Benefits of Learning in Negotiation

Dajun Zeng **Katia Sycara**
The Robotics Institute
Carnegie Mellon University, Pittsburgh, PA 15213
Voice: (412) 268-8815 Fax: (412) 268-5569
Dajun.Zeng@cs.cmu.edu katia@cs.cmu.edu

Abstract

Negotiation has been extensively discussed in game-theoretic, economic, and management science literatures for decades. Recent growing interest in electronic commerce has given increased importance to automated negotiation. Evidence both from theoretical analysis and from observations of human interactions suggests that if decision makers can somehow take into consideration what other agents are thinking and furthermore learn during their interactions how other agents behave, their payoff might increase. In this paper, we propose a sequential decision making model of negotiation, called Bazaar. Within the proposed negotiation framework, we model learning as a Bayesian belief update process. In this paper, we explore the hypothesis that learning is beneficial in sequential negotiation and present initial experimental results.

Introduction

Recent growing interest in autonomous interacting software agents and their potential application in areas such as electronic commerce (e.g., (Sandholm & Lesser 1995)) has given increased importance to automated negotiation. Much DAI and game theoretic research (Rosenschein & Zlotkin 1994; Osborne & Rubinstein 1994; Kraus & Subrahmanian 1995) deals with coordination and negotiation issues by giving pre-computed solutions to specific problems. There has been much research reported on developing theoretical models in which learning plays an eminent role, especially in the area of adaptive dynamics of games (e.g., (Jordan 1992; Kalai & Lehrer 1993)). However, to build autonomous agents that improve their negotiation competence based on learning from their interactions with other agents is still an emerging area.

Learning in negotiation is closely coupled with the issue of how to model the overall negotiation process, i.e., what negotiation protocols are adopted. Standard game-theoretic models (Osborne & Rubinstein 1994) tend to focus on *outcomes* of negotiation in contrast to the *negotiation process* itself. DAI research (Rosenschein & Zlotkin 1994) emphasizes special protocols articulating compromises while trying to minimize the potential interactions or communications of the involved agents. Since we are motivated by a different set of research issues, such as including effective learning mechanisms in the negotiation process, we adopt a different modeling framework, i.e., a sequential decision making paradigm (Bertsekas 1995; Cyert & DeGroot 1987).

The basic characteristic of a sequential decision process is that there is a sequence of decision making points (different stages) which are dependent on each other. In a sequential making process, the decision maker has a chance to update his knowledge after implementing the decision made at a certain stage and receiving feedback. In most negotiation models, however, there is no feedback, therefore no chance of belief updating. In this paper, we propose a sequential decision making model, called Bazaar, which is able to learn. We address multi-agent learning issues in Bazaar by explicitly modeling beliefs about the negotiation environment and the participating agents under a probabilistic framework using a Bayesian learning representation and updating mechanism. We also report our initial experimental results in a simple bargaining scenario. Our ultimate research goal is to develop an adaptive negotiation model capable of exhibiting a rich set of negotiation behaviors with modest computational effort.

Sequential Decision Making with Rational Learning

Our overall research goal is to develop a computational model of negotiation that can handle multi-agent learning and other complicated issues (e.g., multi-issue multi-criteria negotiation) that don't have straightforward and computationally efficient analytic models. We believe that a useful computational model of negotiation should exhibit the following characteristics: (1) The model should support a concise yet effective way to represent negotiation context. (2) The model

[2]*This paper appears in* Proceedings of AAAI-97.

should be prescriptive in nature. (3) The computational resources required for finding reasonable suggestions/solutions should be moderate, sometimes at the cost of compromising the rigor of the model and the optimality of solutions. (4) The model should provide means to model the dynamics of negotiation. (5) The model should also support the learning capability of participating agents. Motivated by these desirable features, we have developed Bazaar, a sequential decision making negotiation model that is capable of learning.

Bazaar: a formal description

▶ In Bazaar, a negotiation process can be modeled by a 10-tuple
$< N, M, \Delta, A, H, Q, \Omega, P, C, E >$, where,

A-1 A set N (the set of players).

