

Designing Automated Agents Capable of Efficiently Negotiating with People - Overcoming the Challenge*

Raz Lin¹ and Sarit Kraus^{1,2**}

¹ Computer Science Department
Bar-Ilan University
Ramat-Gan, 52900 Israel

{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. Yet, most of the time, the research does not focus on automated negotiators capable of negotiating efficiently with people. Many challenges are facing agents designers who aim to design an automated negotiator, even when people are not in the loop. In addition, the fact that people are scarce resources makes the validation process of the efficacy of the automated negotiators an exhausting task. Yet, these challenges can be overcome. By reviewing existing automated negotiators and reporting on experiments we conducted with automated negotiators, we shed some light on how these challenges can be overcome, and thus motivate other researchers to pursue this exciting line of research.

1 Introduction

Research on automated negotiators has flourished in recent years. The negotiation domain is an attractive environment for automated agents, resulting in many benefits. Automated agents can be used with humans in the loop or without them. On the one hand, they can alleviate some of the efforts required of people during negotiations and also assist people that are less qualified in the negotiation process. On the other hand there may be situations in which automated negotiators can even replace human negotiators. Succeeding in developing an automated agent with negotiation capabilities has great advantages and implications.

Building an automated negotiators capable of efficient negotiation with people is not the same as building an automated negotiator capable of efficient negotiation with another automated agent. Special consideration should be made in order to deal with people negotiators. For example, humans tend to make mistakes, and they are affected by cognitive, social and cultural factors, etc. [1]. Even after designing what might be the

* Some of this work will also appear in a review paper to be published in the Communications of the ACM.

** 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.

ultimate automated negotiator, the task of validating it against people is tiresome and exhausting. Mistakes done in the design of the agent or in the design of the experiment might require discarding all results accumulated thus far and the return to square one.

Note that the focus of this paper is not on designing the best or the most efficient automated negotiator, but rather explain and demonstrate the different consideration required by the designer when designing and validating the design of an automated negotiator. We explain this by reviewing existing automated negotiators and reporting on our own experiments with developing agents that negotiate with people. As the negotiation domain is vast, in the paper we concentrate on lessons learned in situations of adversarial *bilateral bargaining*. In addition, we do not focus here on experimental and research design considerations, such as the use of qualitative (e.g., questionnaires) or quantitative approaches. We refer the reader to [3, 6, 39] for surveys and further review on these issues.

This paper contributes to research on automated negotiations in several ways. First, it tackles the problem of negotiating with people. It provides insights that should alleviate the experimental issues with people and assist researches in experimenting with people. By doing so, it also serves as a motivation for future research directions with respect to automated negotiations with people and thereafter allow investigating behavioral and cognitive aspects of negotiations undertaken by human negotiators. Moreover, we provide significant insights that should assist designers in the the design process of automated negotiators.

The rest of this paper is organized as follows. In Section 2 we describe lessons learned with respect to the design of automated negotiators and the different consideration employed when negotiating with people. We continue and provide insights on the evaluation and validation of automated negotiators based on extensive experiments in Section 3. Finally, we provide a summary and discuss future research.

2 The Design of Automated Negotiators

2.1 The Negotiation Environment

When designing an automated agent, the designer needs to take into account the environment in which the agent will operate. Then, the agent can be best tailored to the specific environment to achieve the most efficient outcomes. The environment determines several parameters which dictate the number of negotiators taking part in the negotiation, the time frame of the negotiation and the issues on which the negotiation is being conducted. The number of parties participating in the negotiation process can be two (bilateral negotiations) or more (multilateral negotiations). For example, in a market there can be one seller but many buyers, all involved in negotiating over a certain item. On the other hand, if the item is common, there may also be many sellers taking part in the negotiation process.

The negotiation environment also consists of a set of objectives and issues to be resolved. Various types of issues can be involved, including discrete enumerated value sets, integer-value sets, and real-value sets. A negotiation consists of multi-attribute issues if the parties have to negotiate an agreement which involves several attributes for each issue. Negotiations that involves multi-attribute issues allow making complex decisions while taking into account multiple factors [15]. The negotiation environment

can consist of non-cooperative negotiators or cooperative negotiators. Generally speaking, cooperative agents try to maximize their combined joint utilities (e.g., see [40]) while non-cooperative agents try to maximize their own utilities regardless of the other sides' utilities.

Finally, the negotiation protocol defines the formal interaction between the negotiators: whether the negotiation is done only once (one-shot) or repeatedly, and how the exchange of offers between the agents is conducted. A common exchange of offers model is the alternating offers model [30]. In addition, the protocol states whether agreements are enforceable or not, and whether the negotiation has a finite or infinite horizon. The negotiation is said to have a finite horizon if the length of every possible history of the negotiation is finite. In this respect, time costs may also be assigned and they may increase or decrease the utility of the negotiator.

