A Bayesian classifier for learning opponents' preferences in multi-object automated negotiation

Buffett, Scott; Spencer, Bruce

Publisher’s version / Version de l'éditeur:
https://doi.org/10.1016/j.elerap.2006.06.008
Electronic Commerce Research and Applications, 6, 3, pp. 274-284, 2006-07-14

National Research Council Canada
Institute for Information Technology

Conseil national de recherches Canada
Institut de technologie de l'information

NRC · CNRC

*A Bayesian Classifier for Learning Opponents' Preferences in Multi-Object Automated Negotiation **

Buffett, S., Spencer, B.
2007

* published in the Electronic Commerce Research and Applications Journal. Volume 6, Number 3. pp. 274-284. 2007. NRC 48500.

Canada

Available online at www.sciencedirect.com

ScienceDirect

Electronic Commerce Research and Applications 6 (2007) 274–284

Electronic Commerce Research and Applications

www.elsevier.com/locate/ecra

A Bayesian classifier for learning opponents' preferences in multi-object automated negotiation

Scott Buffett *, Bruce Spencer [1]

Institute for Information Technology – e-Business, National Research Council, 46 Dineen Drive, Fredericton, NB, Canada E3B 9W4
University of New Brunswick, Fredericton, NB, Canada E3B 5A3

Received 2 April 2006; accepted 22 June 2006
Available online 14 July 2006

Abstract

We present a classification method for learning an opponent's preferences during a bilateral multi-issue negotiation. Similar candidate preference relations over the set of offers are grouped into classes, and a Bayesian technique is used to determine, for each class, the likelihood that the opponent's true preference relation lies in that class. Evidence used for classification decision-making is obtained by observing the opponent's sequence of offers, and applying the concession assumption, which states that negotiators usually decrease their offer utilities as time passes in order to find a deal. Simple experiments show that the technique can find the correct class after very few offers and can select a preference relation that is likely to match closely with the opponent's true preferences.

Keywords: Automated negotiation; Multi-issue; Utility; Preference elicitation; Bayesian classification

1. Introduction

Given the speed with which transactions can be negotiated and executed through various electronic services today, research in intelligent agent technology has been focusing increasingly on automated negotiation [8,12,13]. Automated negotiation technology makes it possible for two or more parties to explore a large space of possible outcomes or agreements, with the hope of finding one that is mutually beneficial to all. In multi-agent systems, cooperative agents can exchange proposals for assigning tasks or allocating resources until one is found that satisfies sufficiently or, better yet, optimally, the goals of system functions in terms of time, cost or overall productivity. Alternatively, uncooperative agents may also negotiate with the goal of finding an outcome that best meets their own needs. The advantage of this automated negotiation is that computerized agents can compose, communicate and evaluate proposals quickly in comparison with a human user, and have the processing power to construct effective negotiation protocols and strategies in dynamic environments.

In electronic commerce, automated negotiation can play a pivotal role in the successful completion of transactions. Instituting negotiation capabilities for the price of an item, for example, can increase the likelihood of a sale. This is because the common model of take-it-or-leave-it pricing is far too rigid. One will likely find that, in many situations, a seller would be more willing to accept a price that is slightly lower than the asking price than to have the buyer abandon the transaction altogether. Negotiation is thus necessary for the buyer and seller to determine whether a mutually acceptable price exists. While price negotiation is commonly performed by human buyers and sellers, the case for automating this negotiation becomes much stronger when other factors are introduced to the potential agreements, such as the attributes of the item for sale

* Corresponding author. Address: Institute for Information Technology – e-Business, National Research Council, 46 Dineen Drive, Fredericton, NB, Canada E3B 9W4. Tel.: +1 506 444 0386; fax: +1 506 444 6114.
E-mail addresses: scott.buffett@nrc.gc.ca (S. Buffett), bruce.spencer@nrc.gc.ca (B. Spencer).
[1] Tel.: +1 506 444 0384; fax: +1 506 444 6114.

1567-4223/$ - see front matter.
doi:10.1016/j.elerap.2006.06.008

(e.g. size, quantity, colour, etc.), or other factors related to the transactions such as delivery and warranty. Factors in the transaction that are not directly related to the exchange of goods may also need to be agreed upon. This may include the associated exchange of the buyer's private information. Some information may be required for the completion of the transaction, such as credit card information and home address for delivery, while other information such as age, sex and e-mail address might be requested for the purpose of determining target demographics or marketing. Such information exchange is likely up for negotiation as well. As these new factors are introduced, the number of potential agreements tends to grow exponentially. Thus automated negotiation can greatly help potential transaction partners find agreements that are not only mutually acceptable, but also much more mutually beneficial than they might find on their own.

Much work has been done recently on automated negotiation in the areas of protocol design, strategy computation and user utility elicitation, in various negotiation models such as bilateral negotiation (single-issue and multi-issue) and auctioning. However, not much effort to date has been put into the problem of learning opponents' preferences, particularly in the area of multi-issue negotiation.

In single-issue bilateral negotiation, where typically price is the only issue, there is a clear understanding between the two negotiating parties of the other's preferences over the negotiation domain. The receiver of the money (e.g. the seller in a purchase transaction) typically prefers more to less, while the opposite is true for the giver (buyer). One might not know the shape of the opponent's utility curve over the set of offers, the opponent's concession rate or deadline, but the preference relation over the set is known fully.

In multi-issue bilateral negotiation, on the other hand, there may be some issues under negotiation for which the opponent's preferences are not known. In fact, there may even be preferences that the two sides have in common. This makes negotiation difficult, since a negotiator must have some degree of understanding of the opponent's preferences in order to build effective negotiation strategies. To date, what little work exists in learning about opponents typically assumes that several interactions will take place, over which the preferences will gradually be learned [1,11].

In this paper, we discuss the multi-object negotiation model, where subsets of a set of objects are under negotiation, and show that this is a special case of the multi-issue negotiation model. Under the multi-object model, we demonstrate a technique for learning the opponent's preferences over subsets *during* a negotiation. One setting where such a negotiation might take place is in the realm of privacy. A website might request several items of personal information from a user in order to complete a transaction, and negotiation can take place to determine which subset of those items is suitable to the website and the user. Here, a partial order over the opponent's preferences is known. In particular, the receiver of the items (assuming that the items are desirable) will necessarily prefer offer a over offer b if a is a superset of b. The reverse is true for the giver. However, if neither is a subset of the other, it is not immediately clear which is preferred. To fill in these missing preferences, we can observe or predict that users typically behave in one of several ways. Our method uses a Bayesian classification technique that decides in which of these predefined classes a new opponent's total ordering is likely to reside. This decision is based on the opponent's offers made thus far in the negotiation. The ultimate goal is to learn as much as possible about the total order of the opponent's preferences so that an effective negotiation strategy can be devised.

