic analogue to multi-agent task domains, while not requiring extensive domain modeling [15, 18]. A summary of the agents we developed are specified in Table 2.

<sup>1</sup>See <http://www.eecs.harvard.edu/ai/ct>.

Table 2: CT games

Settings	Agent Name	Agent Properties	Significance vs people
Bilateral, Uncertainty, Revelation Game two-phases: revelation, bargaining	PBE agent	Bayesian perfect equilibrium	No roles
	SIGAL	decision theory, machine learning	Both roles
Bilateral, full information, agreements not enforceable, multiple rounds, three phases: bargaining, resource exchange, movement	PAL	decision theory: Influence Diagram machine learning	Both roles
Contract game, three players two-phases: bargaining, movement	SPE agent	subgame-perfect equilibrium	CS role
	SP-RAP	subgame-perfect equilibrium bounded rational model of opponent risk averse	SP role

### 3.1 Revelation games

We considered negotiation settings in which participants lack information about each other’s preferences, often hindering their ability to reach beneficial agreements [34]. Specifically, we studied a particular class of such settings we call “revelation games”, in which two players are given the choice to truthfully reveal private information before commencing two rounds of alternating negotiation. Revealing this information narrows the search space of possible agreements and may lead to agreement more quickly, but may also cause players to be exploited by others (see examples of such games in Figure 2). Revelation games combine two types of interactions that have been studied in the past in the economics literature: Signaling games [39], in which players choose whether to convey private information to each other, and bargaining [31], in which players engage in multiple negotiation rounds.

We were hopeful that, for revelation games, equilibrium-based agents will interact well with people since behavioral economics work has shown that people often follow equilibrium strategies [7] when deciding whether to reveal private information to others. The question is whether this observation will be stronger than our previous observations reported above that people’s behavior in bargaining settings does not adhere to equilibrium strategies. We formalized the setting as a Bayesian game and computed two types of perfect Bayesian equilibrium: a separating equilibrium where both players reveal their type, and a pooling equilibrium where none of the players reveal their types.

We compared the equilibrium agents with people playing with other people and with the Sigmoid Acceptance Learning Agent (SIGAL) that we developed [34]. The SIGAL agent used classical machine learning techniques to predict how people make and respond to offers during negotiation, how they reveal information and their response to potential revelation actions by the agent. This model is integrated into the agent’s decision tree. We conducted an extensive empirical study spanning hundreds of human subjects.

Results show that the SIGAL agent was able to outperform people and the equilibrium agents. Furthermore, people outperformed the equilibrium agents. It turned out that the negotiation part of the game was more important (with respect to the utility score) than the revelation part. The equilibrium agent made very selfish offers in the last round of the negotiations. Most of these offers were rejected. In the first round, it made offers that were highly beneficial to people and most of these offers were accepted,




Figure 3: An example of a CT Board for multiple negotiation games with unenforceable agreements.

but the small benefit it incurred in these proposals did not aid its performance. The SIGAL agent, on the other hand, (i) learned to make offers that were beneficial to people while not compromising its own benefit; and (ii) incrementally revealed information to people in a way that increased its expected performance. We were able to adjust SIGAL to new, similar settings that varied rules and situational parameters of the game without the need to accumulate new data. However, moving to a completely new setting requires a lot of work collecting data and adjusting the machine learning module to the new setting.

### 3.2 Non-binding agreements

We also studied CT settings of two players in which both participants needed to complete their individual tasks by reaching agreements and exchanging resources, the number of negotiation rounds were not fixed in advance, and the negotiation protocol was an alternating offers protocol that allowed parties to choose the extent to which they kept each of their agreements during the negotiation [19]. That is, there are three phases in each round of the game: negotiation, transfer and movement. The negotiation phase consisted of two rounds of alternating offers in which the players could reach an agreement on resource exchange. After each phase of negotiations, the game moved to the “transfer phase” in which both players could transfer resources to each other. The transfer action was done simultaneously, such that neither player could see what the other player transferred until the end of the phase. A player could choose to transfer more resources than it agreed to, or any subset of the resources it agreed to, including not transferring any resources at all. In the “movement phase” both players could move their icons on the board one step towards the goal square, provided they had the necessary resources. Then, the game moved to the next round, beginning again with negotiation phase. The game ends when one of the players reaches his goal or does not move for two rounds (see an example of one such game in Figure 3).