Figure 1 depicts the different variations in the settings, along with the location of some automated negotiation systems designed for negotiating with people found in the literature (the *Diplomat* agent [19], *AutONA* [2], *Cliff-Edge* [14], the *CT* agent [11], the *Guessing Heuristic* agent [13], the *QO* agent [22], the *KB* Agent [26], and the *Virtual Human* [16]). For example, point D in the cube represents bilateral negotiations with multi-attribute issues and repeated interactions, while point B represents multi-lateral negotiations with a single attribute for negotiation and a one-shot encounter.

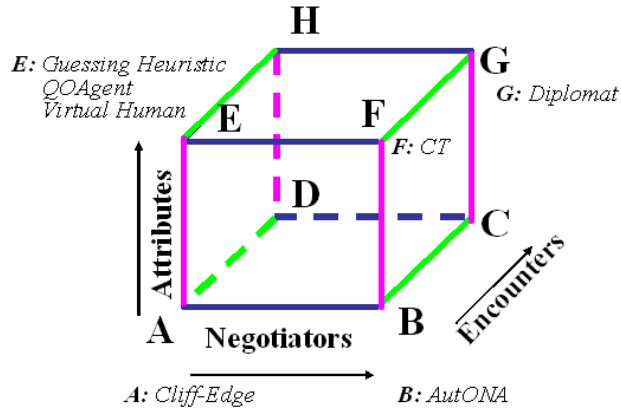


Fig. 1. Variations of the negotiation settings.

The negotiation domain encompasses the negotiation objectives and issues and assigns different values to each. Thus, an agent may be tailored to a given domain (e.g., the *Diplomat* agent [19] is tailored to a specific domain of the Diplomacy game) or domain independent (e.g., the *QOAgent* [22]).

2.2 Opponent Modeling

Automated negotiation agents capable of negotiating efficiently with people must deal with the fact that people are diverse in their behavior and each individual might negoti-

ate in a different manner. Thus, automated agents must rely on a good opponent modeling component to model their counterpart and adapt their behavior to their partner. For example, negotiation over the same issues against parties from different countries, can result in distinct agreements, and it is vital that the negotiator will be aware of these variations. It has been shown that in different countries the attitude regarding negotiations and the actions performed during them are quite different. A Chinese negotiator will appear to concede more often while in the UK it is common to use pressure tactics to impose a deal on the other party. The same tactic, however, against a negotiator from Greece will most likely backfire [5, 20]. Another issue is whether the negotiation is a one-shot negotiation or a repeated one. If the negotiation is done repeatedly with the same people, the automated negotiator can base a model of the counterpart based on the negotiation history of the same person, yet if it is a one-shot negotiation a *general* opponent modeling is required.

In this respect, it is important to know which information model is used in the negotiation. The information model dictates what is known to each agent. It can be a model of complete information, in which each agent has complete knowledge of both the state of the world and the preferences of other agents; or it can be a model of incomplete information, in which agents may have only partial knowledge of either the states of the world or the preferences of other agents (e.g., bargaining games with asymmetric information), or they may be ignorant of the preferences of the opponents and the states of the world [31]. The incomplete information can be modeled in different ways with respect to the uncertainty regarding the preferences of the other party. One approach to modeling the information is to assume that there is a set of different agent types and the other party can be any one of these types.

2.3 People are not Agents

It is important to understand that people do not act as automated agents. Thus, an automated agent which might negotiate efficiently against other automated agents, might do poorly against people, as we demonstrate later in Section 3. The design of automated negotiations is usually followed by assumptions that do not necessarily apply in genuine negotiations with humans, such as assuming complete information or the rationality of the opponent negotiator. In this sense, both parties are assumed to be rational in their behavior (e.g., the decisions made by the agents are described as rational and the agents are considered to be expected utility maximizing agents that cannot deviate from their prescribed behavior). Yet, when dealing with human counterparts, one must take into consideration the fact that humans do not necessarily maximize expected utility or behave rationally.

In particular, results from social sciences suggest that people do not follow equilibrium strategies [7, 24]. Moreover, when playing with humans, the theoretical equilibrium strategy is not necessarily the optimal strategy [36]. 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. Moreover, when people are the ones that design agents, they do not always design them to follow equilibrium strategies [11]. Nonetheless, some assumptions are made, i.e. mainly that the other party will not necessarily maximize its expected utility, however, if given two offers it will prefer the one with the highest utility value.

Lastly, it has been shown that whether the opponent is oblivious or has full knowledge that its counterpart is a computer agent can change the overall result. For example, Grossklags and Schmidt [10] showed that efficient market prices were achieved when human subjects knew that computer agents existed in a double auction market environment. Sanfey et al. [32] matched humans with other humans and with computer agents in the Ultimatum Game and showed that people rejected unfair offers made by humans at significantly higher rates than those made when matched with a computer agent.