A-2 A set M (the set of issues/dimensions covered in negotiation. For instance, in the supply chain management domain, this set could include product price, product quality, payment method, transportation method, etc.)

A-3 A set of vectors $\Delta \equiv \{(D_j)_{j\in M}\}$ (a set of vectors whose elements describe each and every dimension of an agreement under negotiation).

A set A composed of all the possible actions that can be taken by every member of the players set.
- ▷ $A \equiv \Delta \cup \{Accept, Quit\}$

A-4 For each player $i \in N$ a set of possible agreements A_i.
- ▷ For each $i \in N$, $A_i \subset A$.

A-5 A set H of sequences (finite or infinite) that satisfies the following properties:
- ▷ The elements of each sequence are defined over A.
- ▷ The empty sequence Φ is a member of H.
- ▷ If $(a^k)_{k=1,\ldots,K} \in H$ and $L < K$ then $(a^k)_{k=1,\ldots,L} \in H$.
- ▷ If $(a^k)_{k=1,\ldots,K} \in H$ and $a^K \in \{Accept, Quit\}$ then $a^k \overline{\in} \{Accept, Quit\}$ when $k = 1, \ldots, K-1$.

Each member of H is a *history*; each component of a history is an action taken by a player. A history $(a^k)_{k=1,\ldots,K}$ is terminal if there is no a_{K+1} such that $(a^k)_{k=1,\ldots,K+1} \in H$. The set of terminal histories is denoted by Z.

A-6 A function Q that associates each nonterminal history ($h \in H \setminus Z$) to a member of N. (Q is called the *player function* which determines the orderings of agent responses.)

A-7 A set of Ω of relevant information entities. Ω is introduced to represent the players' knowledge and belief about the following aspects of negotiation:
- ▷ The parameters of the environment, which can change over time. For example, in supply chain management, global economic or industry-wide indices such as overall product supply and demand and interest rate, belong to Ω.
- ▷ Beliefs about other players. These beliefs can be approximately decomposed into three categories:
  - (a) Beliefs about the factual aspects of other agents, such as how their payoff functions are structured, how many resources they have, etc.
  - (b) Beliefs about the decision making process of other agents. For example, what would be other players' reservation prices.
  - (c) Beliefs about meta-level issues such as the overall *negotiation style* of other players. Are they tough or compliant? How would they perceive a certain action? What about their risk-taking attitudes? etc.

A-8 For each nonterminal history h and each player $i \in N$, a subjective probability distribution $P_{h,i}$ defined over Ω. This distribution is a concise representation of the knowledge held by each player in each stage of negotiation.

A-9 For each player $i \in N$, each nonterminal history $H \setminus Z$, and each action $a_i \in A_i$, there is an implementation cost C_{i,h,a_i}. C can be interpreted as communication costs or costs associated with time caused by delaying terminal action (*Accept* or *Quit*).

A-10 For each player $i \in N$ a *preference* relation $\succeq_i$ on Z and $P_{h,i}$ for each $h \in Z$. $\succeq_i$ in turn results in an evaluation function $E_i(Z, P_{Z,i})$.

We will present the solution strategy in Bazaar before we discuss the characteristics of the model.

Solution Strategy in Bazaar

Although the role that the players play (e.g., selling or buying) with respect to initiating the negotiation process can have an impact[1] , the decision making process in a negotiation scenario, namely, determining the particular contents of an offer/counter-offer (quit and accept can be viewed as a special offer), is symmetrical for all the players. So the following solution framework is not limited by roles of the players:

▶ For each player i, a negotiation strategy is a sequence of actions $(a_i^k, k = 1, \ldots, K)$, where
- k denotes that a_i^k is the kth action ($k \leq K$) taken by i,
- $a_i^k \in A_i$,
- $a_i^K \in \{Accept, Quit\}$,
- $a_i^k \overline{\in} \{Accept, Quit\}$ when $k = 1, \ldots, K-1$

▶ Before negotiation starts, each player has a certain amount of knowledge about Ω, which may include the knowledge about the environment where the negotiation takes place, and may also include the prior knowledge about other players (from previous experience or from second-hand knowledge, or rumors,

[1]For example, in a 2-player supply chain situation, the supplier often is the first one to initiate a negotiation.

etc.) This prior knowledge is denoted (see A-8) as $P_{\Phi,i}$.