The most important decision of a player in such settings is whether or not to keep the agreements. Another important decision is whether to accept an offer given by the other player. In subgame perfect equilibrium, the players should not keep the agreements. Different equilibria may specify various strategies for the acceptance decision. We ran preliminary experiments and observed that such strategies are not beneficial when the equilibrium agent interacts with people. Most of the time the agent was not able to reach its goal, yielding a low utility score.

Galit et al. [19] present the Personality Adaptive Learning (PAL) agent for negotiating with people from different cultures for the CT game where agreements are not enforceable. The methodology was similar to that of SIGAL (Section 3.1), combining a decision-theoretic model using a decision tree with






Figure 4: An example of a CT Board for the Contract game.

classical machine learning techniques to predict how people respond to offers, and the extent to which they fulfill agreements.

PAL was evaluated empirically in the Colored Trails (CT) environment by playing with people in three countries: Lebanon, the U.S., and Israel, in which people are known to vary widely in their negotiation behavior. The agent was able to outperform people in all three countries.

### 3.3 Contract Game

We studied commitment strategies in a three-player CT game. The game is called Contract Game and is analogous to a market setting in which participants need to reach agreements over contracts and commit to or renege from contracts over time in order to succeed [20]. The game comprises three players, two service providers and one customer. The service providers compete to make repeated contract offers to the customer consisting of resource exchanges in the game (see an example of one such game in Figure 4). We formally analyzed the game to compute subgame perfect equilibrium strategies for the customer and service provider in the game that are based on making contracts containing commitment offers. To evaluate agents that use the equilibrium strategies, we conducted extensive empirical studies in three different countries, the U.S., Israel and China. We ran several configurations in which two human participants played a single agent participant in various role configurations in the game. Our results showed that the computer agent using subgame Nash equilibrium strategies for the customer role was able to outperform people playing the same role in all three countries and obtained statistically significant, higher utility scores than the humans. This was very surprising since it was the first EQ agent after trying many equilibrium agents that was able to achieve such results.

In particular, the customer agent made significantly more commitment type proposals than people did, and requested significantly more resources from service providers than did people. It was quite surprising that people playing with it accepted these offers; in other settings (such as the revelation games) such unfair offers were rejected by people. We hypothesize that the competition between the two service providers made such offers more acceptable. In addition, while in the revelation games the EQ agent had only one opportunity to make an offer, in the contract game it could make offers several times (off the equilibrium path) which we believe also increased the acceptance rate. Also, the customer agent reached one of the goals in all its games and was able to reach the goal significantly more often than people. This is again quite surprising since at the beginning of the game the customer has enough resources to reach both goals. Since reaching one of the goals is very beneficial to the customer it is difficult to understand why human players hadn't always reached the goal.

While the customer EQ agent outperformed people, people outperformed the EQ agent when it played the role of one of the service providers. We believe that this is mainly due to people playing

the customer role not reaching the goal even when they have all the needed resources to do so. To face this problem we then developed an agent termed SP-RAP which extended the EQ agent in the following two ways to handle the uncertainty that characterizes human play in negotiation: First, it employed a risk averse strategy using a convex utility function. Second, it reasoned about a possibly bounded rational customer (CS) player by assigning a positive probability  $p > 0$  for the customer player not reaching the goal. We assigned a separate value for  $p$  for each country by dividing the number of times the CS player reached the goal by the total number of games played. Consequently, SP-RAP outperformed people playing the SP role in all three countries.

## 4 Argumentation Agent

An automated agent can help a human when engaging in an argumentative dialog by utilizing its knowledge and computational advantage to provide arguments to him. Argumentation was studied extensively using the well-established Argumentation Theory (see [41] for a summary). Therefore, in the first step in the development of an automated agent that advised people in such settings, we considered the abilities of Argumentation Theory to predict people's arguments. In [38] we presented extensive studies in three experimental settings, varying in complexity, which show the lack of predictive power of the existing Argumentation Theory. Second, we used Machine Learning (ML) techniques to provide a probability distribution over all known arguments given a partial deliberation. That is, our ML techniques provided the probability of each argument to be used next in a given dialog. Our model achieves 76% accuracy when predicting people's top three argument choices given a partial deliberation. Last, using the prediction model and the newly introduced heuristics of relevance, we designed and evaluated the Predictive and Relevance based Heuristic agent (PRH). Through an extensive human study, we showed that the PRH agent outperforms other agents that propose arguments based on Argumentation Theory, predicted arguments without heuristics or only the heuristics on both axes we examined: people's satisfaction from agents and people's use of the suggested arguments.