The paper is organized as follows. In Section 2 we formalize the framework in which we consider our negotiations, and define the multi-object negotiation model. A protocol from the literature that can be used for such a negotiation model is then discussed. In Section 3 we give a brief introduction Bayesian classification, and introduce the concept of using such a scheme for classifying an opponent's preferences during a negotiation. Next we detail the specifics of our particular classification system in Section 4. To demonstrate the flexibility of our idea, we show how the technique can be extended for use in the more general model of multi-issue negotiation in Section 5. Section 6 then sheds some light on the effectiveness of our technique by describing experimental results, and finally Sections 7 and 8 offer conclusions and discuss plans for future work.

2. Negotiation framework

2.1. The PrivacyPact negotiation protocol

The PrivacyPact protocol [2] was originally developed as a protocol for alternating-offers bilateral negotiation of private information exchanges. However, with simple adjustments the protocol can be used to dictate the rules for exchanges of subsets of objects in general.

The PrivacyPact protocol is a protocol for alternating-offers bilateral negotiation of private information exchanges between a website (the requestor or consumer of private information) and a web user (the provider or producer of private information). Each offer under the protocol consists of two components: a Platform for Privacy Preferences (P3P) [6] statement, which denotes a set of private information units (as well as specifics on how this data will be treated) that the user will provide, and a reward, if any, that the user receives in return. The protocol dictates what offers may and may not be proposed given a negotiation history, in an effort to guide the exchange to efficient convergence. In particular, an actor in the negotiation cannot make an offer that is necessarily worse to the opponent than one the actor had previously made. Specifically, the website cannot ask for a superset of the information requested in a previous offer in exchange for a smaller reward than specified in that offer. Conversely, the user cannot ask to provide a subset of the information requested in a previous offer in exchange for a larger reward than specified in that offer.

The rules of communication as specified by PrivacyPact are as follows. In all cases it is assumed that if either partner wants to discontinue, the communication port can be closed. Initially the participants agree on the domain of negotiation. The first message comes from the consumer, indicating the desired information units and the set of rewards it has to offer. The producer responds by selecting a subset of those information units that are available, and the rewards that it is interested in obtaining. Thereafter messages representing offers alternate from each participant to the other, starting first with an offer from the consumer. All offers remain on the table. Thus any time the consumer asks for a subset of a previous offer made by the producer for less or equal reward, or the producer offers a superset of an offer previously made by the consumer for more or equal rewards, the protocol imposes an end to the negotiation with this final offer being the agreement.

The following section formalizes the negotiation process used for our model, and specifies how the PrivacyPact protocol is used as a framework for our negotiations.

2.2. Mutli-object negotiation formalization

We consider a two-participant bilateral negotiation where each participant is self-interested and has incomplete information about the opponent. Information is incomplete in that a participant is unsure not only about the opponent's reserve limits and deadlines, but also about its preference ranking of possible offers. Negotiation takes place over a set of objects, where the two participants need to agree on which subset of the objects will be traded, if any. Let the participants p and c be the producer and the consumer of the objects, respectively, where the producer is the actor that will send items in the resulting transactions, and the consumer will receive them. Let S be the set of objects under negotiation, and let each offer be a subset $s \subseteq S$. Note that there typically may be other issues under negotiation at the same time, such as the price to pay for the objects. This is the case with the PrivacyPact protocol. However, we assume that these other issues are mutually utility independent with respect to the set of objects (i.e. preferences over sets of objects do not depend outcomes for other issues, and vice-versa), and focus solely on determining the preference relation over the set of subsets, which will remain consistent regardless of the values of other issues. For this reason, the remainder of the paper will only consider that the issue under negotiation is the set of objects. We do however relax this restriction in the discussion on generalizing the technique to multi-issue negotiation.

While the utility functions are private, a partial order of each participant's preference ranking over S is mutually known. Specifically, we assume that the consumer c necessarily values an offer s no more than another offer s' if $s \subseteq s'$. For the producer the reverse is true. Thus the utility function $u_a : 2^S \rightarrow \Re$ for each actor $a \in \{p, c\}$ is such that $\forall s, s' \subseteq S,\ s \subseteq s' \Rightarrow u_p(s) \geqslant u_p(s')$ and $u_c(s) \leqslant u_c(s')$. When $u_p(s)$ is greater than $u_p(s')$, we say that p *prefers* s over s', and denote this by $s \succ_p s'$. We also assume that each actor has a *break-even point* α_a, which is the point at which no offer s such that $u_a(s) < \alpha_a$ is acceptable, and a deadline d_a by which a deal must be made. Utility may also be a function of time, perhaps decreasing in such a way that no deal made past time d_a will have utility greater than α_a. We maintain the utility independence assumption and assume that an actor's preference relation over S stays constant regardless of time.

The PrivacyPact protocol can be adopted for the framework presented in this paper as follows. The first message comes from the consumer, indicating the desired objects. The producer responds by selecting a subset of those elements that are available, giving S. Thereafter, messages representing offers alternate from one participant to the other, starting first with an offer from the consumer. Each offer consists of some $s \subseteq S$. To avoid wasting time, offers are submitted under the constraint that no actor may give an offer that is necessarily worse for the opponent than a previous offer, according to the partial orders defined above. That is, the consumer cannot ask for a superset of an offer it made previously, while the producer cannot offer a subset that it offered previously. All offers remain on the table. Thus any time the consumer asks for a subset of a previous offer made by the producer, or the producer offers a superset of an offer previously made by the consumer, the protocol imposes an end to the negotiation with this final offer being the agreement.

3. Classification of preference relations

Using Bayesian classification is a novel approach to learning a negotiation agent's preferences. In this section we introduce the general Bayesian classification technique, and give a high-level presentation on how the approach can be applied to the problem of learning preferences. Specifics on the operation of our classification system are given in the next section.

3.1. Bayesian classification

Bayesian classification is a technique from the field of supervised machine learning where objects are assigned to classes based on the likelihoods of the observed attributes or *evidence*. Given a set of classes, a Bayesian classifier is provided with information on the attributes of objects that belong to each class. When presented with new unclassified objects, the classifier makes a decision on which class is most likely to include the object. As a simple example, consider a Bayesian classifier that classifies documents into one of two classes: literary and scientific. The classifier uses information on the likely attributes of members of each class (e.g. a document containing the word "hypothesis" is more likely to be from the scientific class). This probabilistic information can be obtained by observing the classification of several objects where the classes are known, and noting the frequency at which objects with certain characteristics are assigned to each class.