Researchers have tried to take some of these issues into consideration when designing agents that are capable of proficiently negotiating with people. For example, dealing only with the bounded rationality of the opponent, several researchers have suggested new notions of equilibria (e.g., the *trembling hand equilibrium* described in Rasmusen ([28], p. 139)). Approximately ten years ago, Kasbah, a seminal negotiation model between agents designed by humans, was presented in the virtual marketplace by Chavez and Maes [4]. Here, the agent's behavior was fully controlled by human players. The main idea was to help users in the negotiation process between buyers and sellers by using automated negotiators. Chavez and Maes's main innovation was, not so much the sophisticated design of the automated negotiators, but rather the creation of a multi-agent negotiation environment. Kraus *et al.* [18] describe an automated agent that negotiates proficiently with humans. Although they also deal with negotiation with humans, there is complete information in their settings. Other researchers have suggested to shift from quantitative decision theory to qualitative decision theory [35]. In using such a model it is not necessary to assume that the opponent will follow the equilibrium strategy or try to be a utility maximizer. Another approach was to develop heuristics for negotiations motivated by the behavior of people in negotiations [19]. However, the fundamental question of whether it is possible to build automated agents for negotiations with humans in open environments has not been fully addressed by these researchers.

Another direction currently under pursuit is the development of virtual humans to train people in interpersonal skills (e.g., Kenny *et al.* [16]). Achieving this requires implementing cognitive and emotional modeling, natural language processing, speech recognition and knowledge representation. This in addition to constructing and implementing the appropriate logic for the task at hand (e.g., negotiation), is in order to make the virtual human into a good trainer.

Nonetheless, human factors and results of laboratory and field experiments reviewed in esteemed publications (e.g., [9, 27]) provide guidelines for the design of automated negotiators. Yet, it is still a great challenge to incorporate these guidelines in the inherent design of an agent to allow it to proficiently negotiate with people.

2.4 Agent Design

Tackling the issues of bounded rationality and incomplete information is a complex task. To achieve this, an automated agent is required to have two inter-dependent mechanisms. The first is a decision making component which works via modeling human factors. This mechanism is in charge of generating offers and deciding whether to accept or reject offers made by the opponent. The challenge behind this mechanism does not lie in the computational complexity of making good decisions, but rather in reasoning about the psychological and social factors that characterize human behavior. The

second component is learning, which allows the agent to infer the opponent's preferences and strategies, based on his actions.

Another inherent problem in the design of the automated agent is the ability to generalize its behavior. While humans can negotiate in different settings and domains, when designing an automated agent a decision should be made whether the agent should be a general purpose negotiator, that is, will be able to successfully negotiate in many settings and be domain-independent (e.g., Lin *et al.* [22]), or the agent will only be suitable for one specific domain (e.g., Kraus and Lehmann [19]). Perhaps the advantage of the agent's specificity is the ability to construct better strategies that could allow it to achieve better agreements, as compared to a more general purpose negotiator. This is due to the fact that the specificity allows the designer to debug the agent's strategy more carefully and against more test cases. By doing so, the designer can fine-tune the agent's strategy and allow for a more proficient automated negotiator. Agents that are domain independent, on the other hand, are harder to test against all possible cases and states.

The issue of trust also plays an important role in negotiations, especially when the other side's behavior is unpredictable. Successful negotiations depend on the trust established between all parties, which can depend on cheap-talk during negotiations (that is, unverifiable information with regard to the other party's private information [8]) and the introduction of unenforceable agreements. Based on the actions and information each party can update its reputation (for better or for worse) with regard to the other party and thus build trust between the sides. Some of the systems that we review below do allow cheap-talk and unenforceable agreements. Building trust can also depend on past and future interactions with the other party (e.g., one-shot interaction or repeated interactions). In this article, though, due to limited space we do not cover the issue of trust in detail. Readers can for example refer to [29] for a comprehensive review on the topic of trust.

Figure 2 depicts a general architecture for an automated agent design. We begin by describing the oldest agent of all of them, i.e. the *Diplomat* agent.

The *Diplomat* Agent Over twenty years ago Kraus and Lehmann developed an agent called *Diplomat* [19], that played the Diplomacy game with the goal to win. The game involves negotiations in multi-issue settings with incomplete information concerning the other agents' goals, and misleading information can be exchanged between the different agents. The negotiation protocol extends the model of alternating offers and allows simultaneous negotiations between the parties, as well as multiple interactions with the opponent agents during each time period. The issue of trust also plays an important role, as commitments might be breached. In addition, as each game consists of several sessions, it can be viewed as repeated negotiation settings.