► Suppose player i has been interacting with another player j for k times. In other words, i has sent exactly k offers or counter-offers to j (presumably received k or $k+1$ offers or counter-offers from j depending on who initiated the negotiation process). Let's assume that neither *Accept* nor *Quit* has appeared in these offers and counter-offers. In Bazaar, the following information is available when i tries to figure out what to do next (the content of its $k+1$th offer):

1. All the actions taken by all the agents up to the current time point when i makes decision about the $k+1$th offer. Formally, each and every history h that is a sequence of k actions is known to i. Let's denote this set of histories by $H_{i,k}$.
2. The set of subjective probability distributions over Ω, $P_{H_{i,k-1},i} \equiv \{P_{h,i} | h \in H_{i,k-1}\}$ is known to i.

A player takes the following steps to decide how to reply to the most recent action taken by other participant(s):

Step 1 Update his subjective evaluation about the environment and other players using Bayesian rules: Given prior distribution $P_{H_{i,k-1},i}$ and newly incoming information $H_{i,k}$, calculate the posterior distribution $P_{H_{i,k},i}$.

Step 2 Select the best action $\widehat{a}_{k+1,i}$ out of A_i according to the following recursive evaluation criteria:
$V_{i,k,H_{i,k}} = E_i(H_{i,k}, P_{H_{i,k},i})$ when $H_{i,k} \in Z$
$V_{i,k,H_{i,k}} = \max_{a_i \in A_i}(-C_{i,H_{i,k},a_i} + \int_\Omega (V_{i,k+1,H_{i,k}|a_i} \times P_{H_{i,k},i}) d\Omega)$
The first equation represents the termination criterion. The second equation can be summarized as "always choose the action that maximizes the expected payoff given the information available at this stage". The implementation cost C at this stage has been deducted from the future (expected) payoff.

Learning in Negotiation

The importance of learning in negotiation has been recently recognized in the game research community as fundamental for understanding human behavior as well as for developing new solution concepts (Osborne & Rubinstein 1994; Harsanyi & Selten 1972). In (Jordan 1992) the author studied the impact of Bayesian learning processes for finite-strategy normal form games. Kalai and Lehrer (Kalai & Lehrer 1993) analyzed infinitely repeated games in which players as subjective utility maximizers learn to predict opponents' future strategies. These theoretical results, however, are available only for the simplest game settings and valid only under very restrictive assumptions such as only a subset of possible negotiation strategies are allowed.

Multi-agent learning has also increasingly drawn research efforts from Distributed AI community. Mor *et. al.* (Mor, Goldman, & Rosenschein 1995) discussed multi-agent learning as a means to reach equilibrium. They modeled agents as finite automata and analyzed the computational complexity of certain classes of learning strategies based on this automaton model. In (Sen & Sekaran 1995) the authors demonstrated that some simple agent adaptive behaviors based on reciprocity allow agents to produce satisfactory global performance. In the context of Bazaar, we are using the Bayesian framework to update the knowledge and belief that each agent has about the environment and other agents.

In this section, we use a buyer-supplier example used before to demonstrate how the Bayesian framework can be utilized in a negotiation setting. For illustrative purposes, we consider the negotiation process only from the viewpoint of the buyer and assume that the relevant information set Ω is comprised of only one item: belief about the supplier's reservation price $RP_{supplier}$. An agent's reservation price is the agent's threshold of offer acceptability. Typically a reservation price is private to each agent, and is different for each agent for each negotiation issue. For example, a supplier's reservation price is the price such that the supplier agent will not accept an offer below this price; a buyer's reservation price is the price such that the buyer will not accept an offer above this price. As shown in Figure 1, when the supplier's reservation price $RP_{supplier}$ is lower than the buyer's reservation price RP_{buyer}, any point within the "zone of agreement" is a candidate solution; while, if RP_{buyer} is lower than $RP_{supplier}$, as shown in Figure 2, the zone of agreement doesn't exist and no deal can be reached via negotiation. If a zone of agreement exists, typically both the buyer and the supplier will make concessions from their initial proposal. The buyer will increase his initial proposal, while the supplier will decrease his. Eventually, a proposal within the zone of agreement will be acceptable to both.

