Detecting Opponent Concessions in Multi-Issue Automated Negotiation

Buffett, Scott; Spencer, Bruce; Comeau, Lucas; Fleming, Michael

The 8th International Conference on Electronic Commerce (ICEC'06) [Proceedings], 2006

National Research Council Canada
Conseil national de recherches Canada

Institute for Information Technology
Institut de technologie de l'information

NRC · CNRC

Detecting Opponent Concessions in Multi-Issue Automated Negotiation *

Buffett, S., Spencer, B., Comeau, L., Fleming, M.M.
August 2006

* published at The 8th International Conference on Electronic Commerce (ICEC'06). Fredericton, New Brunswick, Canada. August 14, 2006.
NRC 48552.

Canada

Detecting Opponent Concessions in Multi-Issue Automated Negotiation *

Scott Buffett, Luc Comeau, Bruce Spencer
National Research Council Canada
Institute for Information Technology - e-Business
Fredericton, New Brunswick, Canada, E3B 9W4
{scott.buffett, luc.comeau, bruce.spencer}@nrc.gc.ca

Michael W. Fleming
University of New Brunswick
Fredericton NB, Canada, E3B 5A3
mwf@unb.ca

ABSTRACT

An agent engaged in multi-issue automated negotiation can benefit greatly from learning about its opponent's preferences. Knowledge of the opponent's preferences can help the agent not only to find mutually acceptable agreements more quickly, but also to negotiate deals that are better for the agent in question. In this paper, we describe a new technique for learning about an opponent's preferences by observing its history of offers in a negotiation. Patterns in the similarity between the opponent's offers and our own agent's offers are used to determine the likelihood that the opponent is making a concession at each stage in the negotiation. These probabilities of concession are then used to determine the opponent's most likely preference relation over all offers. Experimental results show that our technique significantly outperforms a previous method that assumes that a negotiation agent will always make concessions during the course of a negotiation.

Categories and Subject Descriptors

I.2.6 [**Artificial Intelligence**]: Learning; J.4 [**Computer Applications**]: Social and Behavioural Science—*Economics*

Keywords

automated negotiation, preferences, learning, mutli-issue negotiation, negotiation protocols, negotiation strategies

1. INTRODUCTION

One of the main challenges in multi-issue automated negotiation lies in the difficulty in learning the preferences of the opposing negotiation agent. Knowledge of the opponent's preferences is key to effective negotiation. Once an agent has a good model of these preferences, it can work towards quick negotiation convergence by making offers that are similar to the opponent's more highly preferred past offers. The agent may also be able to use this information to infer patterns in the negotiator's behaviour, and exploit this knowledge to negotiate better deals for itself. In single-issue negotiation, where typically the object of negotiation is money, the opponent's preference over two candidate offers is obvious: the receiver of money (i.e. the seller) prefers higher values, while the sender of money (i.e. the buyer) prefers lowers values. This is not to say that constructing effective negotiation strategies is trivial; many unknowns still exist, such as the agent's utility function, deadline and reserve price. However, the problem is simplified somewhat in that a negotiator has full knowledge of the opponent's preference relation over the set of outcomes.

In multi-issue or multi-attribute negotiation, where there may be several matters pertaining to the negotiation that need to be resolved, the opponent's preference over outcomes is not so obvious. It may often be the case that for some issues there is no natural preference ordering over the domain. One example of this is colour. Typically, one cannot easily guess another person's or agent's preferences for an attribute such as this. Even if such a natural preference ordering for each attribute is clear, preferences over the combinations of attribute values might not be. These unknowns make it difficult to negotiate effectively, since devising negotiation strategies requires some knowledge of the opponent's future moves. These moves can only reasonably be predicted with some knowledge of the opponent's likes and dislikes.