The main innovation of the *Diplomat* agent is, most probably, the fact that it consists of five different modules that work together to achieve a common goal. Different personality traits are implemented in the different modules. These personality traits affect the behavior of the agent and can be changed during each run, which allows *Diplomat* to change its 'personality' from one game to another and to act non-deterministically. In addition, the agent has a limited learning capability which allows it to try to estimate the personality traits of its rivals (e.g., their risk attitude). Based on this, *Diplomat* as-

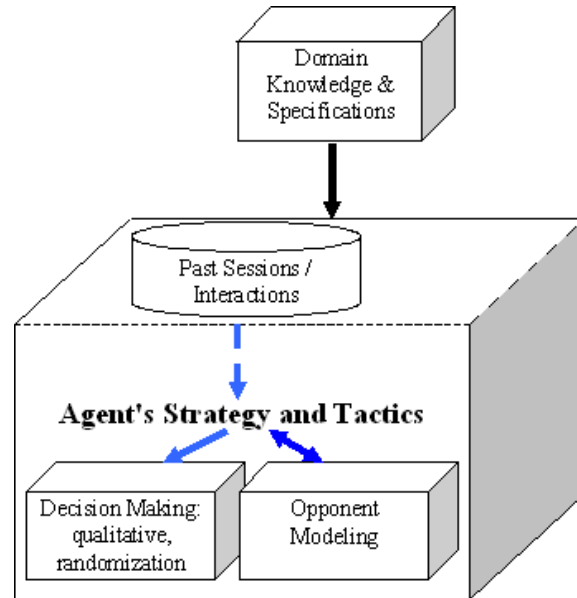


Fig. 2. Architecture of a general agent's design.

sesses whether or not the other players will keep their promises. In addition, *Diplomat* incorporates randomization in its decision making component. This randomization, influenced by *Diplomat*'s personality traits, determines whether some agreements will be breached or fulfilled.

We continue with a more recent agent which is also constrained to a specific domain and involves single-issue negotiations. However, it also takes into account the history of past interactions to model the opponents.

The *AutONA* Agent Byde *et al.* [2] developed *AutONA*, which is an automated negotiation agent. Their problem domain involves multiple negotiations between buyers and sellers over the price and quantity of a given product. The negotiation protocol follows the alternating offers model. Each offer is directed at only one player on the other side of the market, and is private information between each pair of buyers and sellers. In each round, a player can make a new offer, accept an offer, or terminate negotiations. In addition, a time cost is used to provide incentives for timely negotiations. While the model can be viewed as one-shot negotiations, for each experiment, *AutONA* was provided with data from previous experiments.

In order to model the opponent, *AutONA* attaches a belief function to each player, which tries to estimate the probability of a price for a given seller and a given quantity. This belief function is updated based on observed prices in prior negotiations. Several tactics and heuristics are implemented to form the strategy of the negotiator during the negotiation process (e.g., for selecting the opponents with which it will negotiate and

for determining the first offer it will suggest to the opponent). Byde *et al.* also allowed cheap-talk during negotiations, that is, the proposition of offers with no commitments.

After reviewing the *AutONA* agent, we proceed with agents that are applicable to a larger family of domains. The next agent is applicable to the Cliff-Edge family of domains. A Cliff-Edge environment is characterized by a conflict between the desire to maximize profits while preventing the entire deal from falling through. An example of such a domain is the Ultimatum Game. The ultimatum game is an experimental economics game in which two players have to decide how to divide a sum of money between them. The first player can propose a division while the second player can either accept the proposal or reject it. Only if the second player accepts the offer the money is split between the two. The game is played only once.

The Cliff-Edge Agent Katz and Kraus [14] proposed an innovative model for human learning and decision making. Their agent competes repeatedly in one-shot interactions, each time against a different human opponent (e.g., sealed-bid first-price auctions, ultimatum game). Katz and Kraus utilized a reinforcement learning algorithm, which integrates virtual learning with reinforcement learning. That is, offers higher than an accepted offer are treated as successful (virtual) offers, notwithstanding that they were not actually proposed. Similarly, offers lower than a rejected offer are treated as having been (virtually) unsuccessfully proposed. A threshold is also employed to allow for some deviations from this strict categorization. The results of previous interactions are stored in a database, which is used for later interactions. The decision making mechanism of Katz and Kraus's Ultimatum Game agent follows a heuristic based on the qualitative theory of Learning Direction [33]. Simply speaking, if an offer is rejected at a given interaction, then at the next interaction the proposer will offer the opponent a higher offer. In contrast, if an offer is accepted, then during the following interaction the offer will be decreased. Katz and Kraus show that their algorithm performs better than other automated agents. When compared to human behavior, there is an advantage to their automated agent over the human's average payoff.

We continue now with agents that are domain-independent and propose an agent with more generality than the aforementioned agents.

