



OpenAIR@RGU

The Open Access Institutional Repository at The Robert Gordon University

<http://openair.rgu.ac.uk>

This is an author produced version of a paper published in

Decision Support Systems (ISSN 0167-9236)

This version may not include final proof corrections and does not include published layout or pagination.

Citation Details

Citation for the version of the work held in 'OpenAIR@RGU':

LAU, R. Y. K., LI, Y., SONG, D. W. and KWOK, R. C. W., 2008. Knowledge discovery for adaptive negotiation agents in e-marketplaces. Available from OpenAIR@RGU. [online]. Available from: <http://openair.rgu.ac.uk>

Citation for the publisher's version:

LAU, R. Y. K., LI, Y., SONG, D. W. and KWOK, R. C. W., 2008. Knowledge discovery for adaptive negotiation agents in e-marketplaces. Decision Support Systems, 45 (2), pp. 310-323.

Copyright

Items in 'OpenAIR@RGU', The Robert Gordon University Open Access Institutional Repository, are protected by copyright and intellectual property law. If you believe that any material held in 'OpenAIR@RGU' infringes copyright, please contact openair-help@rgu.ac.uk with details. The item will be removed from the repository while the claim is investigated.

Knowledge Discovery for Adaptive Negotiation Agents in e-Marketplaces

Raymond Y.K. Lau

Department of Information Systems

City University of Hong Kong

Tat Chee Avenue, Kowloon Hong Kong SAR

Yuefeng Li

School of Software Engineering and Data Communications

Queensland University of Technology

GPO Box 2434, Brisbane, QLD 4001, Australia

Dawei Song

Knowledge Media Institute

The Open University

Walton Hall, Milton Keynes, MK7 6AA United Kingdom

Ron Chi Wai Kwok

Department of Information Systems

City University of Hong Kong

Tat Chee Avenue, Kowloon, Hong Kong SAR

Abstract

Intelligent software agents are promising in improving the effectiveness of e-marketplaces for e-commerce. Although a large amount of research has been conducted to develop negotiation

protocols and mechanisms for e-marketplaces, existing negotiation mechanisms are weak in dealing with complex and dynamic negotiation spaces often found in e-commerce. This paper illustrates a novel knowledge discovery method and a probabilistic negotiation decision making mechanism to improve the performance of negotiation agents. Our preliminary experiments show that the probabilistic negotiation agents empowered by knowledge discovery mechanisms are more effective and efficient than the Pareto optimal negotiation agents in simulated e-marketplaces.

Key words: Knowledge Discovery, Bayesian Learning, Adaptive Negotiation Agents, e-Marketplaces.

1. Introduction

The number of transactions conducted over e-marketplaces has grown rapidly in recent years. In the context of Business-to-Business (B2B) e-commerce, e-marketplaces are no longer operated in isolation but function as a series of interacting markets along an electronic supply chain (eChain) [33]. It is argued that software agents can provide high level of intelligence and autonomy for enhancing the effectiveness of e-marketplaces [6]. Software agents are encapsulated computer systems situated in some environments such as the Internet and are capable of flexible, *autonomous* actions in that environment to meet their design objectives [39]. These agents can incorporate experiential knowledge of past transactions to streamline the effects of volatile demand and supply conditions across multiple e-marketplaces in the electronic supply chain. Negotiation refers to the process by which group of agents (human or software) communicate with one another in order to reach a mutually acceptable agreement on resource allocation (distribution) [21,36,37]. This paper focuses on the development of a novel knowledge discovery mechanism to enhance negotiation agents' decision making processes in B2B e-marketplaces.

1.1 The Problems

In typical B2B negotiation situations, a negotiator does not know the preferences of its

opponents because each party wants to protect their own business interests. Nevertheless, knowing the preferences of the opponents (e.g., the reservation prices) may help improve the efficiency of the negotiation processes since negotiation agents (human or software) can avoid wasting their time to explore the non-fruitful negotiation options. For cooperative agents, having partial knowledge about their opponents may even help improve the negotiation effectiveness because it becomes easier for the agents to identify the “win-win” outcomes from among the set of feasible solutions. Unfortunately, classical negotiation models [5,32,35,40] do not address the learning issue essential for real-world negotiations. Instead, these models often assume that the preferences (e.g., the utility functions) of the opponents are available as public information. Such an assumption turns out to be invalid for typical e-commerce negotiation situations. Even though agent-based negotiation systems have been developed, these systems still suffer from the problems of supporting only limited types of negotiation scenarios (e.g., bi-lateral negotiations, price only negotiations, availability of opponents’ payoff functions, or static negotiation spaces) [6,10,22,31]. One of the ways to alleviate the weakness of classical negotiation models and provide adequate support for real-world negotiations is to empower negotiation agents with a knowledge discovery mechanism so that they can continuously “mine” the preferences of the opponents based on the histories of negotiation dialogs among the participating agents.

1.2 Contributions

This paper illustrates the design and development of adaptive negotiation agents to enhance the degree of autonomy and the efficiency of e-marketplaces. In particular, the common weaknesses of existing negotiation systems are addressed by introducing a novel knowledge discovery method and a Bayesian learning mechanism to improve the learning autonomy and adaptation power of negotiation agents. These adaptive probabilistic negotiation agents can discover crucial negotiation knowledge such as the opponents’ changing preferences by mining the past negotiation histories and continuously monitoring the current negotiation dialogs with their opponents. Our preliminary experiments show that the probabilistic

negotiation agents empowered by our novel knowledge discovery mechanism outperform a negotiation mechanism which guarantees Pareto optimum. Our research work opens the door to the development of practical intelligent systems to enhance the effectiveness and efficiency of modern e-marketplaces.

1.3 Outline of the paper

The remainder of the paper is organized as follows. A comparative study of previous research work is reported in Section 2. An introduction to the basic negotiation mechanism which guarantees Pareto optimum is given in Section 3. Section 4 illustrates the computational details of the probabilistic negotiation decision making mechanism and the associated knowledge discovery method for adaptive negotiation agents. Section 5 describes the quantitative evaluation of the adaptive negotiation agents and reports our experimental results. Finally, we offer concluding remarks and describe future direction of our research work.

2. Related Work

Fuzzy logic has been applied to develop intelligent negotiation agents in e-Marketplace [6]. Nine pre-defined fuzzy rules are used to generate trade-off for quantitative issues and another nine fuzzy rules are used to generate concession for qualitative issues separately [6]. The proposed negotiation model is somewhat limited since it is developed from the perspective of the supplier agents only. The main weakness of the fuzzy negotiation system is that it is not adaptive; for instance, the system cannot learn and refine the pre-defined fuzzy rules automatically. The probabilistic negotiation agents proposed in this paper are adaptive since they are empowered by a knowledge discovery mechanism to continuously mine the preference information of their opponents.

Non-linear regression has been applied to estimate the specific parameters (e.g., lower/upper bounds of the zone of acceptance of an attribute, negotiation deadline, weight of individual tactic, etc.) of the time-dependent and the behavior-dependent negotiation tactics [4]. It is

assumed that agents' negotiation tactics are static and therefore it is possible to estimate these parameters based on the current negotiation dialogs. Instead of estimating the specific parameters of some pre-defined negotiation tactics, our proposed method adopts a non-parametric negotiation knowledge discovery approach where the opponents' negotiation tactics are not assumed static nor treated as public information. Our probabilistic negotiation agents are evaluated in multi-lateral dynamic negotiation scenarios.

Zeng and Sycara [40] have developed a sequential negotiation model called Bazaar. It was believed that an agent's belief about the opponent's true reservation price could be computed according to the posterior probability $Pr(H_i | o) = \frac{Pr(H_i)Pr(o|H_i)}{\sum_{k=1}^n Pr(o|H_k)Pr(H_k)}$, where $Pr(H_i)$ characterizes the probability distribution of the opponent's reservation prices and was assumed public information in the negotiation system. Moreover, domain knowledge in the form of conditional probabilities $Pr(o | H_i)$ describing the chance of receiving an offering price o given the opponent's true reservation price H_i was assumed available. Similar approach has also been applied to develop negotiation agents in the context of multi-agent co-ordination [5]. Nevertheless these approaches suffer from the problem of assuming the availability of the opponents' private information (e.g., the true reservation price). We illustrate an efficient data mining method of deriving the priori probabilities of offer acceptance without the assumption of the availability of the opponents' private information. Moreover, our proposed Bayesian learning mechanism is extended to deal with multiple negotiation issues.

Mining customers' transaction files to discover their shopping preferences has been conducted [16]. In particular, a Bayesian Belief network (BBN) is constructed to capture the dependency among the preferred shopping items based on the mutual information derived from among these items. The recommender system generates a recommendation set by referring to the customer's current transactional actions and the trained BBN representing the

shopping preferences of a particular customer. Our work is similar in the sense that we mine the negotiators' negotiation histories to discover their corresponding preferences. However, we use a computationally more efficient naive Bayesian approach since we would like the negotiation agents to conduct automated negotiations in real-time.