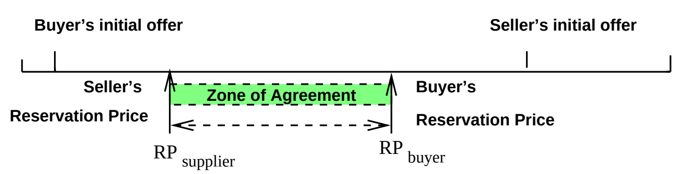

Figure 1: An example of reservation prices and "zone of agreement"

It is obvious that although the buyer knows his own reservation price, the precise value of $RP_{supplier}$ is unknown to him. Therefore, the zone of agreement is not known by either of the agents. Nevertheless, the buyer could update his belief (learn) about $RP_{supplier}$ based on his interactions with the supplier and on his domain knowledge. As a result of learning, the buyer

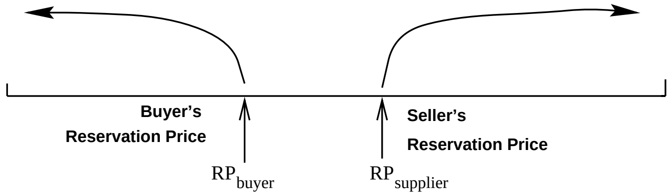

Figure 2: An example in which no "zone of agreement" exists

is expected to gain more accurate expectation of the supplier's payoff structure and therefore make more advantageous offers. In this example, we show how the buyer's belief about $RP_{supplier}$ can be updated during negotiation.

The buyer's partial belief about $RP_{supplier}$ can be represented by a set of hypotheses $H_i, i = 1, 2, \ldots, n$. For instance, H_1 can be "$RP_{supplier} = \$100.00$"; H_2 "$RP_{supplier} = \$130.00$". A priori knowledge held by the buyer can be summarized as probabilistic evaluation over the set of hypotheses $\{H_i\}$ (e.g., $P(H_1) = 0.2$, $P(H_2) = 0.35$, $\ldots$). The Bayesian updating occurs when the buyer receives new signals from the outside environment or from the supplier. Along with domain-specific knowledge, these new signals enable the buyer to acquire new insights about $RP_{supplier}$ in the form of posterior subjective evaluation over H_i. In our case, the offers and counter-offers ($\text{Offer}_{supplier}$) from the supplier comprise the incoming signal; while the domain knowledge can be an observation such as: "Usually in our business people will offer a price which is above their reservation price by 17%", which can be represented by a set of conditional statements of similar form, one of which is shown as follows: $P(e_2 \mid H_2) = 0.95$, where e_2 represents "$\text{Offer}_{supplier} = \$152.1$", and H_2 "$RP_{supplier} = \$130.00$".

Given the encoded domain knowledge in the form of conditional statements and the signal (e) in the form of offers made by the supplier, the buyer can use the standard Bayesian updating rule to revise his belief about $RP_{supplier}$: $P(H_i \mid e) = \frac{P(H_i)P(e|H_i)}{\sum_{k=1}^{n} P(e|H_k)P(H_k)}$

We use a numerical example to show how this updating works. For simplicity, we suppose that the buyer knows that the supplier's reservation price is either \$100.00 or \$130.00. In other words, the buyer has only two hypotheses: H_1: "$RP_{supplier} = \$100.00$" and H_2: "$RP_{supplier} = \$130.00$".

At the beginning of the negotiation, the buyer doesn't have any other additional information. His a priori knowledge can be summarized as: $P(H_1) = 0.5, P(H_2) = 0.5$.