The Guessing Heuristic Agent Jonker *et al.* [13] deal with bilateral multi-issue and multi-attribute negotiations which involve incomplete information. The negotiation follows the alternating offer protocol and is conducted once with each opponent. Jonker *et al.* designed a generic agent that uses a "guessing heuristic" in the buyer-seller domain. This heuristic tries to predict the opponent's preferences based on its offers' history. This is under the assumption that the opponent's utility has a linear function structure. Jonker *et al.* assert that this heuristic allows their agent to improve the outcome of the negotiations. Regarding the offer generation mechanism, they use a concession mechanism to obtain the next offer. In their experiments, the automated agent acts as a proxy for the human user. The user is involved only in the beginning when he inputs the preference parameters. Then the agent generates the offers and the counter-offers. When comparing negotiations involving only automated agents with negotiations involving only humans, the agents usually outperformed the humans (in the buyer's role). Yet, in an additional experiment they matched humans versus agent negotiators. In this

experiment, humans only played the role of the buyer. When comparing the human vs. agent negotiations to that of only automated agents, the humans attained somewhat better results than the agents (in the buyer's role), based on the average utilities. The authors believe this should be accounted to the fact that humans forced the automated negotiators to make more concessions than they themselves did.

The next agent we discuss also deals with bilateral multi-issue negotiations that involve incomplete information. Nonetheless the negotiation protocol is richer than that of the *Guessing Heuristic Agent*.

The *QOAgent* The *QOAgent* [22] is a domain independent agent that can negotiate with people in environments of finite horizon bilateral negotiations with incomplete information. The negotiations consider a finite set of multi-attribute issues and time-constraints. Costs are assigned to each negotiator, such that during the negotiation process, the negotiator might gain or lose utility over time. If no agreement is reached by a given deadline a status quo outcome is enforced. A negotiator can also opt-out of the negotiation if it decides that the negotiation is not proceeding in a favorable manner. Similar to the negotiation protocol in the *Diplomat* agent's domain, the negotiation protocol in the *QOAgent*'s domain extends the model of alternating offers such that each agent can perform up to $M > 0$ interactions with the opponent agent during each time period. In addition, queries and promises are allowed which adds unenforceable agreements to the environment. 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.

As for incomplete information, the *QOAgent* tackles the problem by applying a simple Bayesian update mechanism, which, after each action tries to infer which utility best suits the opponent (whether when receiving an offer or when receiving a response to an offer). For the decision making process, the approach used by the *QOAgent* is more of a qualitative approach [35]. 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 *QOAgent* uses the maximin function and the qualitative valuation of offers. Using these methods the *QOAgent* generates offers and decides whether to accept or reject proposals it has received.

Lin *et al.* [22] tested the *QOAgent* in several distinct domains and their results show that the *QOAgent* reaches more agreements and plays more effectively than its human counterparts, when the effectiveness is measured by the score of the individual utility. They also show that the sum of utilities is higher in negotiations when the *QOAgent* is involved, as compared to human-human negotiations. Thus, they assert, that it is indeed possible to build an automated agent that can negotiate successfully with humans. However, it is also important to state that their agent has certain limitations. They assume that there is a finite set of different agent types and thus their agent cannot generate a dynamic model (and perhaps a more accurate one) of the opponent. In addition, they have

not shown whether their agent can also maintain high scores when matched with other automated agents, which is an important characteristic of open environment negotiations. Moreover, the *QOAgent* does not scale well when numerous offers are proposed, which can cause its performance to deteriorate.

The *KBAgent* The *KBAgent* [26] is an automated agent that was designed to improve the performance of the *QOAgent* [22]. The *KBAgent* incorporates a better learning mechanism which enables it to perform better than the *QOAgent*. In essence, the *KBAgent* uses a database of past negotiation sessions between specific agents (types of negotiators) to allow it to be more efficient in negotiations with agents of that specific type. Based on a database with past negotiation sessions, the agent performs offline learning, which is based on the kernel based density estimation ([37], Chapter 2). From the database the agent estimates the probability of an offer to be accepted, the probability of it to be offered and the expected average utility for the other party. These probabilities are then used in its decision making component, either when accepting an offer or to determine the agent's concession rate. Using the kernel based density estimation the *KBAgent* is capable of using even small databases and does not have to rely on many past negotiation sessions. The approach implemented in the *KBAgent* allows the automated agent negotiate efficiently with people and even perform better than another state-of-the-art automated agent. The results demonstrate that the *KBAgent* achieved significantly higher utility values than the human players. In comparison to the other automated agent (the *QOAgent*) it achieved higher utility values and in the case of one of the two roles, even achieved significantly higher utility values.

After having reviewed all the agents we conclude with a brief discussion on the characteristics and the design of future agents.