The probability model for a Bayesian classifier is as follows. Let $\mathcal{C}$ be the set of classes. The probability of an object belonging to a class $C \in \mathcal{C}$ given the observed evidence E is denoted by $P(C|E)$. This can be computed using Bayes' formula

$$P(C|E) = \frac{P(C) \times P(E|C)}{P(E)} \tag{1}$$

where $P(C)$ is the prior probability an object belonging to class C, $P(E|C)$ is the probability of observing E given that the object belongs to C, and $P(E)$ is the probability of observing E. Returning to the document classification example, consider the initial observations that 60% of the documents are scientific (S) as opposed to 40% literary (L), and 70% of all scientific documents contained the word "hypothesis" as opposed to 5% in literary. Then the probability $P(E)$ of observing the evidence word "hypothesis" is $P(E|S)P(S) + P(E|L)P(L) = (0.7)(0.6) + (0.05)(0.4) = 0.44$, and thus the probability of a document containing "hypothesis" belonging to the scientific class is $\frac{0.6\times0.7}{0.44} = 0.95$, where the probability of such a document belonging to the literary class is $\frac{0.4\times0.05}{0.44} = 0.05$.

The evidence used to perform classification may consist of possibly several attributes or features. For example, we may also observe that 55% of scientific documents contain the word "proof", and that 37% of literary documents contain the word "poetry". The probability of each class can be computed for a document containing "hypothesis" and "poetry", but not "proof", for example. When there are several such features, composing probability distributions over all joint outcomes can become infeasible. In this case, the individual features are assumed to be independent. In most cases it is obvious that such an assumption is unrealistic. However, reasonably favourable results are still often observed in these cases, especially given the savings in time and space. Such classification is referred to as *naïve* Bayesian classification.

In this paper, we partition the possible preference relations into classes, and determine which class is most likely to contain the opponent's preference relation, based on the evidence obtained during a negotiation session. In this case, the evidence used for classification is based on the order in which offers are proposed by the opponent. The idea is that, if there is sufficient evidence to conclude the correct class to which the opponent belongs, then some assumptions can be made about his/her preferences which can be useful when the negotiation engine attempts to build an effective negotiation strategy.

3.2. Bayesian classification of preference relations

The technique for learning opponents' preference relations presented in this paper is based on Bayesian classification. Given a set of classes of preference relations, the idea is to determine the likelihood of the opponent's preference relation being a member of each class. The likelihood of a class is determined by considering the prior probability that a given relation is a member of the class, as well as the likelihood of the evidence obtained during negotiation, given that the relation is a member of the class. The evidence here is based on the set of offers received from the opponent.

All preference relations in a class are relatively similar. That way, we do not need to pinpoint the full preference relation with certainty; we only need to examine the evidence and determine which class is most likely to include the relation. Then any relation in the class should be reasonably close to the opponent's relation and a reasonably effective negotiation strategy can be computed.

Given the order of the offers proposed by the opponent up to a certain point in a negotiation session, some assumptions can be made about the opponent's preferences for various offers. In particular, one idea is to assume that offers made earlier in the negotiation are more preferable to the opponent than offers it made later in the session, since rational negotiators tend to make gradual concessions in an effort to guide negotiation convergence toward a mutually acceptable deal. One cannot assume however that even a rational opponent will concede 100% of the time, so the opponent's most likely preference relation cannot be constructed solely on these assumed preferences. Instead, we answer the question "If these assumed preferences are true, how *likely* is the true preference relation to be in each class?" These probabilities are constructed based on the portion of preference relations that are violated in each class. The best candidate class may hold several preference relations that are violated by some of these assumed preferences, and one of those violated may in fact be the true preference relation. We concede the fact that it is nearly impossible to determine the opponent's complete preference relation during negotiation. The goal is instead to determine a class of relations where each member is likely to be quite close to the opponent's true preferences.

We employ a Bayesian classifier because of its power to learn the likelihood of each class from several patterns in the data. By simply counting the number of preference relations in each class that are violated by the assumed preferences, a human (or some other unsophisticated) classifier may be able to make some simple assumptions. For example, if 20%, 40% and 80% of preference relations from class 1, 2 and 3, respectively, are violated, then one might assume that the likelihood of class 1 should be 50%, since 50% of the non-violated relations are in class 1 (if all classes are the same size). However, a Bayesian classifier may learn that if class 3 has a high violation rate, then class 1 has a much higher likelihood, like 90% for example. This may be for a variety of reasons, such as class 3 relations being very different from those in class 1. The true causes are not really important; all that matters is that the evidence is there to support the classification.

3.3. Inferring preferences from concessions

The classification technique demonstrated in this paper capitalizes on the assumption that the opponent makes concessions during negotiation. Thus, each time a new offer s is received, we assume that the opponent has less utility for s than any of its previous offers. More formally, let s_i and s_j be offers made by a such that $i < j$. Then it is likely that $u_a(s_i) > u_a(s_j)$.

In multi-object negotiation, if an opponent's preference over two offers can be determined, several other preferences can often be inferred. We refer to a set S of items as *utility independent* if, for any subsets s_1, s_2, s_3 and s_4 of S, an agent prefers $s_1 \cup s_3$ over $s_1 \cup s_4$ if and only if it prefers $s_2 \cup s_3$ over $s_2 \cup s_4$. Thus, an agent's preference over two subsets of items does not depend on any other items that are involved. For example, consider the set of items $S = \{a, b, c, d\}$ and two offers $\{a, b\}$ and $\{a, c\}$, where an agent is known to prefer $\{a, b\}$ over $\{a, c\}$. Then, under the assumption of utility independence, we can conclude that (1) $\{a, b, d\}$ is preferred over $\{a, c, d\}$, and (2) $\{b\}$ is preferred over $\{c\}$.

When utility independence can not be assumed to hold (which is quite often the case), preferences such as those stated cannot be inferred with 100% confidence. However, there are scenarios that can provide sufficient reason to believe that certain preferences will hold with a relatively high probability. The following rules can be used to infer such preferences. Let S be the set of negotiation items and s_1, s_2 and s_3 be any subsets of S:

1. *If* $s_1 \succ s_2$ *and* $s_1 \succ s_3$, *then* $s_1 \cup s_3 \succ s_2 \cup s_3$ *and* $s_1 \cup s_2 \succ s_2 \cup s_3$. For example, let $s_1 = \{a\}$, $s_2 = \{b\}$ and $s_3 = \{c\}$. If $\{a\} \succ \{b\}$, then without the utility independence assumption we cannot necessarily conclude that $\{a, c\} \succ \{b, c\}$. However, the added evidence that $\{a\} \succ \{c\}$ increases the likelihood. Assume the opponent a' is the consumer of items (i.e. inclusion is preferred to exclusion) with preference relation $\succ_{a'}$. Then we know with certainty that $\{a, c\} \succ_{a'} \{a\}$, $\{b, c\} \succ_{a'} \{b\}$ and $\{b, c\} \succ_{a'} \{c\}$. Then the lower-bound for $u_c(\{a, c\})$ is $u_c(\{a\})$ and the lower-bound for $u_c(\{b, c\})$ is $\max\{u_c(\{b\}), u_c(\{c\})\}$. Since $u_c(\{a\}) > \max\{u_c(\{b\}), u_c(\{c\})\}$, even in the worst case where $\{a\} \succ_{a'} \{b\}$ has no bearing on whether $\{a, c\} \succ_{a'} \{b, c\}$ and utilities are essentially random (within the constraints), then it is more than 50% likely that $\{a, c\} \succ_{a'} \{b, c\}$ will hold, since $\{b, c\}$ has the lower lower-bound. This worst case is not very realistic, so even when the utility independence assumption does not hold in general, there should be a fairly high probability that $\{a, c\} \succ_{a'} \{b, c\}$ holds. A similar explanation exists for $s_1 \cup s_2 \succ s_2 \cup s_3$, as well as for showing that these inferences can be made when the opponent is the producer of items.
2. *If* $s_1 \cup s_3 \succ s_2 \cup s_3$ *and* $s_1 \succ s_3$, *then* $s_1 \succ s_2$. This can be shown with a reversal of the argument to that above.