## 5 Security Games

The last several years have witnessed the successful application of Bayesian Stackelberg games in allocating limited resources to protect critical infrastructures. These interesting efforts have been led by Prof. Milind Tambe from USC. The first application is the ARMOR system (Assistant for Randomized Monitoring over Routes) that has been deployed at the Los Angeles International Airport (LAX) since 2007 to randomize checkpoints on the roadways entering the airport and canine patrol routes within the airport terminals [33, 35]. Other applications include IRIS, a game-theoretic scheduler for randomized deployment of the US Federal Air Marshal Service (FAMS) requiring significant scale-up in underlying algorithms, which has been in use since 2009 [40]; and PROTECT, which requires further scale up, is deployed for generating randomized patrol schedules for the US Coast Guard in Boston, New York, Los Angeles and other ports around the US [2, 3]. Furthermore, TRUSTS has been evaluated for deterring fare evasion, suppressing urban crime and counter-terrorism within the Los Angeles Metro System [44, 21, 9] and GUARDS was earlier tested by the US Transportation Security Administration (TSA) for security inside the airport [37].

The evaluation of these systems could be quite limited. The only system that was truly evaluated in the field is TRUSTS. We were able to conduct controlled experiments of our game theoretic resource

allocation algorithms. Before this project, the actual evaluation of the deployed security games applications in the field was a major open challenge. The reasons were twofold. First, previous applications focused on counter-terrorism, therefore controlled experiments against real adversaries in the field were not feasible. Second, the number of practical constraints related to real-world deployments limited the ability of researchers to conduct head-to-head comparisons

In TRUSTS we were able to address this challenge and run the largest scale evaluation of security games in the field in terms of duration and number of security officials deployed. We evaluated each component of the system (Fare Evasion, Counter Terrorism and Crime algorithms) by designing and running field experiments. In the context of fare evasion, we ran an extensive experiment, where we compared schedules generated using game theory against competing schedules comprised of a random scheduler, augmented with officers providing real-time knowledge of the current situation. Our results showed that our schedules led to statistically significant improvements over the competing schedules, despite the fact that the latter were improved with real-time knowledge.

In addition, extensive human experiments in the lab were conducted [36, 30]. These experiments showed that incorporating bounded rational models of the adversaries to the Stackelberg games improves the performance of the defenders. These results were observed both when the role of adversaries was played by novices and when it was played by security experts.

## 6 Persuasion Games

A persuasion game involves two players: a sender who attempts to persuade another agent (the receiver) to take a certain action [17]. Persuasion games are similar to both negotiation games and security games [43]. They are similar to negotiation and argumentation games since one player tries to convince another player to do something, as in negotiations. They are also similar to security games in the asymmetry between the players: the sender and the defender are trying to influence the activities of the receiver and the attacker, respectively. So, it is interesting to check if equilibrium strategies will be beneficial in persuasion games. Furthermore, the incorporation of a bounded rational model of the receiver will be beneficial to the sender as the incorporation of bounded rational models of the attacker was beneficial to the defender in security games.

We focused on information disclosure games with two-sided uncertainty [5, 6]. This is a special type of persuasion game in which an agent tries to lead a person to take an action that is beneficial to the agent by providing him with truthful, but possibly partial, information relevant to the action selection. We first computed the subgame perfect Bayesian Nash equilibrium of the game assuming the human receiver is fully rational. We developed a sender agent that follows the equilibrium strategy (GTBA agent).

We also developed a machine learning-based model that effectively predicts people's behavior in these games and we called it Linear weighted-Utility Quantal response (LUQ). The model we provide assumes that people use a subjective utility function which is a linear combination for all given attributes. The model also assumes that while people use this function as a guideline, they do not always choose the action with the greatest utility value, however, the higher an action's utility value is, the more likely they are to choose that action. We integrated this model into our persuasion model and built the LUQA agent.

We ran an extensive empirical study with people in two different games. In a multi-attribute road selection game with two-sided uncertainty, the LUQA agent obtained significantly higher utility points than the GTBA agent. However, in the second game, the Sandwich game, there was no significant advantage to the machine learning-based model, and using the game theory-based agent, GTBA, which assumes that people maximize their expected monetary values is beneficial. We hypothesize that these

different results are due to the nature of the domains. The monetary result plays an important role in the sandwich game. This is because the game is played in an environment where a person's goal is to make a profit. However, in the road selection game the utility scores are associated with time. Thus, it seems that maximizing expected monetary utility is easier for people than maximizing utility scores that are associated with time.