As a summary, there are variety of approaches of negotiation knowledge discovery such as case-based reasoning [3], fuzzy rules [6], time series approximation [24], Bayesian learning [5,40], Markov Chain Process [25], evolutionary learning [19], constraint satisfaction [38], etc. Generally speaking, these learning approaches can be classified into the broad categories of parametric [3,19,24] or non-parametric methods [5,25,40]. The negotiation knowledge discovery method illustrated in this paper is based on non-parametric approach since heterogeneous negotiation agents utilizing various tactics may be deployed to e-Marketplaces. Our non-parametric negotiation learning method is unique in the sense that it can support multi-party multi-issue negotiation situations and it has been tested under dynamic negotiation environment.

3. A Pareto Optimal Negotiation Model

A negotiation space $Neg = \langle P, A, D, U, T \rangle$ is a 5-tuple which consists of a finite set of negotiation parties (agents) P , a set of attributes (i.e., negotiation issues) A understood by all the parties $p \in P$, a set of attribute domains D for A , and a set of utility functions U with each function $U_p^o \in U$ for an agent $p \in P$. An attribute domain is denoted D_{a_i} where $D_{a_i} \in D$ and $a_i \in A$. A utility function pertaining to an agent p is defined by: $U_p^o : D_{a_1} \times D_{a_2} \times \dots \times D_{a_n} \mapsto [0,1]$ [17]. Each agent p has a deadline $t_p^d \in T$. It is assumed that information about P, A, D is provided by the facilitator agents in an e-marketplace. A *multi-lateral* negotiation situation can be modeled as many one-to-one *bi-lateral* negotiations where a negotiation agent p maintains a separate negotiation dialog with each opponent. In

a negotiation round, the agent will make an offer to each of its opponents in turn, and concentrate on the most favorable counter-offer from among the set of incoming offers evaluated according to its own payoff function U_p^o .

An *offer* $\vec{o} = \langle d_{a_1}, d_{a_2}, \dots, d_{a_n} \rangle$ is a n -tuple of attribute values (intervals) pertaining to a finite set of attributes $A = \{a_1, a_2, \dots, a_n\}$. Generally speaking, a finite set of candidate offers O_p acceptable to an agent p (i.e., satisfying its hard constraints) is constructed via the Cartesian product $D_{a_1} \times D_{a_2} \times \dots \times D_{a_n}$. As human agents tend to specify their preferences in terms of a range of values, a more general representation of an offer is a tuple of attribute value intervals such as $o_i = \langle 20-30, 1-2, 10-30, 100-500 \rangle$. The *valuations* of individual attributes and attribute values (intervals) are defined by the valuation functions $U_p^A : A \mapsto [0,1]$ and $U_p^{D_{a_i}} : D_{a_i} \mapsto [0,1]$ respectively, whereas U_p^A is an agent p 's *valuation function* for each attribute $a_i \in A$, and $U_p^{D_{a_i}}$ is an agent p 's valuation function for each attribute value $d_{a_i} \in D_{a_i}$. In addition, the valuations of attributes are assumed normalized, that is, $\sum_{a_i \in A} U_p^A(a_i) = 1$. One common way to quantify an agent's preference (i.e., the utility function U_p^o) for an offer o is by a linear aggregation of the *valuations* [2,13,17,26]: $U_p^o(o) = \sum_{a_i \in A, d_{a_i} \in o} U_p^A(a_i) \times U_p^{D_{a_i}}(d_{a_i})$, where d_{a_i} is the attribute value interval specified in an offer o .

If an agent's initial proposal is rejected by its opponent, it needs to propose an alternative offer with the least utility decrement (i.e., computing a concession). An agent will maintain a set O_p' which contains the offers it has proposed before (including the offer proposed in the current round). In a negotiation round, an alternative offer with a concession can be

determined based on $\exists_{o_{counter} \in \{O_p - O_p'\}} \forall_{o_x \in \{O_p - O_p'\}} : [o_x \preceq_p o_{counter}]$, where $o_x \preceq_p o_y$ denotes that an offer o_y is more preferable than another offer o_x . The preference relation \preceq_p is a total ordering induced by an agent p 's utility function U_p^o over the set of feasible offers O_p . The concession mechanism works by picking an offer from the top of the list ranked by $(\preceq_p, \{O_p - O_p'\})$ in each negotiation round.

The term o_{\equiv} represents agent p 's interpretation about the opponent's proposal o . Once o_{\equiv} is computed, acceptance of the incoming offer o can be determined with respect to p 's own preference (\preceq_p, O_p) . An offer $o_{\equiv} \in O_p$ is equivalent to o iff every attribute interval of o_{\equiv} intersects each corresponding attribute interval of o . The acceptance criteria for an incoming offer o (i.e., the equivalent o_{\equiv}) is defined by:

1. If $\forall_{o_x \in O_p} o_x \preceq_p o_{\equiv}$, an agent p should accept o since it produces the maximal payoff.
2. If $o_{\equiv} \in O_p'$ is true, an agent p should accept o because o_{\equiv} is one of proposals it makes before.

It is shown that if each participating agent $p \in P$ employs their preference ordering (\preceq_p, O_p) to compute concessions and uses the offer acceptability criteria described above to evaluate incoming offers, *Pareto optimal* [29] result is always found if it does exist in a negotiation space [2].

4. The Probabilistic Negotiation Agents

The development of the probabilistic negotiation mechanism for adaptive negotiation agents is driven by the basic intuition that rational negotiators strive for two possibly contradictory objectives [10,19]: (1) maximizing self payoffs, and (2) maximizing the chance of reaching an

agreement. The former can be computed according to a negotiator's private utility function as discussed in Section 3, and the latter can be estimated based on Bayesian learning [8]. The proposed adaptive negotiation agents can refer to the negotiation history files to discover the negotiation preferences of their opponents. Moreover, these agents can monitor the current negotiation dialog with their opponents to identify the possible preferential changes of their opponents.

4.1 Probabilistic Negotiation Decision Making

Our probabilistic negotiation agents' decision making mechanisms are underpinned by a ranking function; this function produces a ranked list of offers according to the potential of individual offers for maximizing self payoff and the chance of offer acceptance by the opponent. In particular, the preference relation \preceq_p of a probabilistic negotiation agent is a total ordering induced by the product of the agent's private utility function U_p^o and the probability function $Pr(accept | o)$ which characterizes the probability of acceptance of an offer o by the opponent. In other words, the feasible offers of an agent p are ranked in descending order according to:

$$Rank(o) = [U_p^o(o)]^\alpha \times [Pr(accept | o)]^{\frac{1}{1-\alpha}} \quad (1)$$

where $\alpha \in [0,1]$ is a trade-off factor for maximizing one's own payoff or maximizing the chance of the offer being accepted by the opponent. According to our current implementation, when $\alpha = 0$ is specified by the human negotiator, the fraction $\frac{1}{\alpha}$ will not be computed and a default value of zero will be returned; this results in instantiating a benevolent agent which only considers the opponent's benefits. On the other hand, if $\alpha = 1$ is specified, the fraction $\frac{1}{1-\alpha}$ will return zero instead of an undefined value; this results in instantiating a strictly self-interest agent. Moreover, a system wide default of 0.5 will be assumed if the α value is not provided by the human negotiator initially. It should be noted that the absolute

numerical value of $Rank(o)$ is not important, but the relative rank of an offer o .

A counter-offer with the least amount of concession (in terms of the least decrement of own payoff and the minimal reduction of offer acceptability) is selected from the top of the list $(\preceq_p, \{O_p - O_p'\})$ ranked by an agent p in each negotiation round. Once the counter-offer is determined, it will be added to the set O_p' . The revised O_p' forms the basis to evaluate the incoming offers. The probability of acceptance of an offer o can be computed according to Bayes theorem [8]:

$$Pr(c_j | o) = \frac{Pr(o | c_j) \times Pr(c_j)}{Pr(o)} \quad (2)$$

where $c_j \in \{accept, reject\}$ and j is the index of a particular class. If the naive assumption of feature (i.e., negotiation issue) independency is made, the prior probability $Pr(o | c_j)$ can be approximated by [23,27]:

$$Pr(o | c_j) = \prod_{i=1}^{|A|} Pr(d_{a_i} | c_j) \quad (3)$$

where d_{a_i} is one of the attribute values of an offer o . By the addition rule of probability theory, $Pr(o) = \sum_{j=1}^n Pr(o | c_j) \times Pr(c_j)$ is held. Therefore, the probability of acceptance of an offer o by the opponent can be estimated according to:

$$\begin{aligned} Pr(accept | o) = & (Pr(accept) \times \prod_{i=1}^{|A|} Pr(d_{a_i} | accept)) \div \\ & [Pr(accept) \times \prod_{i=1}^{|A|} Pr(d_{a_i} | accept) + \\ & Pr(reject) \times \prod_{i=1}^{|A|} Pr(d_{a_i} | reject)] \end{aligned} \quad (4)$$

It should be noted that if only a partial counter-offer (i.e., some attributes are missing in an offer) is evaluated, the corresponding terms such as $Pr(d_{a_i} | accept)$ and $Pr(d_{a_i} | reject)$ are treated as 1 because these negotiation issues are considered not relevant by an agent. As

a result, the probability of offer acceptance is determined by other attribute values. Currently, there are two operating modes of our probabilistic negotiation agents, namely adaptive and non-adaptive. For the non-adaptive mode, the probability negotiation agents estimate the opponents' preferences based on the past negotiation histories only. They operate based on the negotiation mechanism described in Section 3 except that the offer ranking is established according to Eq.(1) instead of based on an agent's own utility function. After a negotiation session begins, the preferences of an agent and its opponents are assumed unchanged.