While these inferred preferences have high likelihood of holding in most cases, determining just how high and whether such likelihood is sufficient requires further investigation. We defer these questions to future work.

3.4. The concession assumption and Pareto optimality

The important aspect of the evidence observed during a negotiation session is the order in which offers are received. Typically during a negotiation session, an actor has a tendency to make concessions as time passes. So if the opponent a follows offer s_i with s_{i+1}, then there is a high probability that $u_a(s_i) > u_a(s_{i+1})$. This information can be used to determine the likelihood of a class. For example, if 71% of the relations in class 1 specify $u_a(s_i) > u_a(s_{i+1})$, while only 37% of the relations in class 2 specify this preference, then this could indicate that the class 1 is more likely to be the correct class. Note that we do not consider $u_a(s_i) > u_a(s_{i+1})$ to be 100% true and subsequently disregard all preference relations where this does not hold. We merely accept that there is a sufficient likelihood that it is true, and consider that classes where this preference is fairly common will have increased likelihood, even though these classes will contain preference relations where this does not hold.

There are various reasons why an agent will choose not to concede at a particular stage in a negotiation. For example, an agent may choose to make more random offers, in an effort to confuse the negotiation partner. Alternatively, an agent could purposely make offers that are progressively worse for the other party, in an effort to wear on the other party's patience and hopefully force him to offer or accept a sub-standard deal. However, given the protocol that we employ, this can be viewed as somewhat irrational behaviour. Since the PrivacyPact protocol dictates that all previous offers remain on the table, there is incentive for the negotiators to make gradual concessions. A larger concession followed by a return back up the utility scale will likely hamper the negotiator's leverage, since that large concession will always be available to the opponent.

Even agents that display more rational behaviour will choose not to concede on occasion. Consider two agents a_1 and a_2 with utility functions u_1 and u_2 over the set of offers. Agent a_1 submits offer s. Offer s is said to be *Pareto optimal* if there exists no other offer s' such that $u_1(s') \geqslant u_1(s)$ and $u_2(s') \geqslant u_2(s)$ and either $u_1(s') > u_1(s)$ and $u_2(s') > u_2(s)$. That is, no other offer exists that is better for at least one of the parties and is worse for none. In zero-sum negotiations where the subject of negotiation is a single issue such as money or some other divisible good, all offers are Pareto optimal. However, this is not the case in most multi-attribute negotiations. Consider two attributes

X and Y, each with possible values from $\{x, x'\}$ and $\{y, y'\}$, respectively. If agent a_1 most prefers an outcome with attribute values x and y, while agent a_2 most prefers an outcome with attribute values x' and y', then it is possible that both would prefer an outcome with attribute values x and y' over x' and y. Then the outcome with x' and y is Pareto non-optimal.

After agent a_1 above submits s, he will observe a_2's counteroffer (assuming that a_2 does not accept s or quit). If a_2's counteroffer closely resembles an offer s', a_1 may conclude that s' is more preferable to a_2 than s. If a_1 also prefers s' to s, then it can be quite beneficial for a_1 to offer s'. So the fact that s was Pareto non-optimal played a part in a_1 making an offer that was not a concession. Since non-concessions can deteriorate the effectiveness of our technique, we are currently working on methods for to identifying likely occurrences so they can be removed from consideration. This is however outside the scope of this particular paper, so we defer it to future work. See Section 8 for preliminary details on this effort.

4. The classification system

In this section we detail the specifics our classification system. Each time a new offer is received, the probability that a class holds the true preference relation is computed using Bayes' Rule (Eq. (1)). Here we demonstrate how the initial classes are constructed, and describe our methodology for computing the components of Eq. (1), namely the prior probability $P(C)$ of each class, the probability $P(E)$ of observed evidence, and the probability $P(E|C)$ of observed evidence given each class. We then bring these components together and present our classification mechanism.

4.1. The initial set of classes

Determining the initial set of classes that will work optimally with the proposed classification mechanism is a problem in itself. Techniques from the area of data mining, such as collaborative filtering, may work well here. Different techniques may work better than others depending on the specifics of the problem and perhaps any extra information known about the opponent. Regardless, the main idea is that classes should be constructed in such a way that all preference relations in the same class should be fairly similar to each other, and preferably much less similar to relations in different classes. Since we expect only to determine the correct class of preference relations during the negotiation, then these preference relations will need to be similar to the true preference relation in order to have any real value.

In this paper we utilize a technique based on k-means clustering. Let $\mathcal{C}$ be the set of k clusters (classes) in which to partition the set PR of possible preference relations. Let the *distance* $d(pr, pr')$ between two preference relations pr and pr' be the number of preferences that hold in pr, but do not hold in pr'. Thus,

$$d(pr, pr') = |\{(a, b) | a \preceq b \text{ in } pr \wedge a \npreceq b \text{ in } pr'\}| \tag{2}$$

Let the set $PR_c \subseteq PR$ of centroids be a subset of the preference relations with $PR_c = k$. As is the case with k-means clustering, the initial centroids should be chosen such that their distances from each other are relatively high. Each iteration in the clustering phase is as follows. Assign each relation $pr \in PR$ to the cluster C_i such that the distance $d(pr, c_i)$ to the centroid c_i associated with cluster C_i is minimized over all centroids. Once the assignments are completed, a new centroid is chosen from each cluster. In traditional k-means clustering, the new centroid for a cluster is taken to be the point such that mean distance from that point to the points in the cluster is minimized. We do something similar. The new centroid for a cluster C is chosen to be the preference relation $pr \in C$ such the mean distance $d(pr, pr')$ for all $pr' \in C$ is minimized. All relations in PR are then reassigned as stated above. The process continues until there are two consecutive iterations that result in the same set of clusters.

The prior probability $P(C)$ of a class can then be assessed either by examining the class's preference relations and using expert knowledge to determine the likelihood of a particular negotiator possessing such relations, or simply by treating the probabilities across relations uniformly and assigning the class's probability based on the proportion of preference relations that it contains. Historical information might be used to determine the prior probabilities as well.