2.5 The Rule of Thumb for Designing Automated Agents

We should probably begin with the conclusion. Despite the title of this section, there may not be a good rule of thumb for designing automated negotiators with human negotiators. Table 1 summarizes the main contributions made by each of the reviewed agents. If we look into the design of all the aforementioned agents, we cannot find one specific feature that connects them or can account for their good negotiation skills. Nonetheless, we can note several features that have been used in several agents. Agent designers might take these features into consideration when designing their automated agent, while also taking into account the settings and the environment in which their agent will operate.

The first feature is *randomization*, which was used in *Diplomat*, *QOAgent* and also (though not explicitly) in the *Cliff-Edge* agents. The randomization factor allows these agents to be more resilient (or robust) to adversaries that try to manipulate them to gain better results on their part. In addition, it allows them to be more flexible, rather than strict, in accepting agreements and ending negotiations.

The second feature can be viewed as a *concession strategy*. Both the *AutONA* agent and the *Guessing Heuristic* agent implemented this strategy, which influenced the offer generation mechanism of their agent. A concession strategy might also have a psychological effect on the opponent which would make it more comfortable for the opponent to accept agreements or to make concessions on his own as well.

Agent	Main contribution
<i>Diplomat</i>	Changing the agent's personality heuristics Non deterministic behavior / randomization
<i>AutONA</i>	Tactics and heuristics Incorporating data from past interactions Concession mechanism
<i>Cliff-Edge</i>	Virtual learning Incorporating data from past interactions Gender-sensitive approach Non deterministic behavior / randomization (implicitly)
<i>Guessing Heuristic</i>	Generic agent / domain independent Concession mechanism
<i>QOAgent</i>	Generic agent / domain independent Qualitative decision making Non deterministic behavior / randomization
<i>KBAgent</i>	Generic agent / domain independent Qualitative decision making Non deterministic behavior / randomization Incorporating data from past interactions

Table 1. Main contributions of each agent.

The last feature which was common in several agents is the use of a *database*. The database can be built based on previous interactions with the same human opponent or it can be built for all opponents. The agent consults the database to better model the opponent, to learn about possible behaviors and actions and to adjust its behavior to the specific opponent. A database of the history can also be used to obtain information about the behavior of the opponents, if such information is not known, or cannot be characterized, in advance.

Lastly, though not exactly a feature, but worth mentioning, is that none of the agents we reviewed implemented *equilibrium strategies*. This is an interesting observation and most likely is due to the fact that these strategies have been shown to behave poorly when implemented in automated negotiators matched with human negotiators, mainly due to the complex environment and the bounded rationality of people.

3 Validating the Efficacy of the Automated Negotiator

There is a great challenge in validating the efficacy of the automated negotiator. Such an evaluation is important in order to select the most appropriate agent for the task at hand. Yet, no single criteria is defined. The answer to the questions of “what constitutes a good negotiator agent?” is multifaceted. For example, is a good agent an agent that:

- achieves a maximal payoff when matched with human negotiators? But will it also generate these payoffs when matched with other automated agents, which might

be more accessible than human negotiators, and which also exist in open environments?

- generates a maximal combined payoff for both negotiators, that is, the agent is more concerned with maximizing the combined utilities than its own reward?
- allows most negotiations to end with an agreement, rather than one of the sides opting-out or terminating the negotiations with a status-quo outcome?
- is domain dependent and its technique is suitable only for that domain or one that is domain independent and can be adapted to several domains? This might be an important factor if an agent is required to adapt to dynamic settings, for example.
- behave in such a manner that would leave its counterpart speculating whether it is an automated negotiator or a human one?

There probably does not exist a single solution to the question. Yet, researchers should take these and other measures into consideration when designing their agents. Perhaps, certain criteria and benchmarks are in order to allow an adequate comparison between automated agents. In this respect, researchers should sought using a standardized objective measures for the evaluation of their agents (as, for example, proposed by [23]).

After choosing the criteria, the automated negotiators should be matched with other negotiators. As people are scarce resources it is quite tempting to match the automated negotiators with other automated negotiators. We continue to provide results of experimentation with automated negotiators and show that there is no real replacement for validating the automated negotiator against people.

It is important to note though that when experimenting with people, as people are scarce resources, one should maximize the data obtained from them. Using questionnaires to gain subjective information, thoughts and suggestions from the subjects is important, even if it will not be later used for academic purposes. The thoughts of the subjects can provide insights with respect to the experiment design, the user interface and others aspects and could motivate the researcher in improving these issues. If the experiments involve several stages, one should consider whether several questionnaires should be used, between each and every stage. Nonetheless, the design of the questionnaires should be made carefully. Omitting important questions and remembering of them later on, will make it hard to get back to the subjects and request them to fill in the missing data, merely due to the effect of time and memory of the subjects.