On the other hand, for the adaptive probabilistic agents, the probability function $Pr(accept | o)$ is revised in each negotiation round based on the most recent negotiation dialog. Therefore, the adaptive probabilistic negotiation agents are sensitive to the opponents' recent preferential changes. After updating the priori probabilities based on the current negotiation dialog, the set of feasible offers O_p for the agent $p \in P$ will be re-ranked again according to Eq.(1). As a result, more sensible negotiation decision making can be conducted from time to time according to the most recent preferences of agent p and its opponents. As the time dimension is always an important issue for practical negotiations [21], our probabilistic negotiation agents are extended to take into account the time pressure:

$$Rank(o) = [U_p^o(o)]^{\frac{1}{\alpha \times TP(t)}} \times [Pr(accept | o)]^{\frac{1}{(1-\alpha) \times TP(t)}} \quad (5)$$

The term $TP(t)$ represents the time pressure function. The basic intuition is that when the negotiation deadline is approaching, an agent is more likely to concede in order to make a deal [14,30]. However, different agents may have different attitudes towards deadlines. An agent may be eager to reach a deal and so it will concede quickly (*Conceder agent*). On the other hand, an agent may not give ground easily during negotiation (*Boulware agent*) [28].

Therefore, a time pressure function $TP(t) = 1 - (\frac{t}{t_d})^{\frac{1}{e_p}}$ is developed to approximate a wide spectrum of agents' concession attitude. Our TP function is similar to the negotiation

decision function referred to in the literature [9,11]. The term t_p^d indicates the deadline for an agent p , and e_p is used to model the “concession attitude” of the agent p .

4.2 Mining Negotiation Knowledge

Data mining refers to the non-trivial process of identifying valid, novel, potentially useful, and ultimately understandable patterns in data [12]. In the context of knowledge discovery for automated negotiations, the novel patterns are the negotiation preferences (i.e., the frequently requested issues and their values). This kind of patterns is ultimately understandable and potentially useful because they can be applied to improve both negotiation effectiveness (e.g., joint payoffs) and negotiation efficiency (e.g., reducing the amount of time to reach agreements). In association rule mining, the measures of rule support and rule confidence are used to evaluate the quality of the association rules extracted from frequent item-sets [1]. In fact, rule support and rule confidence correspond to the joint probability and the conditional probability of the appearance of items (e.g., consumer products) in transactions. Our approach of discovering the preferences of negotiators is also based on computing the priori probabilities of the frequently requested items (negotiation options) appearing in some offers. The prior probabilities such as $Pr(accept)$, $Pr(reject)$, $Pr(d_{a_i} | accept)$, and $Pr(d_{a_i} | reject)$ can be estimated based on the negotiation histories. Figure 1 depicts a segment of a negotiation history file.

[INSERT FIGURE 1 HERE]

The basic assumption of our negotiation knowledge discovery method is that each counter-offer from the opponent is considered an acceptable offer (i.e., a positive training example). Moreover, if an agent proposes an offer and it is rejected by the opponent, it is treated as a negative training example. As agents’ preferences may change in real-world negotiation situations, the most recently archived negotiation sessions are more useful than the sessions archived long time ago in terms of estimating the opponent’s current preferences. Moreover, a negotiation agent will maintain a separate history for the negotiation processes it

conducted with each of its partners. A negotiation session refers to a particular negotiation process. The negotiation process ends when an agreement could be reached or all the parties decide to quit. For each negotiation process (session), a negotiation agent will make series of offers (i.e., entries). Based on the above assumptions, the training examples generated from the past negotiation sessions and the current negotiation dialog should be weighted. The weight factor w_i^S is computed and assigned to a negotiation session i according to a linear function:

$$w_i^S = w_{max} - step \times \frac{w_{max} - w_{min}}{|session| - 1} \quad (6)$$

where w_i^S is the highest weight assigned to a particular negotiation session i ; the terms $w_{max} > 0$ and $w_{min} > 0$ represent the maximal and the minimal weights assigned to valuate all the negotiation sessions. The term $|session|$ is the total number of archived negotiation sessions including the current negotiation session for knowledge discovery purpose. The term $step = \langle 0, 1, \dots, |session| - 1 \rangle$ represents the sequence of negotiation sessions. For example, the step value of the most current negotiation session is 0, and the second most current session is 1, so on so forth.

In addition, the weight of each offer (or counter-offer) within a negotiation session i varies. For instance, a counter-offer proposed by the opponent at the earlier stage is more preferable (for the opponent) than the one proposed at a later stage. Therefore, each entry in a negotiation session i is also weighted in chronological order. The second weight factor w_{ij}^E for the j th negotiation entry (i.e., an event) in the i th archived negotiation session is computed according to:

$$w_{ij}^E = w_i^S - (j-1) \times \frac{w_i^S - w_{i+1}^S}{|E|} \quad (7)$$

where E is the total number of entries of an archived negotiation session i . w_i^S and w_{i+1}^S

are the highest session weights assigned to the i session and the session immediately preceding it respectively. For the oldest session (i.e., $(i+1) > |session|$) in a negotiation history file, the value of w_{i+1}^S is assumed zero.

[INSERT TABLE 1 HERE]

Table 1 shows an example of the sample space which consists of 3 past negotiation sessions and 1 current negotiation session. The maximal weight $w_{max} = 500$ and the minimal weight $w_{min} = 200$ are set. The entry depicted at the bottom of Table 1 represents the current negotiation session between an agent and its opponent. The weight of the second most current negotiation session is computed according to Eq.(6), that is $w_2^S = 500 - 1 \times \frac{500-200}{4-1} = 400$. In addition, the weight of the second negotiation entry in this session is computed according to Eq.(7), that is $w_{22}^E = 400 - (2-1) \times \frac{400-300}{4} = 375$. In fact, the weights can be interpreted as the additional sample points attached to each event (i.e., a training example). According to Table 1, there are 3650 sample points in the sample space, and $Pr(accept) = \frac{1500}{3650} = 0.41$ is estimated. Similarly, $Pr(reject) = \frac{2150}{3650} = 0.59$, $Pr(price = 25 - 30 | accept) = \frac{675}{1500} = 0.45$, $Pr(qty = 50 - 50 | accept) = \frac{825}{1500} = 0.55$.

5. The Experiments

5.1 General Procedure

The simulated e-marketplaces were characterized by multi-lateral negotiations among some buyer agents ($B1, \dots, Bn$) and some seller agents ($S1, \dots, Sn$). These agents negotiated over some virtual services or products described by five attributes (i.e., $|A|=5$) with each attribute domain containing five discrete values $D_{a_i} = \{1, 2, 3, 4, 5\}$. For each agent p , the size of the feasible offer set is: $|O_p| = 5^5 = 3,125$. The valuation of an attribute or a discrete attribute value fell in the unit interval of $(0, 1]$. For each negotiation case, an agreement zone

always existed since the difference between the buyers and the sellers only lay on their valuations against the same set of negotiation issues (e.g., attributes and attribute values). The simulated e-marketplaces were symmetric where the same number of buyers and sellers participated.

At the beginning of every negotiation round, each agent would invoke its own decision making mechanism to generate an offer for that round. The order of deliberation of out-going offers among the agents was randomly chosen by the facilitator agent. At the message exchange phase, each agent sent the offer messages to all the opponents (e.g., $S1 \rightarrow B1$, $S1 \rightarrow B2$, $S1 \rightarrow B3$, etc. for the seller $S1$). After the message exchange phase, the facilitator agent randomly selected a sequence of agents such as $\langle B2, B1, S1, S2, \dots, S_n \rangle$ for incoming offer evaluation. For instance, with reference to the above sequence, agent $B2$ would evaluate its incoming offers first, then agent $B1$ would evaluate its incoming offers, etc. If agreements could be made, an agent always selected the best deal (evaluated according to its private utility function). If there was a tie, an opponent would be randomly selected by an agent. Once an agreement was made between a pair, they would be removed from the e-marketplace immediately by the facilitator agent, and the remaining agents would continue their negotiations until either agreement was made or the negotiation deadline was due. Our e-marketplaces were instantiated on a PC with a Pentium-4 2.2GHz single processor and 1 GB main memory. To avoid the communication overheads, all the experiments were conducted under our Intranet environment. All the agents were developed using Java SDK 1.5.0.

5.2 Evaluation Measures

Both the effectiveness (in terms of average joint payoff) and the efficiency (in terms of average number of negotiation rounds) of the negotiation processes were evaluated. We adopt the relative measure of “negotiation rounds” to assess the negotiation time involved in a negotiation process (and indirectly measuring the computational/communication costs) so that

it becomes easier to compare our results with others which may be conducted in different computational environments. Moreover, average weighted Euclidean distance [8] was also used to measure how far away the solutions obtained by our probabilistic agents from the Pareto optimum:

$$AvgDist = \frac{\sum_{j=1}^{|PS|} \sum_{i=1}^{|PA|} dist(\vec{o}_i, \vec{o}_j)}{|PA| \times |PS|} \quad (8)$$

$$dist(\vec{o}_x, \vec{o}_y) = \sqrt{\sum_{i=1}^{|A|} w_i (d_{a_i}^x - d_{a_i}^y)^2} \quad (9)$$

where $PA \subseteq P$ is the set of agents reaching an agreement in an e-marketplace, and PS is the set of Pareto optimal solutions. As each agent has its own preference w_i for an attribute a_i , the average distance is computed among the agents PA by Eq.(8). Since the Pareto optimum set PS may contain more than one optimal solution, the mean distance between the agents' solution and every Pareto optimal solution is computed. The weight factor $w_i = U_p^A(a_i)$ in Eq.(9) is an agent's valuation for a particular attribute $a_i \in A$. An offer vector \vec{o}_x contains an attribute value $d_{a_i}^x$ along the i th dimension (issue) in a negotiation space. If an attribute interval instead of a single value is specified for an offer, the mid-point of an attribute interval is first computed.