4.2. Negotiation evidence

Our classifier works by learning and inferring particular preferences from the negotiation history that likely belong to the opponent, and observing the percentage of candidate preference relations that are *violated* as a result of these preferences. Let pr_H be a partial preference relation over the set of offers S, representing the preferences learned or inferred from the negotiation history H, and let pr be a candidate preference relation for the opponent. If there exist $s, s' \in S$ such that $s \succ s'$ in pr_H and $s' \succ s$ in pr, then pr is said to be violated by H, since if the preferences in pr_H are true, then pr cannot possibly be the opponent's true preference relation. The percentage of preferences from each class that are violated by H make up the *evidence* of the negotiation. Formally, let the evidence $E = (v_{C_1}, \ldots, v_{C_n})$ resulting from a negotiation history H be an n-tuple where v_{C_i} is the percentage of preference relations in class C_i that are violated by H. The probability $P(E)$ of observing E is the likelihood of observing this combination of violation percentages in a negotiation.

It may appear to be a difficult task to accurately pin down these probabilities, since they are quite dependent on various aspects of the negotiation, such as length for example. That is, if negotiations are relatively short, evidence with relatively low violation percentages should have higher probabilities than when negotiations are longer, since the more offers there are in the negotiation history, the more violations there will be in the preferences relations. While this is true, it turns out that it is not important. What is important is the relative size of $P(E)$ *given* each class C, or $P(E|C)$. This is because the essential role of $P(E)$ in Bayes' formula is normalization. For example, consider two classes C_1 and C_2. Having $P(E|C_1) = 0.4$ and $P(E|C_2) = 0.2$ would yield the same results as $P(E|C_1) = 0.02$ and $P(E|C_2) = 0.01$. What is important here is the fact that the probability of observing the evidence E is *twice as likely* if C_1 is the true class than if C_2 is the true class (refer to Section 3.1 for more explanation).

4.3. Learning classification probabilities

An effective way to determine class probabilities for a Bayesian classifier is to undergo a learning phase, where the classifier learns the properties of objects for which classes are known. This information is then used to classify objects for which true classes are unknown. We perform prior learning for our classifier in much the same way. Given a set PR of candidate preference relations and the set $\mathscr{C}$ of classes over PR, a preference relation $pr \in PR$ is selected at random. Several possible offers of varying sizes are then selected at random and ordered according to pr. This ordering of offers is then compared with each preference relation in the classes, and the percentage violated in each class by the ordering, as well as the true class, are noted. This process is repeated until sufficient data is obtained, yielding the probability function

$$P((v_{C_1},\ldots,v_{C_N})|C), \quad C \in \mathscr{C}, \quad \mathscr{C} = \{C_1,\ldots,C_N\},$$
$$v_{C_1},\ldots,v_{C_N} \in \Re$$

So $P((v_{C_1},\ldots,v_{C_N})|C)$ denotes the probability of observing evidence $(v_{C_1},\ldots,v_{C_N})$ given that the preference relation is in class C.

4.4. Classification mechanism

The probability of an opponent's preference relation being a member of a class C, given a negotiation history H, is computed as follows. Let $\mathscr{C}$ be a partition over the set of preference relations where each $C \in \mathscr{C}$ has prior probability $P(C)$ of holding an opponent's relation. Note that $P(C)$ may vary for different opponents, based on their business, culture, geography, etc. Let the evidence E be based on the negotiation history (i.e. the sequence of offers given by each party) thus far in H. Then the probability $P(C|E)$ that the opponent's preference relation is in C given the evidence E is given by Bayes' formula.

We summarize the entire classification method as follows. Initially a partial order (as defined in Section 2) exists for the opponent over the set of possible offers. This induces a number of candidate preference relations over the set of possible offers that are consistent with this partial order. Let the evidence E obtained from a negotiation session represent the percentage of initial candidate preference relations that have been violated in each class as a result of the negotiation. That is, E gives the percentage of initial relations in each class that are inconsistent with the orderings observed in the sequence of offers received from the opponent. Let $P(C)$ be the prior probability that each $C \in \mathscr{C}$ includes the opponent's true preference relation as determined in the initial construction of the classes, let $P(E|C)$ be the probability of evidence E being observed given that the opponents relation is a member of C, and let $P(E)$ be overall probability of observing E. The posterior probability $P(C|E)$ that the opponent's preference relation lies in class C given the evidence E is then computed using Eq. (1). This process may continue perhaps until one class has sufficiently high probability to make some reasonable assumptions about the opponent's preferences, and effective strategy development can follow as a result.

5. Multi-object negotiation as multi-issue negotiation

To complete the treatment of this technique described in this paper, we conclude with a section which shows that multi-object negotiation is just a special case of multi-attribute negotiation, and show that the technique can often be applied effectively in the more general model. Let S be the set of objects under negotiation where subsets of S constitute the possible offers. This negotiation model is simply a special case of multi-issue negotiation where there are $|S|$ issues, and each issue has two possible values. More formally, let 2^S be the set of possible offers in a multi-object negotiation of objects S. This corresponds to a multi-issue negotiation over the set S of issues, where each offer $S' \in 2^S$ in the multi-object model corresponds to the offer $S'' \in \{0,1\}^S$ where $s'_i \in S' \Rightarrow s''_i = 1$ in S'' and $s'_i \notin S' \Rightarrow s''_i = 0$ (or other binary values).

It is usually quite clear in multi-object negotiation whether each actor prefers more or less objects. Typically the consumer prefers to receive more (and the producer prefers to give less), but there may also be scenarios where less is preferred, such as negotiating sets of tasks to complete a job, or penalties to follow as a result of a guilty plea. Either way, an actor knows which value (i.e. in or out of the transaction) that the opponent prefers. The same assumption would be needed for the technique discussed in the previous subsection to apply in the general case of multi-issue negotiation. That is, given any issue, each actor's preferences over the set of values for the issue must be common knowledge. Thus a partial order of the opponent's preferences would be known, and this partial order could potentially increase with each offer received. Candidate preference relations could then be eliminated at each step.

Let I be the set of issues under negotiation, with $v(i)$ the set of values for issue $i \in I$. Then according to the above assumption, a total order is known as to the preferences over $v(i)$, given values for the other issues. If the utility independence assumption holds, then this total ordering holds for any values for the other issues. Let $v'_j,\ v''_j \in v(i_j)$

be values for issue i_j. Then offer $s' = \langle v'_1, \ldots, v'_{|I|}\rangle$ is preferred by opponent a at least as much as $s'' = \langle v''_1, \ldots, v''_{|I|}\rangle$ if v'_j is preferred at least as much as v''_j for all $j = 1, \ldots, |I|$. Thus there exists an initial partial ordering.