We begin with automated negotiators designed specifically to serve as efficient negotiators with people. To this end, we designed the *QOAgent* [22] and the *KBAgent* [26] (refer to Section 2.4 for an overview on these agents). Results reported in the cited papers validate the efficacy of these agents in negotiations with people. Yet, their behavior against other automated agents was quite different. Matching the automated agent and validating its efficacy when matched with other automated agents, as the ones described in this paper, was an infeasible task. Most agents are designed for a specific environment and protocol (with regard to the negotiation protocol and the communication protocol), requiring re-programming the agents and writing adapters for each to try and matching all in the same environment (for example, as proposed by [12]). In this respect, and not in the scope of this paper, we can also mention that there is an ongoing work to alleviate these difficulties and an effort to create a negotiation environment enabling the design of agents and their empiric evaluation, by objectively comparing the agents with other

agents designed in this system (e.g., the GENIUS environment [21]). Thus, we took a different approach. We conducted an experiment in which we matched the *KBAgent* along with 60 other peer designed agents (PDAs). Each agent was matched with all other agents. The PDAs were designed by computer science students that were given a task to implement an efficient automated agent for a domain in which an employee and a job candidate are negotiating to formalize the hiring terms and conditions of the job candidate. The students were provided skeleton classes to help them implement their agents. This also allowed them to focus on the strategy and the behavior of the agent, and eliminate the need to implement the communication protocol or the negotiation protocol. In addition, it provided them with a simulation environment in which they could test their agents and their strategies.

The results were quite surprising. While the *KBAgent* was shown to be an efficient automated negotiator with people, it performed worse than some of the PDAs in some of the roles in the negotiation. On the other hand, taking an automated negotiator which achieves good results against other automated agent does not guarantee it will achieve comparable results against people. To this end we took the *CT* agent developed by Talman *et al.* [34]. The *CT* agent was developed for the Colored Trails game environment [11], 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.

Talman *et al.*'s agent was shown to outperform other automated agents as well as increased the social welfare of the automated agents which participated in the game. Yet, when this agent was matched with other people its performance degraded poorly. When analyzing the results the conclusion as for the variance in its behavior is mainly accounted for its offer proposal mechanism. The offers proposed by the agent were reasonable offers, yet not "intuitive". While other automated agents applied rational considerations and accepted the offers, people tend to discard the offers and it even generated antagonism from the people when the agent continued and proposed these kind of offers. This eventually caused the negotiation to terminal unexpectedly, resulting with poor outcomes. It is worth mentioning that this result is not unique. For example, it was shown that an agent following equilibrium (thus, when playing with other automated agents might reach Pareto optimal outcomes), when matched with people, results with the people becoming frustrated, mainly since the automated agent repeatedly proposes the same offer, and the negotiation often ends with no agreement [18]. This has been shown in cases in which the complexity of finding the equilibrium is low and the players have full information. The solution to both cases was to apply some rules, constraints and heuristics to the agent's behavior to adapt it to people it is matched with.

4 Conclusions

The importance of designing an automated negotiator that can negotiate efficiently with humans cannot be understated. In this papers we presented lessons learned by several agents designed to specifically negotiation with people. There are probably more lessons to be sought, yet this paper tried to focus on the most important ones. By pursu-

ing non-classical methods of decision making and a learning mechanism for modeling the opponent it could be possible to achieve greater flexibility and effective outcomes.

To this date, it seems that research in Multi-agent systems (MAS) has neglected the issue of proficiently negotiating with people, at the expense of designing automated agents aimed to negotiate with rational agents or other automated agents (e.g., [38]). Others have focused on improving different heuristics and strategies and the analysis of game theory aspects (e.g., [17, 25]). Nonetheless it is noteworthy that these are important aspects in which the MAS community has certainly made an impact. Unfortunately, not much progress has been made with regard to automated negotiators with people, leaving many unfaced challenges. We hope that this paper could alleviate some of the concerns involved in the design process of automated negotiators and experimenting with people and motivate researches in dwelling to this exciting line of research.