5.3 Experiment One

Hypothesis One: The probabilistic negotiation agents empowered by the knowledge discovery mechanisms are more efficient than the Pareto optimal agents which are not equipped with the knowledge discovery mechanisms.

5.3.1 The Experimental Procedures

The first experiment aimed at developing a basic test to see if the proposed negotiation knowledge mining method could improve the negotiation processes. Two buyer agents and two seller agents were involved (i.e., $|P| = 4$) in this experiment. There were six negotiation

groups which are characterized by various levels of conflict among the buyers and the sellers. Each group contained ten negotiation cases (i.e., totally $6 \times 10 = 60$ simulated e-marketplaces). For the first negotiation group, buyers and sellers had exactly the same utility functions (i.e., no conflict). Two utility functions are the same if both the valuation of the attributes and the valuation of the corresponding attribute values are the same. Table 2 shows an example of the utility functions for a buyer agent p_1 and a seller agent p_2 used in this experiment.

[INSERT TABLE 2 HERE]

For each succeeding negotiation group, buyers and sellers were characterized by having common weighting from one (small conflict group) to five attributes (highest conflict group) respectively [26]. For these negotiation groups, opposing valuations of the attribute values were created between the buying side and the selling side. In this experiment, no negotiation deadline was imposed in the e-marketplaces. The control group consisted of the Pareto optimal negotiation agents developed according to the negotiation mechanism described in Section 3. These agents could find Pareto optimal solutions when time constraint was not present. After running a simulated e-marketplace, the average joint-payoffs and the average negotiation time were recorded.

The experimental group comprised of the same number of non-adaptive (i.e., the priori probabilities about the opponents' preferences were not updated) probabilistic negotiation agents as defined in Section 4. The same set of negotiation cases attempted by the Pareto optimal agents were applied to the probabilistic agents. The negotiation histories of the Pareto optimal agents were made available to the probabilistic agents as the training set. As a result, each probabilistic agent had some knowledge about its opponents before the negotiation process began. In particular, only the first 60% of the entries captured in a negotiation session were used to train the probabilistic agents. We employed a heuristic $w_{min} \geq |E|_{MAX}$ and $w_{max} \geq |session| \times w_{min}$ to derive various combinations of w_{max} and w_{min} , whereas

$|E|_{MAX}$ is the number of entries of the largest archived session, and $|session|$ is the number of sessions archived in the negotiation history file. Based on the empirical testing for typical negotiation scenarios, we found that the parameters $w_{max} = 20,000$, $w_{min} = 2,000$, $\alpha = 0.6$ produced good performance and so they were adopted in this experiment. After a negotiation process began, a probabilistic agent could estimate the posteriori probability $Pr(accept|o)$ for each of its opponent based on the training data. In this experiment, the preference of each agent remained static.

5.3.2 The Experimental Results

According to the experimental results depicted in Table 3, the Pareto optimal agents achieved an overall average joint utility of 2.31 by using 938.25 negotiation rounds on average. On the other hand, the probabilistic agents achieved an overall average joint utility of 2.23 in 724.75 negotiation rounds on average. There was a $\frac{2.31-2.23}{2.31} \times 100 = 3.2\%$ decrement of the overall average joint utility when the probabilistic agents were engaged in the same negotiation situations as the Pareto optimal agents. The overall average distance from the solutions found by the probabilistic agents to the Pareto optimum is 0.16 which is considered a small distance. However, the improvement in terms of reduced average negotiation time of the probabilistic agents was $\frac{938.25-724.75}{938.25} \times 100 = 22.8\%$. Except the first negotiation group, the probabilistic agents consistently consumed less negotiation time than that of the Pareto optimal agents. For each test case in the first negotiation group, both buyers and sellers had exactly the same utility function. Therefore, an agreement could always be found in the first negotiation round. According to paired one tail t-test, the average negotiation time consumed by probabilistic agents is significantly less than that of the Pareto optimal agents, $t(5) = -3.48, p < .01$. Therefore, we conclude that the probabilistic negotiation agents empowered by knowledge discovery mechanisms can identify negotiation solutions faster than the Pareto optimal agents do. Hypothesis One is supported according to our experiment.

[INSERT TABLE 3 HERE]

[INSERT FIGURE 2 HERE]

5.3.3 *The Impact of the Trade-off Factor*

The impact of the trade-off factor α on the performance of our probabilistic negotiation agents was evaluated. In particular, we would like to observe how the various levels of α affect the quality of the solutions (e.g., average distance from the Pareto optimum) and the time required to search for those solutions. Figure 2 plots the overall average distance between the solutions found by the probabilistic agents and that produced by the Pareto optimal agents over the six negotiation groups listed in Table 3. It should be noted that the overall average distances plotted in Figure 2 was scaled up by a factor of one thousand. It is shown that a lower rate of decrement of the overall average distance occurs beyond $\alpha = 0.6$, and at the same time a higher rate of increment of the overall average negotiation time (e.g., a larger angle of the slop) occurs. Therefore, we set the trade-off factor to $\alpha = 0.6$ to strive for a better balance between the quality of the negotiation solutions and the time required to identify those solutions. According to our testing, the average joint payoffs of the agents do not vary with respect to the choices of different w_{max} and w_{min} as long as our heuristic of how to estimate these parameters was followed.

5.3.4 *The Impact of the Negotiation History*

In addition, we examined the impact of the availability of various amount of negotiation history data (i.e., the training set) on the effectiveness of the probabilistic negotiation agents. The negotiation histories were obtained by repeatedly invoking the Pareto optimal agents to attempt the 60 negotiation cases we developed before. Each negotiation history file contained certain number of recorded sessions (i.e., negotiation processes) and each session contained

certain number of entries (i.e., offers and counter-offers). In particular, we made the first x percentage of the negotiation entries in a session and the first y percentage of negotiation sessions in a history file available to train the probabilistic agents in each run. After the training process, the probabilistic agents would start to negotiate as before.

Figure 3 highlights the overall average joint payoff achieved by the probabilistic negotiation agents when various number of entries and sessions are used to train them. It is shown that using more than 60% of the top entries and more than 2 negotiation sessions to train the probabilistic agents cannot improve the maximal average joint utility. In fact, employing a large number of sessions (e.g., 10 sessions) and all the entries (100%) of a negotiation session to train a probabilistic agent may even lead to slightly degraded performance because the final offers do not represent the actual preference of the opponent due to the concession making process. According to our empirical testing, it only took 6.1 seconds to train an agent (i.e., computing all the priori probabilities) with a negotiation history file containing 10 sessions and each session containing 500 entries on average. This shows a positive sign for the computational efficiency of our negotiation knowledge discovery method.

[INSERT FIGURE 3 HERE]

5.3.5 Discussion

As a summary, this experiment confirms that the probabilistic negotiation agents empowered by the knowledge discovery mechanism can make use of the knowledge about their opponents to find agreements faster. With better knowledge about their opponents, the probabilistic negotiation agents can by-pass some of the non-fruitful offers (e.g., chance of acceptance is low) from the set of feasible offers. Even though the probabilistic agents are not fully self-interested, they can achieve near optimal joint payoffs. This represents a win-win negotiation strategy. Such a strategy is desirable for negotiations in B2B e-commerce because it helps maintain long-term relationships among business partners.

5.4 Experiment Two

Hypothesis Two: Under time pressure, the probabilistic negotiation agents outperform the Pareto optimal agents in terms of negotiation effectiveness.

5.4.1 The Experimental Procedures

The second experiment tries to evaluate the effectiveness of our probabilistic agents under realistic negotiation condition such as the presence of time pressure. In addition, we would like to test the agents' time adjustment mechanisms defined according to Eq.(5). The same set of negotiation cases employed in experiment one was re-used with a negotiation deadline of 500 rounds. This experiment was still based on a control group vs. experimental group design. The first simulation run involved the Pareto optimal agents, and then the non-adaptive probabilistic agents participated in the second simulation run. The third simulation run involved the non-adaptive probabilistic agents with the time adjustment mechanisms Eq.(5) activated. The concession attitude $e_p = 0.4$ was set for all the time sensitive probabilistic agents. If an agent could not find a deal before the negotiation deadline, its payoff would be zero.