With the assumption of conceding offers and utility independence, each time an offer s is received, it is assumed that the opponent a's utility function is such that $u_a(s) < u_a(s')$ for all previous offers s'. Then, with the assumption of utility independence, several preferences can be induced:

$$u\left(\left\langle v_1^1, \ldots, v_{|I|}^1 \right\rangle\right) < u\left(\left\langle v_1^2, \ldots, v_{|I|}^2 \right\rangle\right)$$
$$\Rightarrow u\left(\left\langle v_1^3, \ldots, v_{|I|}^3 \right\rangle\right) < u\left(\left\langle v_1^4, \ldots, v_{|I|}^4 \right\rangle\right)$$
$$\text{where } \forall i \; v_i^1 = v_i^2 \Rightarrow v_i^3 = v_i^4, \;\; v_i^1 = v_i^3 \Rightarrow v_i^2 = v_i^4 \tag{3}$$

6. Experimentation

To test the ideas put forth in this paper, we considered the case where S contained five objects v, w, x, y, z and took the point of view of the producer. Thus the idea was to determine the consumer's preference relation. The goals of the experiments were to show that: (1) the classification technique can determine the correct class with high accuracy relatively quickly; and (2) the predicted class will contain preference relations that are relatively close to the opponent's true preference relation.

Initially 120 valid preference relations were randomly generated and divided into four classes using the k-means clustering method. The classifier was then put through a learning phase, where sets of offers were chosen at random and ordered according to a randomly chosen preference relation from a randomly chosen class. The classifier would then learn the probability distributions for the percentage of violations observed in each class as a result of these ordered offers, given the class of the chosen preference relation.

In each run during the experiments, a valid preference relation was randomly generated to represent the opponent's preferences. A negotiation was then simulated, with the opponent following a random concession strategy. That is, the number of ranks to move down the ordered list of offers to find the next offer was determined at random, independently of offers received from our agent. After each offer was received, the classifier updated the probabilities of each class. Certain statistics were then noted. Each negotiation lasted a maximum of eight rounds, and 500 negotiation sessions were run for each experiment.

The goal of the first experiment was to determine the accuracy of the classifier. To demonstrate this, after each round of negotiation in the experiments, the predicted probability of the class that truly contained the opponent's preference relation was noted. The average predicted probability of the correct class after each round is depicted in Fig. 1. This figure displays the results for when the opponent's random concession strategy is to concede either 1 or 2 steps (uniformly) down the ordered list (left figure), as well as the results when the opponent concedes between 1 and 3 steps (right figure). For the purpose of comparison, the probability of the correct class predicted by a simple method, which simply measures the probability of a class by the percentage of the overall preference relations that it holds, is also given in each figure.

Results are statistically significant when the number of offers is greater than 2. So our classification technique works better than simply basing probabilities on the percentage of preference relations in each class. Note that the difference in the performance when the random concession strategy is perturbed is not statistically significant at any point.

Simply predicting the correct class, however, is not quite enough. In practice, when predicting opponent preferences during a negotiation with the intent of building a negotiation strategy, the class with the highest likelihood will be selected and a preference relation will be chosen from that

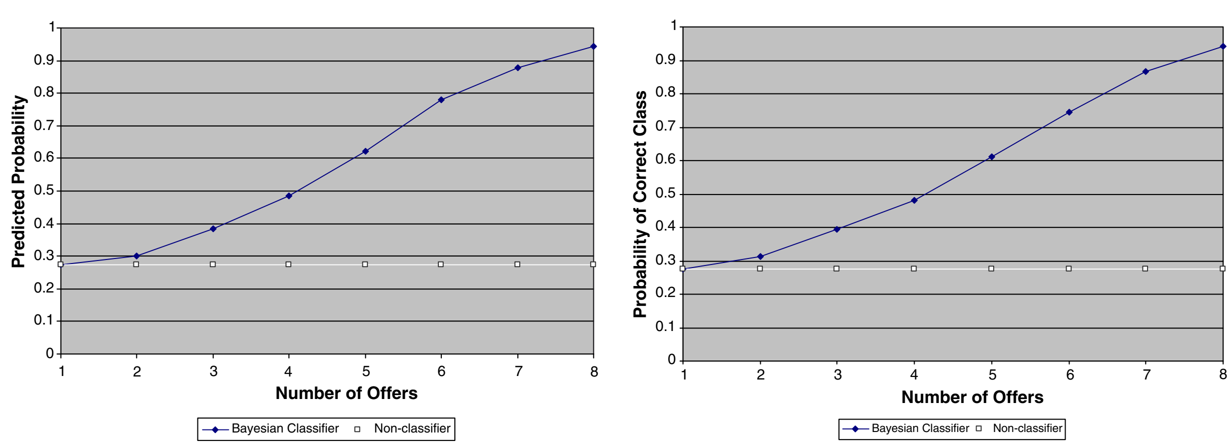

Fig. 1. Average likelihood of selecting the correct class in negotiations where the opponent concedes either 1 or 2 steps down the ordered list for each offer (left), and between 1 and 3 steps (right).

class. It is the hope that any of the preference relations in this class is relatively similar to the true preference relation. The second experiment measures the average distance between the true preference relation and all preference relations in the class with highest probability. The results are depicted in Fig. 2. Again, results are given for the two random concession strategies described above. For the purpose of comparison, the average distance between the true preference relation and a randomly chosen relation, is also given in each figure.

The final experiment compared the performance of our technique with another learning technique that infers opponent preferences by analyzing its responses to our agent's offers. Here the assumption is made that if the opponent follows our offer s with counteroffer s', then the opponent most likely prefers s' over s, since s was deemed to be unacceptable but it was willing to agree on s'. These preferences are accumulated over the course of negotiation, and the opponent is assumed to have some random preference relation that is consistent with these learned preferences. Fig. 3 depicts the results of this test by comparing the average distance between the true preference relation and these consistent preference relations, with the average distance between the true preference relation and the relations in the class deemed most probable by the Bayesian classifier. As usual, results are given for the two opponent concession strategies. Results show that our technique outperforms the simple technique (with statistical significance) after two or more offers for each opponent concession strategy, until the number of offers reaches 7 with the 1–3 concession strategy. This sudden success for the simple technique is due to the small size of the example we are using (five objects, 31 possible offers). With larger concessions, the two sides can become close to achieving a deal after only 5 or 6 offers. Thus these later rounds of offers give valuable information to the learning technique on which of these relatively few borderline offers were preferred by the opponent. Thus in a more realistic negotiation session where there are thousands or millions of possible offers, the technique will not likely perform so