References

1. M. H. Bazerman and M. A. Neale. Negotiator rationality and negotiator cognition: The interactive roles of prescriptive and descriptive research. In H. P. Young, editor, *Negotiation Analysis*, pages 109–130. The University of Michigan Press, 1992.
2. 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.
3. P. J. Carnevale and C. K. W. De Dreu. Laboratory experiments on negotiation and social conflict. *International Negotiation*, 10(1):51–66, 2005.
4. A. Chavez and P. Maes. Kasbah: An agent marketplace for buying and selling goods. In *Proceedings of the first international Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology*, pages 75–90, 1996.
5. R. Cohen. *Negotiating Across Cultures: Communication Obstacles in International Diplomacy*. United States Institute of Peace Press, Washington, D.C., 1991.
6. J. W. Creswell. *Research design: qualitative, quantitative, and mixed methods approaches*. Sage Publishing, Thousand Oaks, CA, 2003.
7. 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.
8. J. Farrell and M. Rabin. Cheap talk. *Journal of Economic Perspectives*, 10(3):103118, 1996.
9. R. Fisher and W. Ury. *Getting to Yes: Negotiating Agreement Without Giving In*. Penguin Books, 1991.
10. J. Grossklags and C. Schmidt. Software agents and market (in) efficiency: a human trader experiment. *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews*, 36(1):56–67, 2006.
11. 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.
12. K. Hindriks, C. Jonker, and D. Tykhonov. Towards an open negotiation architecture for heterogeneous agents. In *CIA*, volume 5180 of *LNAI*, pages 264–279. Springer, 2008.
13. 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.
14. R. Katz and S. Kraus. Efficient agents for cliff edge environments with a large set of decision options. In *AAMAS*, pages 697–704, 2006.
15. R. Keeney and H. Raiffa. *Decisions with Multiple Objective: Preferences and Value Trade-offs*. John Wiley & Sons, NY, USA, 1976.
16. P. Kenny, A. Hartholt, J. Gratch, W. Swartout, D. Traum, S. Marsella, and D. Piepol. Building interactive virtual humans for training environments. In *IITSEC*, 2007.

17. S. Kraus. *Strategic Negotiation in Multiagent Environments*. MIT Press, Cambridge MA, USA, 2001.
18. S. Kraus, P. Hoz-Weiss, J. Wilkenfeld, D. R. Andersen, and A. Pate. Resolving crises through automated bilateral negotiations. *Artificial Intelligence*, 172(1):1–18, 2008.
19. S. Kraus and D. Lehmann. Designing and building a negotiating automated agent. *Computational Intelligence*, 11(1):132–171, 1995.
20. M. LeBaron and V. Pillay. *Conflict Across Cultures: A Unique Experience of Bridging Differences*. Nicholas Brealey Publishing, Boston, MA, 2006.
21. R. Lin, S. Kraus, D. Tykhonov, K. Hindriks, and C. M. Jonker. Supporting the design of general automated negotiators. In *ACAN*, 2009.
22. R. Lin, S. Kraus, J. Wilkenfeld, and J. Barry. Negotiating with bounded rational agents in environments with incomplete information using an automated agent. *Artificial Intelligence*, 172(6-7):823–851, 2008.
23. R. Lin, Y. Oshrat, and S. Kraus. Investigating the benefits of automated negotiations in enhancing people’s negotiation skills. In *AAMAS*, pages 345–352, 2009.
24. R. D. McKelvey and T. R. Palfrey. An experimental study of the centipede game. *Econometrica*, 60(4):803–836, 1992.
25. A. Muthoo. *Bargaining Theory with Applications*. Cambridge University Press, 1999.
26. Y. Oshrat, R. Lin, and S. Kraus. Facing the challenge of human-agent negotiations via effective general opponent modeling. In *AAMAS*, 2009.
27. H. Raiffa. *The Art and Science of Negotiation*. Harvard University Press, 1982.
28. E. Rasmusen. *Games and Information: An Introduction to Game Theory*. Blackwell Publishers, 2001.
29. W. Ross and J. LaCroix. Multiple meanings of trust in negotiation theory and research: A literature review and integrative model. *International Journal of Conflict Management*, 7(4):314–360, 1996.
30. A. Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):97–109, 1982.
31. A. Rubinstein. A bargaining model with incomplete information about preferences. *Econometrica*, 53(5):1151–1172, 1985.
32. A. Sanfey, J. Rilling, J. Aronson, L. Nystrom, and J. Cohen. The neural basis of economic decision-making in the ultimatum game. *Science*, 300:1755–1758, 2003.
33. R. Selten and R. Stoecker. End behavior in sequences of finite prisoner’s dilemma supergames: A learning theory approach. *Economic Behavior and Organization*, 7(1):47–70, 1986.
34. S. Talman, Y. Gal, M. Hadad, and S. Kraus. Adapting to agents’ personalities in negotiation. In *AAMAS*, pages 383–389, 2005.
35. M. Tennenholtz. On stable social laws and qualitative equilibrium for risk-averse agents. In *KR*, pages 553–561, 1996.
36. A. Tversky and D. Kahneman. The framing of decisions and the psychology of choice. *Science*, 211:453–458, 1981.
37. M. Wand and M. Jones. *Kernel Smoothing*. Chapman & Hall, London, 1995.
38. M. P. Wellman, A. Greenwald, and P. Stone. *Autonomous Bidding Agents: Strategies and Lessons from the Trading Agent Competition*. MIT Press, Cambridge MA, USA, 2007.
39. J. Wilkenfeld. Reflections on simulation and experimentation in the study of negotiation. *International Negotiation*, 9(3):429–440, 2004.
40. X. Zhang, V. Lesser, and R. Podorozhny. Multi-dimensional, multistep negotiation for task allocation in a cooperative system. *JAAMAS*, 10(1):5–40, 2005.