[INSERT TABLE 4 HERE]

5.4.2 The Experimental Results

The average joint-payoffs obtained by different types of agents from six negotiation groups are tabulated in Table 4. The overall average joint payoffs achieved by the Pareto optimal agents, the probabilistic agents, and the time sensitive probabilistic agents are 1.42, 1.58, and 2.09 respectively. Except the first negotiation group where agents could always find the best agreements in the first negotiation round, the probabilistic agents consistently performed better than the Pareto optimal agents, and the time sensitive probabilistic agents also performed better than their non time sensitive counterparts. According to paired one tail t-test, the average joint payoffs achieved by probabilistic agents are significantly higher than that of

the Pareto optimal agents (e.g., $t(5) = 4.72, p < .01$ for the non time sensitive probabilistic agents and $t(5) = 3.13, p = .01$ for the time sensitive probabilistic agents). Therefore, we can conclude that our proposed probabilistic agents are more effective than the Pareto optimal agents under time pressure. Hypothesis Two is supported according to our experiment. The reason for such a difference is that the Pareto optimal agents could not be able to find solutions before the deadline in many cases where negotiation conflicts existed. On the other hand, the probabilistic agents were able to carry out the search faster (e.g., by ignoring some less promising deals). As a result, they were able to seal some deals even though a tough deadline was imposed.

5.4.3 The Difference of Agents' Concession Behavior

Figure 4 shows the difference of the concession making processes conducted by a Pareto optimal agent (PO), a probabilistic agent (PR), and a time sensitive probabilistic agent (PRT) respectively. There was a significant performance boost of the probabilistic agents who were empowered by the time adjustment mechanism because these agents were sensitive to the negotiation deadlines. When the deadline was approaching, these agents tended to propose offers which were more likely to be accepted by their opponents (e.g., the α value drop to a very low value). As a consequence, the time sensitive probabilistic agents could find agreements for all the negotiation cases in this experiment. The Y axis in Figure 4 represents the potential utility value brought to an agent if the corresponding offer is really accepted by the opponent. The comparison is based on one of the negotiation cases from negotiation group 4, and the potential payoffs of the offers are computed from the perspective of the sellers. It is not difficult to find that the concession making behavior of the time sensitive probabilistic agent is different from the other two agents. For instance, there was a bigger drop of the utility values of the agent's offers when the deadline was approaching. In this case, the time sensitive probabilistic agent found an agreement at the 486th negotiation round, whereas the other two agents could not find solution before the deadline of the 500th negotiation rounds. According to paired one tail t-test, the average payoffs of the time sensitive probabilistic

agents is significantly higher than that of their non time sensitive counterparts, $t(5) = 2.59, p = .02$. As can be seen, the proposed time adjustment mechanism for probabilistic agents is effective since it can improve negotiation outcomes in general.

[INSERT FIGURE 4 HERE]

5.4.4 *The Impact of the Agents' Concession Attitude*

We further tested the concession attitude of the time sensitive probabilistic negotiation agents by varying the parameter value $e_p = 0.1$ (extreme Boulware agents) and $e_p = 10$ (Conceder agents) while keeping the other experimental conditions unchanged. For the extreme Boulware agents, the overall average joint payoff achieved is 2.06. For the Conceder agents, the overall average joint payoff achieved is 1.95. The extreme Boulware agents actually failed to reach an agreement in one case with high conflict and so their performance was not as good as the little Boulware agents. On the other hand, the conceder agents conceded too quickly even for the neutral negotiation situations, and therefore their performance was not as good as the little Boulware agents either.

5.5 *Experiment Three*

Hypothesis Three: Under dynamic negotiation environment (e.g., the presence of preferential changes of the negotiators), the adaptive probabilistic negotiation agents can achieve near Pareto optimal negotiation results.

5.5.1 *The Experimental Procedures*

Under realistic negotiation situations, the preferences of negotiation agents may change over time. This experiment tries to test if the adaptive probabilistic negotiation agents (i.e., their priori probabilities were updated after every negotiation round) can achieve good negotiation outcomes given the preferential changes of themselves and their opponents. At the beginning of the simulations, we employed the same set of negotiation cases used in experiment one. However, the utility functions of the agents would be modified n times after the negotiation

processes started. As a result, the utility functions of these agents may not be the same as that examined in experiment one when the negotiation processes completed. The final outcomes in terms of average joint payoffs and the average distances from the Pareto optimum would be recorded. The final negotiation outcomes were computed according to the last modified utility functions of the agents. For instance, after 100 rounds of negotiations, the valuation values of m attributes pertaining to each agent would be randomly selected and modified. After another 100 rounds of negotiations, the agents' valuation functions would be changed again (i.e., $n = 2$). To ensure the required number of preferential changes could be injected into each agent, these agents were forced not to accept a deal until the last change was injected. No negotiation deadline was imposed in this experiment.

To evaluate the adaptiveness of the probabilistic agents, we invoked the adaptive evolutionary negotiation agents [20] under the same conditions (e.g., the same utility function and the same preferential changes of the agents). The evolutionary agents were developed based on a genetic algorithm (a heuristic search approach) and they were not equipped with a knowledge discovery mechanism. The negotiation agents discussed in this paper were empowered by a knowledge discovery mechanism underpinned by Bayesian learning. For this experiment, the parameters $n = 2$ and $m = 2$ were used, whereas n and m stand for the frequency of changes and the number of attributes modified respectively. It should be noted that the first negotiation group was a reference group where no preferential changes was injected to the agents. As the preferences of the buyers and the sellers were the same for this negotiation group, an agreement was always reached in the first negotiation round in each case.

5.5.2 The Experimental Results

The comparison between the performance of the evolutionary agents and that of the adaptive probabilistic agents is tabulated in Table 5. By ignoring the reference group (negotiation group 1), the overall average distance of the solutions found by the adaptive probabilistic

agents from the Pareto optimum is $\frac{0.14+0.22+0.19+0.23+0.25}{5} = 0.21$, which is close to that (i.e., 0.19) achieved by the non-adaptive probabilistic agents in experiment one. This demonstrates that the average performance of our adaptive probabilistic agents can be maintained even though they operate under a more challenging dynamic negotiation environment.

[INSERT TABLE 5 HERE]

Based on our simulations, the difference between the average joint payoff of the evolutionary agents and that of the adaptive probabilistic agents is only marginal (by paired one tail t-test, $t(5) = 1.17$, $p = .15$). On the other hand, significant difference of the average weighted Euclidean distances between the evolutionary agents and the probabilistic agents was not found (by paired one tail t-test, $t(5) = -0.54$, $p = .31$). Therefore, we conclude that the adaptive probabilistic agents are able to adapt to the dynamic negotiation environment, and the performance of these agents is comparable to that achieved by the adaptive evolutionary agents whose effectiveness was tested in a previous study [20]. The adaptive probabilistic agents can produce negotiation outcomes close to the Pareto optimum (e.g., 0.21 point away from the optimum). In general, Hypothesis Three is supported according to our study.

Figure 5 plots the overall average distances from the Pareto optimum given the various values of n and m . The overall average distances shown in Figure 5 excluded the negotiation group one since no preferential change was injected to the agents in this group. It is not difficult to observe that the agents can adapt to the preferential changes presented in the negotiation environment. As a result, the overall average distances from the Pareto optimum do not vary much given more frequent changes and greater extent of changes. However, if the changes occurred more frequently, the negotiation efficiency will be affected. Figure 6 shows that more frequent preferential changes of the agents will generally increase the overall average negotiation time. The reason is that the agents need to take time to learn the opponents' new preferences and adapt to these preferences. Nevertheless, the probabilistic negotiation agents seem robust enough in responding to the frequent changes because the

negotiation time is only increased linearly with respect to the number of preferential changes as shown in Figure 6.

[INSERT FIGURE 5 HERE]

[INSERT FIGURE 6 HERE]

6. Conclusions and Future Work

Intelligent software agents are promising for supporting business negotiations in e-marketplaces. Since real-world negotiation spaces are complex and dynamic, it is desirable to empower negotiation agents with effective knowledge discovery mechanisms so that these agents can automatically uncover essential negotiation knowledge to improve negotiation outcomes. A novel knowledge discovery method and the corresponding probabilistic negotiation decision making mechanism are developed for adaptive negotiation agents. These agents can continuously learn the preferences of their opponents based on the negotiation dialogs recorded in history files. Our preliminary experiments show that the probabilistic negotiation agents empowered by knowledge discovery mechanisms can make a better balance between maximizing self-payoff and improving offer acceptability, and therefore they are more effective and efficient than the Pareto optimal agents under realistic negotiation conditions. Our research opens the door to the development of intelligent software tools to enhance the autonomy and effectiveness of e-marketplaces. As naive Bayesian learning is adopted in our negotiation knowledge discovery framework, dependencies among negotiation issues cannot be taken into account. Future research will explore Bayesian belief network to model the dependency of negotiation issues in complex negotiation spaces. Since our current probabilistic negotiation decision making mechanism only takes into account the opponent agents' concession patterns, an extended decision making mechanism which also considers the opponents' reputation will be examined in the future.

Acknowledgments

The work reported in this paper has been funded in part by the Australian Research Council

(ARC) Discovery grant, grant number DP0556455. This work is also funded in part by the UK's Engineering and Physical Sciences Research Council (EPSRC), grant number EP/E002145/1. The authors would like to thank three anonymous reviewers for their insightful comments and suggestions that help a lot in improving the quality of the paper.