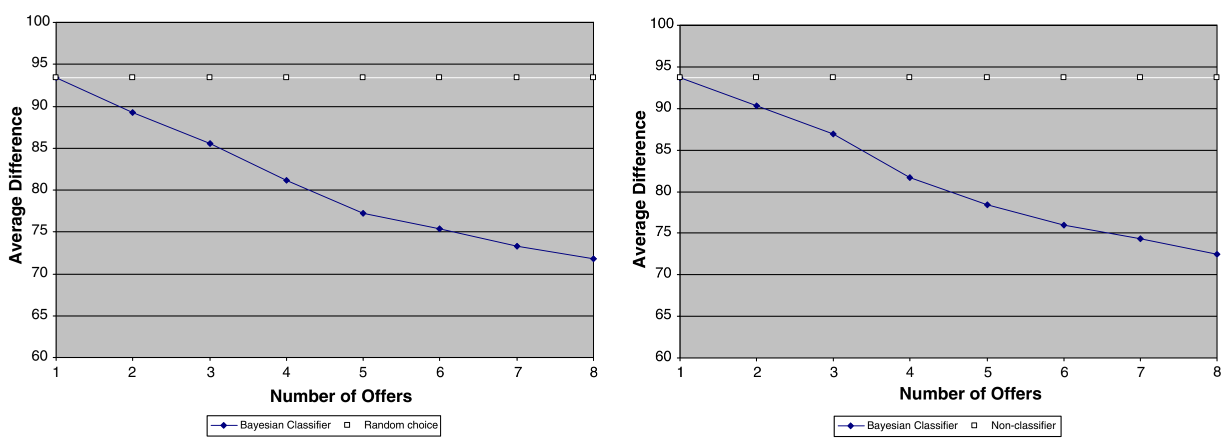

Fig. 2. Average distance between the true preference relation and the preference relations in the class with highest likelihood, where the opponent concedes either 1 or 2 steps down the ordered list for each offer (left), and between 1 and 3 steps (right).

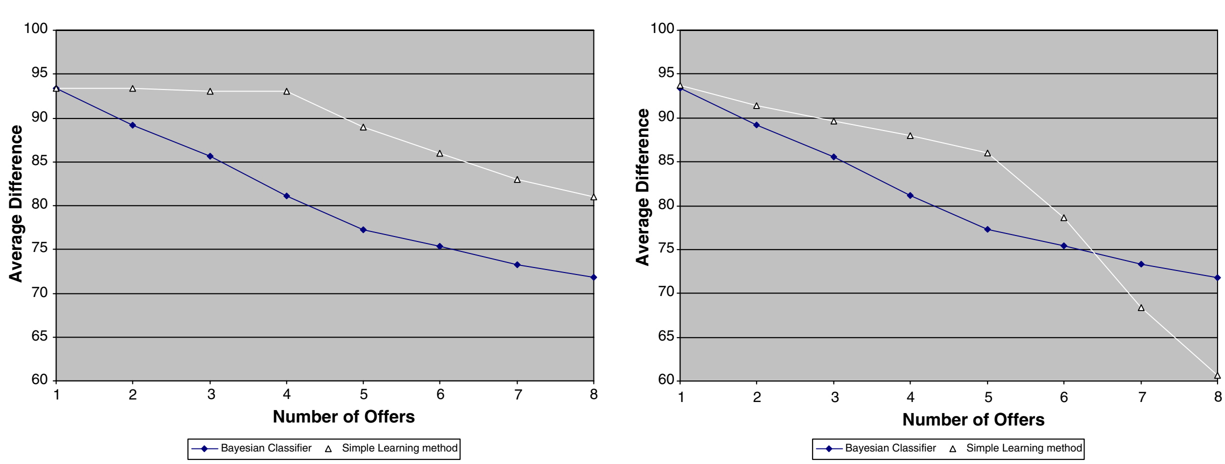

Fig. 3. Performance comparison of the Bayesian classifier and a simple learning technique, where the opponent concedes either 1 or 2 steps down the ordered list for each offer (left), and between 1 and 3 steps (right).

Table 1
Summary of results

Experiment	Opponent concedes 1–2 steps	Opponent concedes 1–3 steps
Likelihood of selecting correct class	Outperformed simple technique that selects class according to size. Stat. sig. for offers > 2	Outperformed simple technique that selects class according to size. Stat. sig. for offers > 2
Average distance from selected class *pr*'s to true *pr*	Outperformed simple technique that selects a relation at random. Stat. sig. for offers > 1	Outperformed simple technique that selects a relation at random. Stat. sig. for offers > 1
Average distance from selected class *pr*'s to true *pr* (learning experiment)	Outperformed simple technique that builds a class of relations that are consistent with learned preferences. Stat. sig. for offers > 1	Outperformed simple technique that builds a class of relations that are consistent with learned preferences. Stat. sig. for 1 < offers < 7

well. However, there is merit to using the approach perhaps when a negotiator feels that a deal is imminent. In this case the technique could be incorporated with our approach by first selecting the most likely class, and then examining the preference relations from that class that are consistent with these learned preferences.

We summarize all of the results in Table 1. The rows describe the three different experiments, while the columns differentiate the two concession strategies used by the opponent. Each cell then indicates how our technique performed in each subexperiment.

7. Conclusions and related work

In this paper we have presented a classification technique for determining close approximations of a negotiation opponent's preference relation over the domain of possible offers. Techniques for learning an opponent's preferences in automated multi-object or multi-issue negotiation framework are vital to the success of a negotiator or negotiation agent, since the formulation of effective negotiation strategies depends on having some indication of the opponent's possible future moves. We define our area of interest as the multi-object negotiation model, and demonstrate a technique for learning the opponent's preferences by observing its behaviour during the negotiation process. The technique uses Bayesian classification to determine a class in which the opponent's preference relation likely fits. While we concede that it might be impossible to determine the total order of the opponent's preferences in practice, we show that narrowing the preference relation to a specific class can be sufficient. Since all preference relations in a class are relatively similar, any relation in the class should be reasonably close to the opponent's relation and a reasonably effective negotiation strategy can be computed. Our restrictive model of multi-object negotiation is then relaxed to show how the technique would work in the more general multi-issue negotiation model. Finally we conclude with a few experiments that demonstrate the effectiveness of our techniques in a particular negotiation scenario. Results show that our classification technique is able to determine the correct class with high accuracy after few offers, and that the difference between the opponent's true preference relation and those in the most probable class as deemed by the Bayesian classifier is relatively small. Our technique is also shown to outperform another technique than learns from the opponent's offers during the negotiation.