## 7 Discussions

The state-of-the-art agent, NegoChat, for multi-issue negotiations integrates methods from several disciplines: qualitative decision-making, machine learning and heuristics based on psychological theories. None of the equilibrium agents that were developed were successful when interacting with people. The reliance on heuristic and machine learning makes the transfer of NegoChat from one scenario to the other and from one culture to the other problematic. This was evident recently when we tried to run experiments with NegoChat, which was developed based on data collected in Israel and Egypt. We had to spend a lot of time and effort until this transfer was possible.

Similarly, in most of the cases, the equilibrium agent was not successful in the CT game settings. The only exception is the contract game. We believe that the success of the equilibrium agent in the contract game has to do with the specifics of the game: the competition between the two SPs. In the contract game, it was extremely difficult to predict people's behavior, thus the success of the equilibrium agent is even more significant.

The same observations were seen in argumentation – the argumentation theory-based agent was not very successful. Therefore, in all these cases the development of new negotiation agents to new settings requires the collection of data and the adjustment of the agent to the new settings. Therefore, we strongly believe that the development of theoretical models for the design and implementation of agents that negotiate in multi-issue negotiation settings can be very useful. However, this is still an open question.

On the other hand, it seems that in security games the deployment of Stackelberg equilibrium is beneficial (possibly with the incorporation of a bounded rational model of the attacker) and similarly in persuasion games where using subgame perfect Bayesian equilibrium is beneficial (possibly with the incorporation of a bounded rational model of the receiver).

We hypothesize that this is the case since in security games and persuasion games the interactions between the agent and the human is quite limited. The attacker or the receiver needs to choose one action compared to many decision-making activities that are required from a human negotiator. Nevertheless, even in security games and to some extent in persuasion games it was shown that taking the limitations of the other player into consideration is beneficial.

## 8 Acknowledgment

This work was supported in part by ERC grant #267523.

## References

- [1] O. Amir, B. Grosz, E. Law & R. Stern (2013): *Collaborative health care plan support*. In: *Proceedings of the Eleventh International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2013)*, St Paul, MN, pp. 793–796.