References

- [1] R. Agrawal and R. Srikant. Fast algorithms for mining association rules in large databases. In Jorge B. Bocca, Matthias Jarke, and Carlo Zaniolo, editors, *VLDB'94, Proceedings of 20th International Conference on Very Large Data Bases*, pages 487–499, Santiago de Chile, Chile, September 12–15 1994. Morgan Kaufmann Publishers.
- [2] Mihai Barbuceanu and Wai-Kau Lo. Multi-attribute utility theoretic negotiation for electronic commerce. In Frank Dignum and Ulises Cortés, editors, *Agent-Mediated Electronic Commerce III — Current Issues in Agent Based Electronic Commerce Systems*, number 2003 in *Lecture Notes in Artificial Intelligence*, pages 15–30. Springer-Verlag: Heidelberg, Germany, 2001.
- [3] Jakub Brzostowski and Ryszard Kowalczyk. On possibilistic case-based reasoning for selecting partners in multi-agent negotiation. In Geoffrey I. Webb and Xinghuo Yu, editors, *Proceedings of the 17th Australian Joint Conference on Artificial Intelligence*, volume 3339 of *Lecture Notes in Computer Science*, pages 694–705, Cairns, Australia, December 4–6 2004. Springer.
- [4] Jakub Brzostowski and Ryszard Kowalczyk. Adaptive negotiation with on-line prediction of opponent behaviour in agent-based negotiations. In *Proceedings of the 2006 IEEE/WIC International Conference on Intelligent Agent Technology*, pages 263–269. IEEE Computer Society, 2006.
- [5] H. H. Bui, S. Venkatesh, and D. Kieronska. Learning other agents' preferences in multiagent negotiation using the bayesian classifier. *International Journal of Cooperative Information Systems*, 8(4):275–293, 1999.
- [6] Chi-Bin Cheng, Chu-Chai Henry Chan, and Kun-Cheng Lin. Intelligent agents for e-marketplace: Negotiation with issue trade-offs by fuzzy inference systems. *Decision Support Systems*, 42(2):626–638, 2006.
- [7] D. Chiu, S. Cheung, P. Hung, S. Chiu, and A. Chung. Developing e-negotiation support with a meta-modeling approach in a web services environment. *Decision Support Systems*, 40(1):51–69, 2005.
- [8] R. Duda and P. Hart. *Pattern Classification and Scene Analysis*. John Wiley & Sons, New York, New York, 1973.
- [9] P. Faratin, C. Sierra, and N. R. Jennings. Negotiation decision functions for autonomous agents. *Journal of Robotics and Autonomous Systems*, 24(3-4):159–182, 1998.

- [10] P. Faratin, C. Sierra, and N. R. Jennings. Using similarity criteria to make issue trade-offs in automated negotiations. *Artificial Intelligence*, 142(2):205–237, 2002.
- [11] S. Fatima, M. Wooldridge, and N. R. Jennings. An agenda based framework for multi-issues negotiation. *Artificial Intelligence*, 152(1):1–45, 2004.
- [12] U.M. Fayyad, G. Piatetsky-Shapiro, and P. Smyth. From Data Mining to Knowledge Discovery: An Overview. In U.M. Fayyad, G. Piatetsky-Shapiro, P. Smyth, and R. Uthurusamy, editors, *Advances in Knowledge Discovery and Data Mining*, pages 1–34. The MIT Press, 1996.
- [13] E. Gerding, D. van Bragt, and H. La Poutré. Multi-issue negotiation processes by evolutionary simulation, validation and social extensions. *Computational Economics*, 22:39–63, 2003.
- [14] F. Harinck and C. De Dreu. Negotiating interests or values and reaching integrative agreements: the importance of time pressure and temporary impasses. *European Journal of Social Psychology*, 34(5):595–611, 2004.
- [15] M. He, N. R. Jennings, and H. Leung. On agent-mediated electronic commerce. *IEEE Trans. on Knowledge and Data Engineering*, 15(4):985–1003, 2003.
- [16] Junzhong Ji, Chunnian Liu, Jing Yan, and Ning Zhong. Bayesian Networks Structure Learning and Its Application to Personalized Recommendation in B2C Portal. In *Proceedings of the 3rd IEEE/WIC International Conference on Web Intelligence*, pages 179–183, Beijing, China, September 20–24 2004. IEEE Computer Society.
- [17] R. Keeney and H. Raiffa. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. Cambridge University Press, Cambridge, UK, 1993.
- [18] J. Kim and A. Segev. A web services-enabled marketplace architecture for negotiation process management. *Decision Support Systems*, 40(1):71–87, 2005.
- [19] R. Krovi, A. Graesser, and W. Pracht. Agent behaviors in virtual negotiation environments. *IEEE Transactions on Systems, Man, and Cybernetics*, 29(1):15–25, 1999.
- [20] R.Y.K. Lau, M. Tang, O. Wong, S. Milliner, and Y. Chen. An Evolutionary Learning Approach for Adaptive Negotiation Agents. *International Journal of Intelligent Systems*, 21(1):41–72, 2006.
- [21] A. R. Lomuscio and N. R. Jennings. A classification scheme for negotiation in electronic commerce. *Journal of Group Decision and Negotiation*, 12(1):31–56, 2003.
- [22] P. Maes, R. Guttman, and A. Moukas. Agents that buy and sell. *Communications of the ACM*, 42(3):81–91, March 1999.
- [23] T. Mitchell. *Machine Learning*. McGraw Hill, New York, New York, 1996.
- [24] Wilson Wai Ho Mok and R. P. Sundarraj. Learning algorithms for single-instance electronic negotiations using the time-dependent behavioral tactic. *ACM Transactions on Internet Technology*, 5(1):195–230, 2005.

- [25] Vidya Narayanan and Nicholas R. Jennings. An adaptive bilateral negotiation model for E-commerce settings. In *Proceedings of the Seventh IEEE International Conference on E-Commerce Technology (CEC'05)*, pages 34–41, 2005.
- [26] Jim R. Oliver. A machine learning approach to automated negotiation and prospects for electronic commerce. *Journal of Management Information Systems*, 13(3):82–112, 1996.
- [27] M.J. Pazzani and D. Billsus. Learning and revising user profiles: The identification of interesting web sites. *Machine Learning*, 27(3):313–331, 1997.
- [28] H. Raiffa. *The Art and Science of Negotiation*. Harvard University Press, 1982.
- [29] J. Rosenschein and G. Zlotkin. Task oriented domains. In *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*, pages 29–52. MIT Press, Cambridge, Massachusetts, 1994.
- [30] A. Roth, J. Murnighan, and F. Schoumaker. The deadline effect in bargaining: Some experimental evidence. *American Economic Review*, 78(4):806–823, 1988.
- [31] Bonnie Rubenstein-Montano and Ross A. Malaga. A Weighted Sum Genetic Algorithm to Support Multiple-Party Multi-Objective Negotiations. *IEEE Transactions on Evolutionary Computation*, 6(4):366–377, August 2002.
- [32] A. Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):97–109, 1982.
- [33] Rahul Singh, A. F. Salam, and Lakshmi Iyer. Agents in e-supply chains. *Communications of the ACM*, 48(6):108–115, 2005.
- [34] Paul Tremblett. Java and UDDI registries. *Dr. Dobb's Journal of Software Tools*, 27(9):34, 37–40, September 2002.
- [35] J. von Neumann and O. Morgenstern. *The Theory of Games and Economic Behaviour*. Princeton University Press, 1994.
- [36] Rudolf Vetschera. Preference structures and negotiator behavior in electronic negotiations. *Decision Support Systems*, 44(1): 135-146, 2007.
- [37] Jacques Wainer, Paulo Roberto Ferreira Jr. and Everton Rufino Constantino. Scheduling meetings through multi-agent negotiations. *Decision Support Systems*, 44(1): 285-297, 2007
- [38] Huaiqing Wang, Stephen Liao, and Lejian Liao. Modeling constraint-based negotiating agents. *Decision Support Systems*, 33(2): 201-217, 2002.
- [39] M. Wooldridge and N. Jennings. Intelligent Agents: Theory and Practice. *Knowledge Engineering Review*, 10(2):115–152, 1995.
- [40] D. Zeng and K. Sycara. Bayesian learning in negotiation. *International Journal of Human-Computer Studies*, 48(1):125–141, 1998.