Interest in multi-issue automated negotiation has been growing in recent years. Much work in utility elicitation [3,4,10] has recently focused on determining utilities of the user on whose behalf the negotiation agent works, but little has been done to determine the opponent's preferences. Fatima et al. [8,9] break the multi-issue negotiation problem into several negotiations where some issues are settled together and some separately, and determine optimal agendas for those negotiations. Faratin et al. [7] and Coehoorn and Jennings [5] attempt to learn the opponent's preferences and construct counteroffers that are likely to be of interest to the opponent. This is done by making tradeoffs that do not lower the agent's utility, but match more closely with the opponent's previous offers. While this method is likely to allow the negotiators to come to a deal more quickly, it is a cooperative approach and not meant to reveal information about the opponent that can be exploited. Our work differs from this as we provide the negotiation strategy engine with the opponent's preferences. These give insight into the opponent's possible future moves in the negotiation, allowing for a more game-theoretical analysis. This gives the negotiator the ability to make more strategic decisions about what to offer and what to accept, which will help to achieve higher utility. Restificar and Haddawy [15], on the other hand, attempt to gauge the opponent's utility function by paying attention to offers that are rejected and how they are countered. They exploit the fact that if an opponent counters offer a with offer b, then they believe that the opponent's expected utility of offering b (given the chance that they might end up with nothing) is higher than the utility of a for sure. However, they consider only single-issue (money) negotiation. The focus is more on modeling the opponent's attitudes toward risk in such negotiations, since simply determining preferences is straightforward (receivers of money always prefer more to less, while the givers prefer less to more). Similarly, Mudgal and Vassileva [14] examine the idea of learning opponent preferences during a negotiation to in an attempt to determine attitude toward risk, urgency to make a deal and importance of money. Based on previous offers, these factors are modeled using an influence diagram. If subsequent offers differ greatly from the predicted behaviour, the conditional probability distributions are updated. Our work differs greatly from this as we focus on determining the opponent's preferences for outcomes over several attributes (not just money), where preferences are much more difficult to ascertain.

8. Future work

One main area for future work is to improve on the current clustering method for determining an effective initial set of classes. In addition to the clustering literature, we plan to consider using other techniques such as machine learning or collaborative filtering. Collaborative filtering techniques are used to help find similar-minded candidates based on limited information received on agents' preferences. This can help us not only to determine a class structure where preference relations within a structure have maximum similarity, but also to incorporate this structure into the classification mechanism itself. Methods such as this can also be used to help determine where utility independence is likely to hold, and where conditional utilities are likely to exist.

We are also working on the problem of detecting nonconceding offers proposed by the opponent. If a negotiation agent is acting rationally, we can often assume that it prefers to find an acceptable deal over not finding a deal, and thus always tries to make progress toward a deal. If this is the case, a negotiator should only make an offer that is better for itself than its previous offer if it believes that this offer is also better for its opponent (and is thus Pareto superior). We are currently investigating methods that exploit this idea to help predict these occurrences. Detecting these occurrences (and thus disregarding them when making inferences about preferences) should help the performance of our classifier.

Also worthy of deeper examination are the rules for inferring new preferences as described in Section 3.3. The results given in this paper provide supporting evidence that they work well in general, but a deeper understanding of exactly how well they work and under what circumstances is needed in order to realize the full potential.

Another focus for future work is to develop techniques for devising effective negotiation strategies, given the information on likelihoods of classes that can be extracted using our method. This may involve constructing a game tree containing a limited selection of future moves for each actor, where perhaps the moves for the opponent are only those deemed best (from the opponent's point of view) using our beliefs about the opponent's preference relation.

Finally, we also hope to combine the process of determining negotiation strategies with the process of eliciting preferences from the user. If our negotiation engine works on behalf of this user, the user's preferences should be no secret; however they are still difficult to extract nonetheless. Our goal is to interleave the utility elicitation process with the negotiation process, so we can determine which elicitation questions to ask by considering which offers we may give (and receive) as a result, and choose which offers to give by determining what preference information will subsequently be obtained from the user. Only once maximum preference information is extracted from both parties can optimal negotiation strategies be developed.

References

[1] C. Boutilier, M. Goldszmidt, B. Sabata, Sequential auctions for the allocation of resources with complementaries, in: Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI-99), Stockholm, 1999, pp. 527–534.

[2] S. Buffett, K. Jia, S. Liu, B. Spencer, F. Wang, Negotiating exchanges of P3P-labeled information for compensation, Comput. Intell. 20 (4) (2004) 663–677.

[3] S. Buffett, N. Scott, B. Spencer, M.M. Richter, M.W. Fleming, Determining internet users' values for private information, in: Proceedings of the Second Annual Conference on Privacy, Security and Trust (PST04), 2004, pp. 79–88.

[4] U. Chajewska, D. Koller, R. Parr, Making rational decisions using adaptive utility elicitation, in: AAAI-00, Austin, TX, USA, 2000, pp. 363–369.

[5] R.M. Coehoorn, N.R. Jennings, Learning an opponent's preferences to make effective multi-issue negotiation tradeoffs, in: Proceedings of the Sixth International Conference on Electronic Commerce (ICEC2004), Delft, The Netherlands, 2004, pp. 113–120.

[6] L. Cranor, M. Langheinrich, M. Marchiori, M. Presler-Marshall, J. Reagle, The Platform for Privacy Preferences (P3P) 1.0 Specification. http://www.w3.org/tr/p3p/, 16 April 2002. W3C Recommendation.

[7] P. Faratin, C. Sierra, N.R. Jennings, Using similarity criteria to make issue trade-offs in automated negotiations, Artif. Intell. 142 (2002) 205–237.

[8] S.S. Fatima, M. Wooldridge, N.R. Jennings, An agenda-based framework for multi-issue negotiation, Artif. Intell. 152 (1) (2004) 1–45.

[9] S.S. Fatima, M. Wooldridge, N.R. Jennings, Optimal negotiation of multiple issues in incomplete information settings, in: Proceedings of the Third International Conference on Autonomous Agents and Multi-Agent Systems, New York, NY, 2004, pp. 1080–1087.

[10] P. Haddawy, V. Ha, A. Restificar, B. Geisler, J. Miyamoto, Preference elicitation via theory refinement, J. Mach. Learn. Res. 4 (2003) 313–337.

[11] S. Hon-Snir, D. Monderer, A. Sela, A learning approach to auctions, J. Econ. Theory 82 (1998) 65–88.

[12] N.R. Jennings, P. Faratin, A. Lomuscio, S. Parsons, C. Sierra, M. Wooldridge, Automated negotiation: prospects, methods and challenges, Int. J. Group Decis. Negotiation 10 (2) (2001) 199–215.

[13] N.R. Jennings, S. Parsons, C. Sierra, P. Faratin, Automated negotiation, in: Proceedings of the Fifth International Conference on the Practical Application of Intelligent Agents and Multiagent Systems (PAAM-2000), Manchester, UK, 2000, pp. 23–30.

[14] Chhaya Mudgal, Julita Vassileva, Bilateral negotiation with incomplete and uncertain information: a decision-theoretic approach using a model of the opponent, in: Cooperative Information Agents, 2000, pp. 107–118.

[15] A. Restificar, P. Haddawy, Inferring implicit preferences from negotiation actions, in: Proceedings of the International Symposium on Artificial Intelligence and Mathematics, Fort Lauderdale, USA, 2004.