In addition, we suppose that the buyer is aware of "Suppliers will typically offer a price which is above their reservation price by 17%", part of which is encoded as: $P(e_1 \mid H_1) = 0.95$ and $P(e_1 \mid H_2) = 0.75$, where e_1 denotes the event that the supplier asks \$117.00 for the goods under negotiation.

Now suppose that the supplier offers \$117.00 for the product the buyer wants to purchase. Given this signal and the domain knowledge, the buyer can calculate the posterior estimation of $RP_{supplier}$ as follows: $P(H_1 \mid e_1) = \frac{P(H_1)P(e_1|H_1)}{P(H_1)P(e_1|H_1)+P(H_2)P(e_1|H_2)} = 55.9\%$; $P(H_2 \mid e_1) = \frac{P(H_2)P(e_1|H_2)}{P(H_2)P(e_1|H_1)+P(H_2)P(e_1|H_2)} = 44.1\%$

Suppose that the buyer adopts a simple negotiation strategy: "Propose a price which is 10% below the estimated $RP_{supplier}$". Prior to receiving the supplier's offer (\$117.00), the buyer would propose \$115.00 (the mean of the $RP_{supplier}$ subjective distribution). After receiving the offer from the supplier and updating his belief about $RP_{supplier}$, the buyer will propose \$113.23 instead. Since the new offer is calculated based on a more accurate estimation of the supplier's utility structure, it might result in a potentially more beneficial final outcome for the buyer and may also help both sides reach the agreement more efficiently.

Experimental Study: Learning in Bargaining

In our initial experiments, we consider a simple bargaining scenario with the following characteristics:

- The set of players N is comprised of one buyer and one supplier
- The set of dimensions M contains only one issue, *price*
- For simplicity, the range of possible prices is from 0 to 100 units
- The set of possible actions (proposed prices by either the buyer or the supplier) A equals to $\{0, 1, 2, \ldots, 100\}$
- The player function Q is defined in such a way that the buyer and the supplier make alternate proposals. Who will be proposing first is decided by coin-tossing
- For simplicity, the relevant information set Ω contains only the supplier's reservation price RP_s and the buyer's reservation price RP_b
- Reservation prices are private information. In other words, each player only knows his own reservation price
- The range of possible prices is public information
- Each player's utility is linear to the final price (a number between 0 and 100) accepted by both players
- Each agent is allowed to propose only strictly monotonically. For example, the supplier's subsequent offers will increase monotonically, while the buyer's offers will decrease monotonically. They are not allowed to propose the same value more than once

Since the negotiation process is symmetrical for the buyer and the supplier, the following discussions about the strategies with or without learning apply to both agents. In our experiments, by the *non-learning* *agents*, we mean the agent that makes his decision based solely on his own reservation price. For instance, the supplier may start proposing 100 initially. The buyer deems it unacceptable and proposes another value. Since the non-learning supplier does not have a model of the buyer (in terms of the buyer's reservation price), the supplier just compares the buyer's offer with his own reservation price RP_s. If the buyer's offer exceeds RP_s, the supplier will accept the offer and the negotiation process ends. If not, the supplier will propose a value which is below his previous offer by a fixed percentage [2] but above RP_s (the actual value will be rounded up to an integer value). The *learning* agent's negotiation strategy is fundamentally different. Decisions will be made based on both the agent's own and the opponent's reservation price. Note that reservation prices are private information and there is no way that the agent can know the exact value of his opponent's reservation price, even after an agreement has been reached. However, each learning agent can have some a priori estimation about his opponent's reservation price and update his estimation during the negotiation process using the Bayesian updating mechanism. In our implementation, an agent represents his subjective beliefs about his opponent's reservation price using a piecewise probability distribution function. This function is implemented as a vector with 101 elements $\vec{\Pi} = [P_0, P_1, \ldots, P_{100}]$. In this vector, P_i represents the agent's current estimation of the probability that his opponent's reservation price is i. The current estimation of his opponent's reservation price itself is calculated as the mean $\sum_{i=0}^{100} i * P_i/101$. In general, the buyer and the supplier will have different initial set of subjective belief vectors $\vec{\Pi}_s^0$ and $\vec{\Pi}_b^0$.