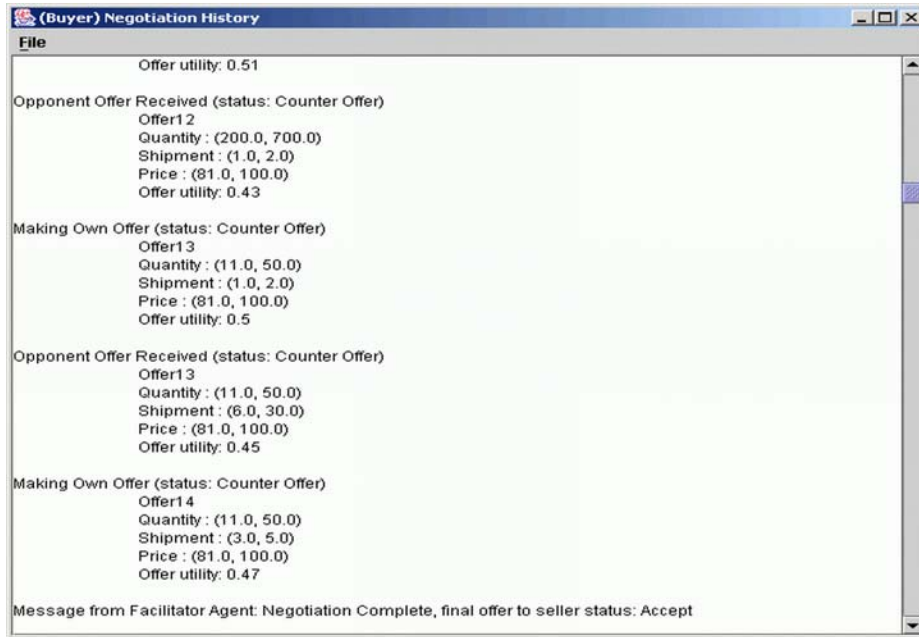


Figure 1. A Segment of Negotiation History

Table 1. A Weighted Sample Space

Session	Offers	Price (d_{a_1})	Shipment Time (d_{a_2})	QTY (d_{a_3})	Opponent Accept (c_j)	Weights
4	o_1	5-10	1-2	20-30	N	200
	o_2	15-20	3-4	50-50	Y	175
	o_3	1-2	2-2	10-20	N	150
	o_4	25-30	5-8	60-100	Y	125
3	o_1	5-10	1-2	20-30	N	300
	o_2	15-20	3-4	50-50	Y	275
	o_3	1-2	2-2	10-20	N	250
	o_4	25-30	5-8	60-100	Y	225
2	o_1	5-10	1-2	20-30	N	400
	o_2	15-20	3-4	50-50	Y	375
	o_3	1-2	2-2	10-20	N	350
	o_4	25-30	5-8	60-100	Y	325
1	o_1	5-10	1-2	20-30	N	500

Table 2. The Utility Functions of Two Agents

Agent	Buyer: p_1			
$U_{p_1}^A(\text{price}) = 0.3$	$U_{p_1}^A(\text{qty}) = 0.3$	$U_{p_1}^A(\text{size}) = 0.2$	$U_{p_1}^A(\text{delivery}) = 0.1$	$U_{p_1}^A(\text{warranty}) = 0.1$
$U_{p_1}^{D_{price}}(1) = 0.9$	$U_{p_1}^{D_{qty}}(1) = 0.9$	$U_{p_1}^{D_{size}}(1) = 0.9$	$U_{p_1}^{D_{delivery}}(1) = 0.9$	$U_{p_1}^{D_{warranty}}(1) = 0.9$
$U_{p_1}^{D_{price}}(2) = 0.8$	$U_{p_1}^{D_{qty}}(2) = 0.8$	$U_{p_1}^{D_{size}}(2) = 0.8$	$U_{p_1}^{D_{delivery}}(2) = 0.8$	$U_{p_1}^{D_{warranty}}(2) = 0.8$
$U_{p_1}^{D_{price}}(3) = 0.6$	$U_{p_1}^{D_{qty}}(3) = 0.6$	$U_{p_1}^{D_{size}}(3) = 0.6$	$U_{p_1}^{D_{delivery}}(3) = 0.6$	$U_{p_1}^{D_{warranty}}(3) = 0.6$
$U_{p_1}^{D_{price}}(4) = 0.2$	$U_{p_1}^{D_{qty}}(4) = 0.2$	$U_{p_1}^{D_{size}}(4) = 0.2$	$U_{p_1}^{D_{delivery}}(4) = 0.2$	$U_{p_1}^{D_{warranty}}(4) = 0.2$
$U_{p_1}^{D_{price}}(5) = 0.1$	$U_{p_1}^{D_{qty}}(5) = 0.1$	$U_{p_1}^{D_{size}}(5) = 0.1$	$U_{p_1}^{D_{delivery}}(5) = 0.1$	$U_{p_1}^{D_{warranty}}(5) = 0.1$
Agent	Seller: p_2			
$U_{p_2}^A(\text{price}) = 0.3$	$U_{p_2}^A(\text{qty}) = 0.3$	$U_{p_2}^A(\text{size}) = 0.2$	$U_{p_2}^A(\text{delivery}) = 0.1$	$U_{p_2}^A(\text{warranty}) = 0.1$
$U_{p_2}^{D_{price}}(1) = 0.1$	$U_{p_2}^{D_{qty}}(1) = 0.1$	$U_{p_2}^{D_{size}}(1) = 0.1$	$U_{p_2}^{D_{delivery}}(1) = 0.1$	$U_{p_2}^{D_{warranty}}(1) = 0.1$
$U_{p_2}^{D_{price}}(2) = 0.2$	$U_{p_2}^{D_{qty}}(2) = 0.2$	$U_{p_2}^{D_{size}}(2) = 0.2$	$U_{p_2}^{D_{delivery}}(2) = 0.2$	$U_{p_2}^{D_{warranty}}(2) = 0.2$
$U_{p_2}^{D_{price}}(3) = 0.5$	$U_{p_2}^{D_{qty}}(3) = 0.5$	$U_{p_2}^{D_{size}}(3) = 0.5$	$U_{p_2}^{D_{delivery}}(3) = 0.5$	$U_{p_2}^{D_{warranty}}(3) = 0.5$
$U_{p_2}^{D_{price}}(4) = 0.8$	$U_{p_2}^{D_{qty}}(4) = 0.8$	$U_{p_2}^{D_{size}}(4) = 0.8$	$U_{p_2}^{D_{delivery}}(4) = 0.8$	$U_{p_2}^{D_{warranty}}(4) = 0.8$
$U_{p_2}^{D_{price}}(5) = 0.9$	$U_{p_2}^{D_{qty}}(5) = 0.9$	$U_{p_2}^{D_{size}}(5) = 0.9$	$U_{p_2}^{D_{delivery}}(5) = 0.9$	$U_{p_2}^{D_{warranty}}(5) = 0.9$

Table 3. The Impact of Bayesian Learning on Negotiations

Group	Pareto Optimal		Probabilistic		
	Avg. Joint-Util.	Avg. Time	Avg. Joint-Util.	Avg. Time	Avg. Dist.
1	2.74	1.0	2.74	1.0	0.00
2	2.41	562.0	2.35	459.0	0.15
3	2.38	813.0	2.31	622.5	0.18
4	2.29	1036.0	2.23	804.0	0.17
5	2.15	1483.5	2.04	1110.0	0.22
6	1.87	1734.0	1.73	1352.0	0.26
Mean	2.31	938.25	2.23	724.75	0.16

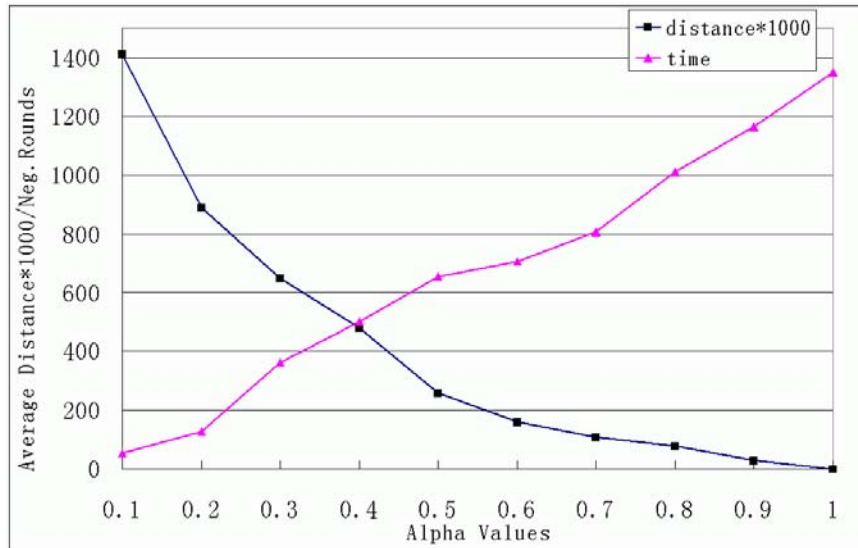


Figure 2. The Impact of α on Agent's Performance

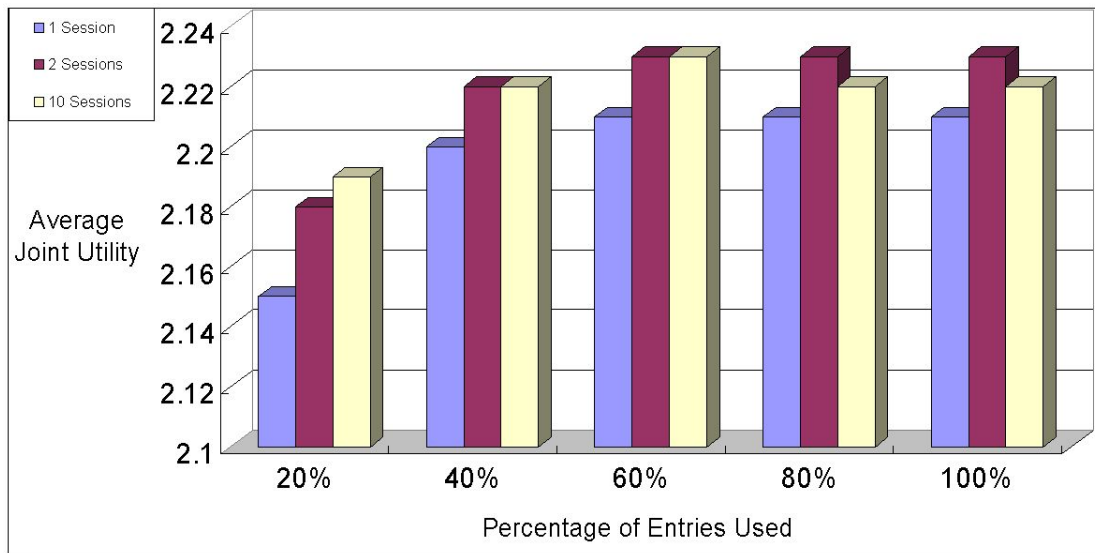


Figure 3. The Impact of Negotiation History on Agent's Performance

Table 4. Agent Performance Under Time Pressure

Group	Pareto Optimal	Probabilistic Eq.(1)	Probabilistic Eq.(5)
1	2.74	2.74	2.74
2	1.93	2.16	2.31
3	1.66	1.81	2.02
4	1.15	1.34	1.95
5	0.86	1.05	1.88
6	0.19	0.36	1.63
Mean	1.42	1.58	2.09