- [2] Bo An, Manish Jain, Milind Tambe & Christopher Kiekintveld (2011): *Mixed-Initiative Optimization in Security Games: A Preliminary Report*. In: *AAAI Spring Symposium: Help me help you: Bridging the gaps in human-agent collaboration*.
- [3] Bo An, Fernando Ordóñez, Milind Tambe, Eric Shieh, Rong Yang, Craig Baldwin, Joseph DiRenzo III, Kathryn Moretti, Ben Maule & Garrett Meyer (2013): *A deployed quantal response-based patrol planning system for the US Coast Guard*. *Interfaces* 43(5), pp. 400–420, doi:10.1287/inte.2013.0700.
- [4] D. Ariely (2008): *Predictably Irrational*. Harper Collins, doi:10.1002/bdm.673.
- [5] Amos Azaria, Zinovi Rabinovich, Claudia V Goldman & Sarit Kraus (2014): *Strategic information disclosure to people with multiple alternatives*. *ACM Transactions on Intelligent Systems and Technology (TIST)* 5(4), p. 64, doi:10.1145/2558397.
- [6] Amos Azaria, Zinovi Rabinovich, Sarit Kraus & Claudia V. Goldman (2011): *Strategic Information Disclosure to People with Multiple Alternatives*. In: *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2011, San Francisco, California, USA, August 7-11, 2011*.
- [7] Jeff Banks, Colin F Camerer & David Porter (1994): *Experimental Tests of Nash Refinements in Signaling Games*. *Games and Economic Behavior* 6, pp. 1–31, doi:10.1006/game.1994.1001.
- [8] M. H. Bazerman & M. A. Neale (1992): *Negotiator Rationality and Negotiator Cognition: The Interactive Roles of Prescriptive and Descriptive Research*. In H. P. Young, editor: *Negotiation Analysis*, The University of Michigan Press, pp. 109–130.
- [9] Francesco M Delle Fave, Albert Xin Jiang, Zhengyu Yin, Chao Zhang, Milind Tambe, Sarit Kraus & John P Sullivan (2014): *Game-Theoretic Patrolling with Dynamic Execution Uncertainty and a Case Study on a Real Transit System*. *Journal of Artificial Intelligence Research*, pp. 321–367, doi:10.1613/jair.4317.
- [10] E. Durenard (2013): *Professional Automated Trading: Theory and Practice*. John Wiley & Sons, doi:10.1002/9781118755846.
- [11] Avshalom Elmalech, David Sarne & Barbara J. Grosz (2015): *Problem restructuring for better decision making in recurring decision situations*. *Autonomous Agents and Multi-Agent Systems* 29(1), pp. 1–39, doi:10.1007/s1045801492473.
- [12] Avshalom Elmalech, David Sarne, Avi Rosenfeld & Eden Shalom Erez (2015): *When Suboptimal Rules*. In: *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, pp. 1313–1319.
- [13] I. Erev & A. E Roth (1998): *Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria*. *American Economic Review* 88(4), pp. 848–881.
- [14] Shaheen Fatima, Sarit Kraus & Michael Wooldridge (2014): *Principles of Automated Negotiation*. Cambridge University Press, doi:10.1017/CBO9780511751691.
- [15] Ya'akov Gal, Barbara Grosz, Sarit Kraus, Avi Pfeffer & Stuart Shieber (2010): *Agent decision-making in open mixed networks*. *Artificial Intelligence* 174(18), pp. 1460–1480, doi:10.1016/j.artint.2010.09.002.
- [16] Adam D Galinsky & Thomas Mussweiler (2001): *First offers as anchors: the role of perspective-taking and negotiator focus*. *Journal of personality and social psychology* 81(4), p. 657, doi:10.1037/0022-3514.81.4.657.
- [17] Jacob Glazer & Ariel Rubinstein (2006): *A study in the pragmatics of persuasion: a game theoretical approach*. *Theoretical Economics* 1(4), pp. 395–410.
- [18] Barbara J Grosz, Sarit Kraus, Shavit Talman, Boaz Stossel & Moti Havlin (2004): *The influence of social dependencies on decision-making: Initial investigations with a new game*. In: *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems-Volume 2*, IEEE Computer Society, pp. 782–789.
- [19] G. Haim, Y. Gal, M. Gelfand & S. Kraus (2012): *A cultural sensitive agent for human-computer negotiation*. In: *Proceedings of the Eleventh International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2012)*, Valencia, Spain, pp. 451–458.