The buyer and the supplier have a different set of conditional probability functions. Let's take the buyer's standpoint. One of the conditional distribution function DK_i^b represents the distribution of the possible proposals made by the supplier given that $RP_s = i$. Figure 3 shows the shape of DK_i^b. In essence, this function says that with a very high probability the supplier will propose some percent[3] above his true reservation price b. Higher or lower than that is less probable. Similar functions are defined for the supplier as well.

Experimental Results

We conducted experiments in three different settings:

1. non-learning buyer vs. non-learning supplier
2. learning buyer vs. learning supplier
3. learning buyer vs. non-learning supplier

For each configuration, we ran 500 random experiments. Each experiment instance corresponds to a complete bargaining scenario which involves multiple rounds of exchanging proposals and counter-proposals. We generated these 500 random experimental instances by creating 500 pairs of random numbers. Out of each pair, the lower end, representing the supplier's reservation price, was a realization of a random number that is uniformly distributed in the interval [0..49]. The upper end, representing the buyer's reservation price, was a realization of a random number that is uniformly distributed in the interval [50..100]. In this way, we ensured that the zone of agreement always exists. Note that learning takes place *within* each run of the experiment rather than between runs.

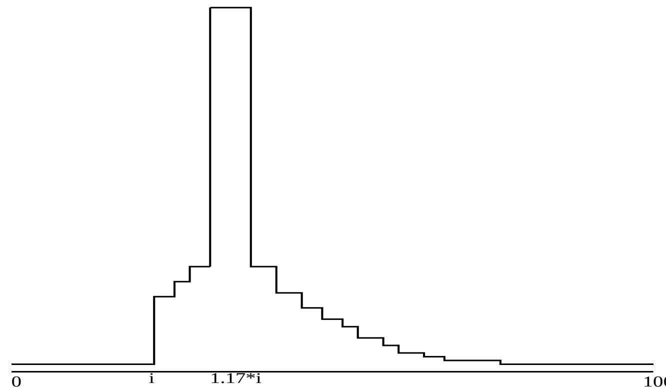

Figure 3: An example of conditional probability function DK_i^b

Configuration	Joint Utility	# of Proposals exchanged
both learn	0.22	24
neither learn	0.18	34
only buyer learns	0.15	28

Table 1: Average Performance of Three Experimental Configurations

We measured the quality of a particular bargaining process using the normalized joint utility fashioned after the Nash solution(Luce & Raiffa 1957). Suppose the buyer and the supplier agree on a particular price P_*, the joint utility is then defined as: $\frac{(P_* - RP_s) \times (RP_b - P_*)}{(RP_b - RP_s)^2}$ It can be easily shown that the joint utility reaches the maximum 0.25 when P_* is the arithmetic average of RP_b and RP_s. Note that in our experimental setting this theoretic maximum might not be reached, for RP_b and RP_s are not known to both agents.

The cost of a bargaining process is measured by the number of proposals exchanged before reaching an agreement. We report in Table 1 the average performance of all three configurations. Our observations about these experimental results are as follows.

- We noticed that in terms of overall bargaining quality and number of proposals exchanged to reach a

[2] In our experiments, the percentage was arbitrarily set to 1.5%.

[3] In our experiments, we arbitrarily set this percent 30%.

compromise, the “both learn” configuration outperformed the other two. This confirmed our intuition that building learning capability into agents' decision making helps agents form more accurate model of the opponent and results in better performance and less expensive process.

- Judged from the viewpoint of the joint utility, the “only buyer learns” configuration does less well compared with “both learn”. In effect, it is even worse than “neither learn”. A careful examination of data reveals that although the joint utility suffers, the buyer (the only learning agent) actually did consistently better for himself (in terms of maximizing his own individual utility) than he did in the “both learn” configuration. We suspect the reason is that the buyer has formed better estimation of his non-learning opponent's reservation price and therefore takes advantage of the “dummy” supplier. Since the optimal Nash solution requires an even split in the zone of the agreement, the buyer-dominant solution leads to lower joint utility. The “neither learn” configuration doesn't show any consistent bias either in favor of the buyer or the supplier.