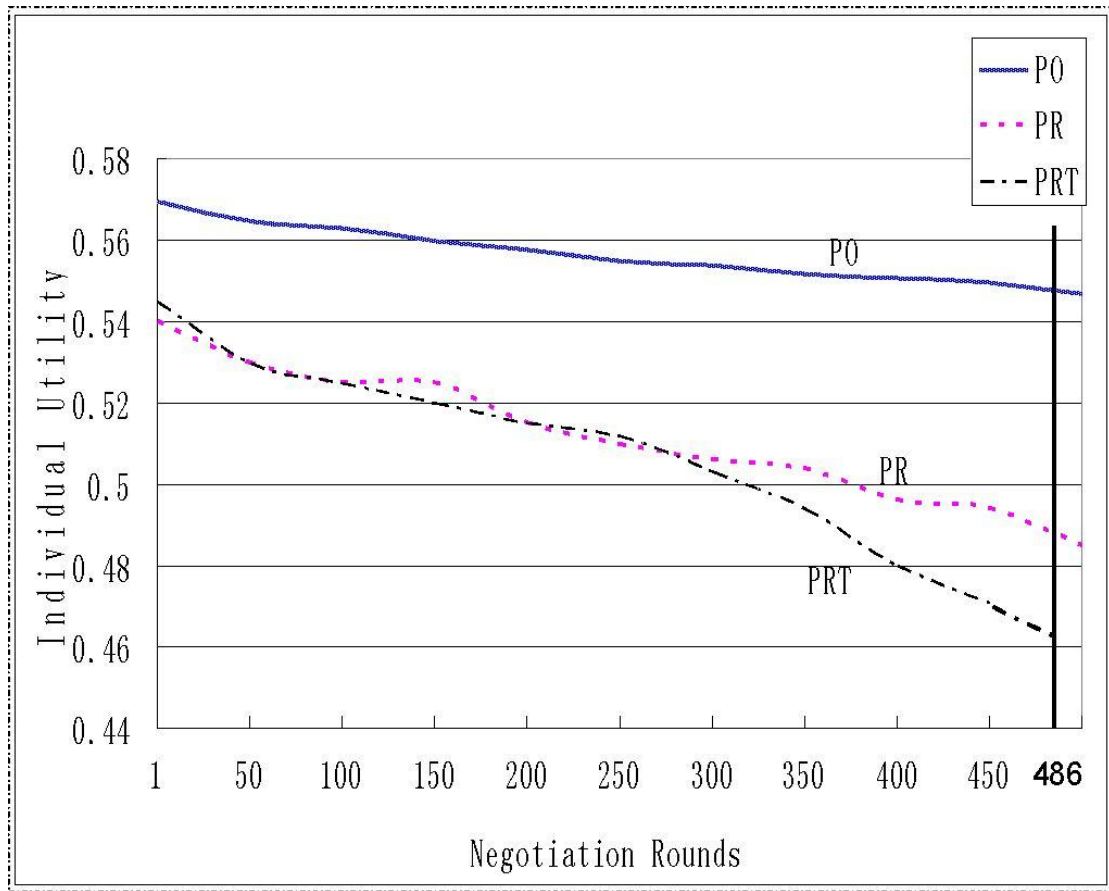


Figure 4. Comparative Concession Behavior Among Three Types of Agents

Table 5. Performance of Adaptive Negotiation Agents

Group	Evolutionary Agents		Probabilistic Agents	
	Avg. Joint Utility	Avg. Dist.	Avg. Joint Utility	Avg. Dist.
1	2.74	0.00	2.74	0.00
2	2.41	0.15	2.42	0.14
3	2.16	0.21	2.14	0.22
4	2.12	0.18	2.11	0.19
5	1.96	0.23	1.96	0.23
6	1.79	0.25	1.78	0.25

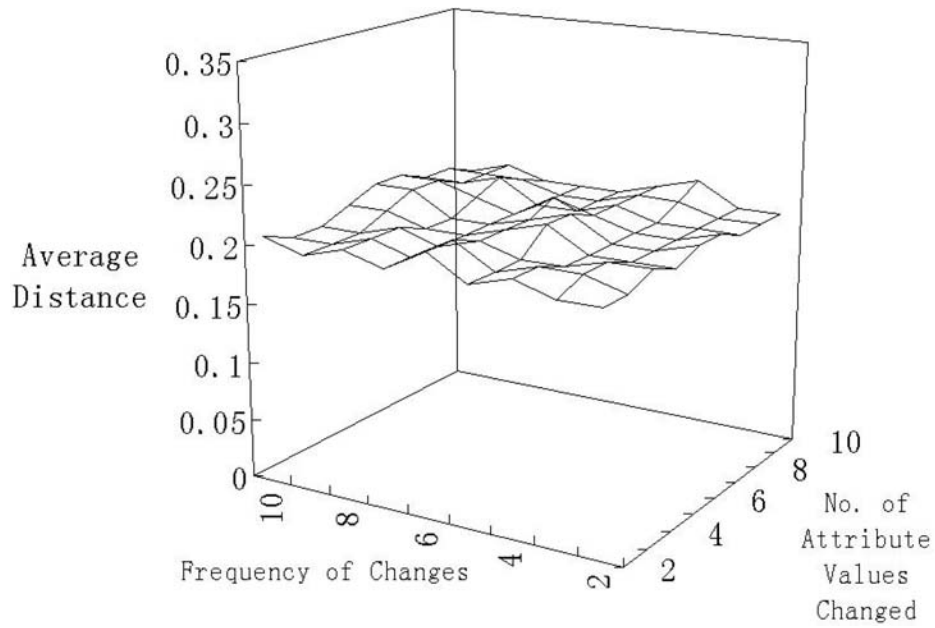


Figure 5. The Impact of Preferential Changes on Average Distance from Pareto Optimum

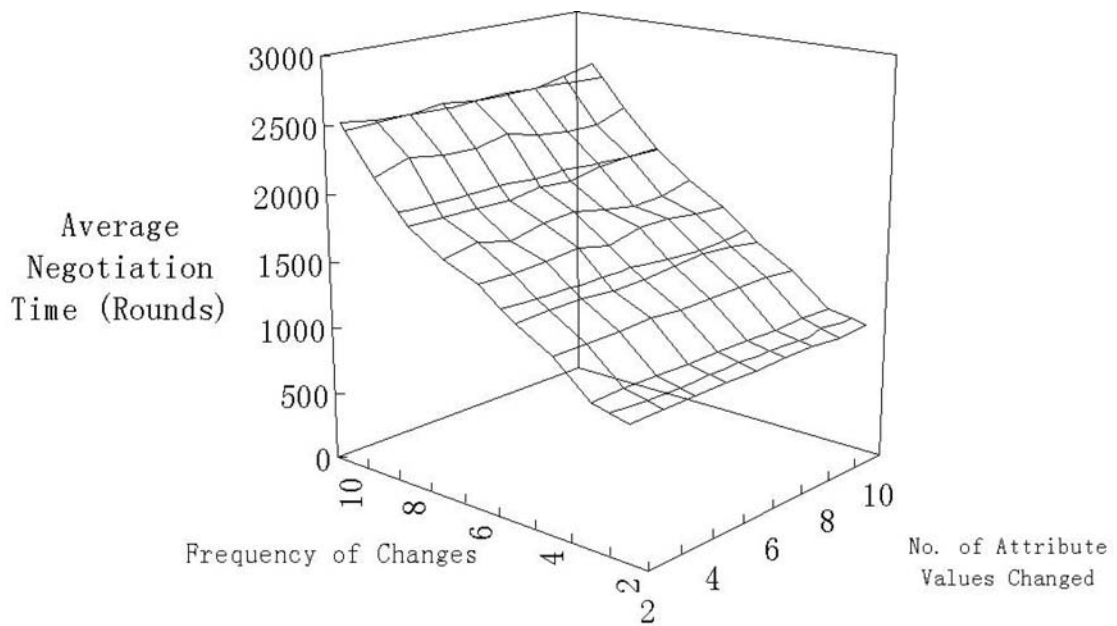


Figure 6. The Impact of Preferential Changes on Average Negotiation Time