- [20] Galit Haim, Ya'akov (Kobi) Gal, Sarit Kraus & Bo An (2014): *Human-Computer Negotiation in Three-Player Market Settings*. In: *Proc. of ECAI14*, pp. 417–422.
- [21] Albert Xin Jiang, Zhengyu Yin, Chao Zhang, Milind Tambe & Sarit Kraus (2013): *Game-theoretic randomization for security patrolling with dynamic execution uncertainty*. In: *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*, International Foundation for Autonomous Agents and Multiagent Systems, pp. 207–214.
- [22] K. Kauppi, A. Brandon-Jones, S. Ronchi & E. van Raaij (2013): *Tools without skills: Exploring the moderating effect of absorptive capacity on the relationship between e-purchasing tools and category performance*. *International Journal of Operations & Production Management* 33(7), pp. 828–857, doi:10.1108/IJOPM-12-2011-0445.
- [23] G. Kersten & H. Lai (2007): *Negotiation support and e-negotiation systems: an overview*. *Group Decision and Negotiation* 16(6), pp. 553–586, doi:10.1007/s10726-007-9095-5.
- [24] Sarit Kraus, Penina Hoz-Weiss, Jonathan Wilkenfeld, David R Andersen & Amy Pate (2008): *Resolving crises through automated bilateral negotiations*. *Artificial Intelligence* 172(1), pp. 1–18, doi:10.1016/j.artint.2007.05.007.
- [25] D. A. Lax & J. K. Sebenius (1992): *Thinking coalitionally: party arithmetic, process opportunism, and strategic sequencing*. In H. P. Young, editor: *Negotiation Analysis*, The University of Michigan Press, pp. 153–193.
- [26] R. Lin, Y. Oshrat & S. Kraus (2009): *Investigating the benefits of automated negotiations in enhancing people's negotiation skills*. In: *Proceedings of the Eighth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2009)*, Budapest, Hungary, pp. 345–352.
- [27] Raz Lin, Sarit Kraus, Jonathan Wilkenfeld & James Barry (2008): *Negotiating with bounded rational agents in environments with incomplete information using an automated agent*. *Artificial Intelligence* 172(6), pp. 823–851, doi:10.1016/j.artint.2007.09.007.
- [28] M. Maschler, E. Solan & S. Zamir (2013): *Game Theory*. Cambridge University Press: Cambridge, England, doi:10.1017/CBO9780511794216.
- [29] R. D. McKelvey & T. R. Palfrey (1992): *An experimental study of the centipede game*. *Econometrica* 60(4), pp. 803–836, doi:10.2307/2951567.
- [30] Thanh Hong Nguyen, Rong Yang, Amos Azaria, Sarit Kraus & Milind Tambe (2013): *Analyzing the Effectiveness of Adversary Modeling in Security Games*. In: *Proc of AAAI 2013*.
- [31] Martin J Osborne & Ariel Rubinstein (1994): *A course in game theory*. MIT press.
- [32] Y. Oshrat, R. Lin & S. Kraus (2009): *Facing the Challenge of Human-Agent Negotiations via Effective General Opponent Modeling*. In: *Proceedings of the Eighth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2009)*, Budapest, Hungary, pp. 377–384.
- [33] Praveen Paruchuri, Jonathan P Pearce, Janusz Marecki, Milind Tambe, Fernando Ordonez & Sarit Kraus (2008): *Playing games for security: an efficient exact algorithm for solving Bayesian Stackelberg games*. In: *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 2*, International Foundation for Autonomous Agents and Multiagent Systems, pp. 895–902.
- [34] N. Peled, K. Gal & S. Kraus (2015): *A study of computational and human strategies in revelation games*. *Autonomous Agents and Multi-Agent Systems* 29(1), pp. 73–97, doi:10.1007/s10458-014-9253-5.
- [35] James Pita, Manish Jain, Janusz Marecki, Fernando Ordóñez, Christopher Portway, Milind Tambe, Craig Western, Praveen Paruchuri & Sarit Kraus (2008): *Deployed ARMOR protection: the application of a game theoretic model for security at the Los Angeles International Airport*. In: *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems: industrial track*, International Foundation for Autonomous Agents and Multiagent Systems, pp. 125–132.
- [36] James Pita, Manish Jain, Milind Tambe, Fernando Ordóñez & Sarit Kraus (2010): *Robust solutions to stackelberg games: Addressing bounded rationality and limited observations in human cognition*. *Artificial Intelligence* 174(15), pp. 1142–1171, doi:10.1016/j.artint.2010.07.002.

- [37] James Pita, Milind Tambe, Christopher Kiekintveld, Shane Cullen & Erin Steigerwald (2011): *GUARDS - innovative application of game theory for national airport security*. In: *IJCAI Proceedings-International Joint Conference on Artificial Intelligence*, p. 2710.
- [38] Ariel Rosenfeld & Sarit Kraus (2015): *Providing arguments in discussions based on the prediction of human argumentative behavior*. In: *Proc. of AACL*.
- [39] Andrew Michael Spence (1974): *Market signaling: Informational transfer in hiring and related screening processes*. 143, Harvard Univ Pr.
- [40] Jason Tsai, Christopher Kiekintveld, Fernando Ordonez, Milind Tambe & Shyamsunder Rathi (2009): *IRIS-a tool for strategic security allocation in transportation networks*. In: *Proc. of AAMAS09*.
- [41] Douglas Walton (2009): *Argumentation theory: A very short introduction*. In: *Argumentation in artificial intelligence*, Springer, pp. 1–22, doi:10.1007/978-0-387-98197-0-1.
- [42] Jonathan Wilkenfeld, Sarit Kraus, Kim M Holley & Michael A Harris (1995): *GENIE: A decision support system for crisis negotiations*. *Decision Support Systems* 14(4), pp. 369–391, doi:10.1016/0167-9236(94)00027-P.
- [43] Haifeng Xu, Zinovi Rabinovich, Shaddin Dughmi & Milind Tambe (2015): *Exploring Information Asymmetry in Two-Stage Security Games*. In: *Proc of AACL*.
- [44] Zhengyu Yin, Albert Xin Jiang, Milind Tambe, Christopher Kiekintveld, Kevin Leyton-Brown, Tuomas Sandholm & John P Sullivan (2012): *TRUSTS: Scheduling randomized patrols for fare inspection in transit systems using game theory*. *AI Magazine* 33(4), p. 59.