We examined the data of “neither learn and “both learn” in more detail by further dividing all the 500 experiment instances (1000 instances altogether for both configurations) into different categories according to the size of the zone of agreement. Then, we calculated the differences of the corresponding joint utilities between “neither learn” and “both learn” and plotted the percentage difference in joint utility improvement against the size of the zone of agreement. The result is shown in Figure 4. We observed that there seems to be a positive correlation between these two variables. An intuitive explanation could be that the greater the room for agreement flexibility (greater the zone of agreement), the better the learning agents seize the opportunity.

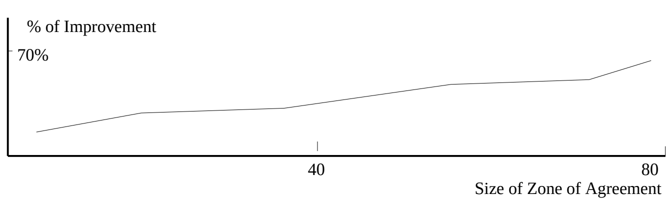

Figure 4: Relations between the Size of the Zone of Agreement & Percent Improvement of Joint Utility

Concluding Remarks and Future work

In this paper, we presented Bazaar, a sequential decision making model of negotiation in which multiagent learning is an integral construct of the model. This model is motivated by providing a computational framework for negotiation which satisfies the following features: (1) the model provides a solution strategy to guide offers instead of only prescribing the final outcome, and (2) learning can be easily incorporated into the model. Initial experiments show that learning is beneficial in this sequential negotiation model. Current work focuses on conducting more extensive experiments and theoretical analysis of the impact of learning under various conditions. Future work will investigate the application of the Bazaar framework on nontrivial negotiation scenarios such as supply chain management.

Acknowledgements

This research has been sponsored in part by ONR Grant #N-00014-96-1-1222, by ARPA Grant #F33615-93-1-1330, and by NSF grant #IRI-9612131.

References

Bertsekas, D. P. 1995. *Dynamic Programming and Optimal Control.* Belmont, MA: Athena Scientific.

Cyert, R. M., and DeGroot, M. H. 1987. *Bayesian Analysis and Uncertainty in Economic Theory.* Rowman & Littlefield.

Harsanyi, J. C., and Selten, R. 1972. A generalized nash solution for two-person bargaining games with incomplete information. *Management Science* 18:80–106.

Jordan, J. S. 1992. The exponential covergence of bayesian learning in normal form games. *Games and Economic Behavior* 4:202–217.

Kalai, E., and Lehrer, E. 1993. Rational learning leads to nash equilibrium. *Econometrica* 61(5):1019–1045.

Kraus, S., and Subrahmanian, V. S. 1995. Multiagent reasoning with probability, time, and beliefs. *International Journal of Intelligent Systems* 10(5):459–499.

Luce, R. D., and Raiffa, H. 1957. *Games and Decisions: Introduction and Critical Survey.* New York, Wiley.

Mor, Y.; Goldman, C. V.; and Rosenschein, J. S. 1995. Learn your opponent's strategy (in polynomial time). In *Proceedings of IJCAI-95 Workshop on Adaptation and Learning in Multiagent Systems.*

Osborne, M. J., and Rubinstein, A. 1994. *A Course in Game Theory.* The MIT Press.

Rosenschein, J., and Zlotkin, G. 1994. *Rules of Encounter.* Cambridge, Mass.: MIT Press.

Sandholm, T. W., and Lesser, V. R. 1995. Coalition fromation among bounded rational agents. In *Proceedings of 14th International Joint Conference on Artificial Intelligence.*

Sen, S., and Sekaran, M. 1995. Multiagent coordination with learning classifier systems. In *Proceedings of IJCAI-95 Workshop on Adaptation and Learning in Multiagent Systems.*