Previous work has attempted to exploit the information gathered from the opponent's offer history during the negotiation to infer likely preferences over future offers. One example of such efforts is the use of similarity criteria [2, 6, 8]. Here, the opponent's more preferred of two possible outcomes is deemed to be the one that is most similar to the opponent's previous offers. This idea is based on the assumption that the opponent likely has high utility for offers that it submitted, and thus offers that are similar to those should also have high utility, relative to offers that are less similar. A common strategy is then to examine several candidate offers (that have high utility for the offerer), and submit one that is most similar to the opponent's previous offers. One problem with this approach is that all of the previous offers are treated equally. For the approach to be effective, more weight should be given to similarities to highly preferred offers. Another problem lies in the possibility that the opponent could be using a similar strategy, and is subsequently making offers that it believes to be good for the agent in question. Thus some offers might have relatively low utility for the opponent, and should be discounted. Another example of work that considers offer history uses Bayesian classification [4]. These efforts capitalize on the fact that a negotiator often makes concessions over the course of a negotiation in search of a deal. Thus an offer given earlier is likely to have higher utility for the offerer than one given later. This is not always the case however, especially in domains where the two agents may agree on several preferences, since agents can make offers that are more and more enticing to the opponent which are at the same time better and better for themselves.

In this paper, we demonstrate a technique for learning a similarity-maximizing negotiation agent's preferences over its offer history. Such an agent is one that attempts to find agreements relatively quickly by making offers that are similar to its opponent's previous offers, under the belief that such similar offers are likely to have higher utility for the opponent and thus have higher probability of acceptance. Learning agents such as this are less apt to make steady concessions; as preferences are learned, the agent will tend to make offers that are both good for itself and the opposing negotiator. The main focus of this paper is thus to challenge the assumption that agents make constant concessions during negotiations, and to determine the opposing agent's likely preferences over its previous offers. To accomplish this, patterns in the similarity between the agent's offers and our own agent's previous offers are examined. Based on these patterns, the likelihood of concession is determined, which is the probability that the agent has higher utility for the earlier offer than that for the later offer. These probabilities are then used to find the most likely preference relation over its offer history. This preference relation can then be used by the negotiation strategies mentioned above to make more accurate inferences about the opponent's preferences over potential future offers.

The paper is organized as follows. In section 2 we lay out the required background information by formalizing the problem description and our negotiation protocol, as well as offer a deeper discussion of the relevant negotiation strategies from the literature upon which we build. Section 3 then specifies our assumptions and presents our technique for computing concession likelihoods based on these assumptions. We then demonstrate the performance of our technique when used against the previously discussed negotiation strategies, and discuss the results. Section 5 then offers a few conclusions and discusses related work, while section 6 rounds out the paper with a few ideas for future work.

2. BACKGROUND

2.1 Problem Definition

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 an agent a is unsure not only about the reserve limits and deadlines of its opponent a', but also about the opponent's preference ranking of possible offers. Each agent has a utility function $u_a : S \rightarrow \Re$ over the set S of possible outcomes. An agent's utility function therefore induces a preference relation $\succ_a$, where $s \succ_a s' \Leftrightarrow u_a(s) > u_a(s')$. The subscript a will be omitted when it is obvious to which agent we are referring. When $u_a(s) = u_a(s')$, $s \succ_a s'$ is chosen arbitrarily. Each agent 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.

2.2 Convergent Negotiation Protocol

This paper depends upon the existing PrivacyPact [2] protocol, which we now rename to the **CON**nvergent **NEG**otiation protocol (CONNEG). The protocol is designed for two opponents who will negotiate over the two-way exchange of multiple items and/or items with multiple attributes. It defines the set of well formed offers; it also defines the set of allowable offers in the following sense: Depending on the offers previously sent by one agent, some offers are not permitted; for example a seller cannot raise his asking price.

CONNEG defines the set of well-formed offers and the rules for exchange. The set of well-formed offers is defined based on a given set of items to be negotiated and for each item a set of attributes. An offer consists of a selection of attributes for each item. Each partner is assumed to have its own utility, or evaluation, for each possible offer, and these utilities are not shared with the opponent. However, there are mutually known orderings; both participants know something about the opponent's preferences.

The rules of communication as specified by CONNEG are as follows. In all cases it is assumed that if either partner wants to discontinue, the communication port can be closed. Messages representing offers alternate from each participant to the other. All offers remain on the table. An offer from one participant is not allowed if it is mutually known to be worse for the opponent than some previous offer from the participant.

CONNEG has several important properties:

- It is finite. The set of allowable offers decreases monotonically so negotiations cannot continue indefinitely.
- It is complete. If a mutually acceptable offer exists, the negotiation will find it, unless prematurely halted by one of the agents.
- It is reasonable from the perspective of modeling how

cooperative humans negotiate. Offers given in previous rounds continue to be available to the opponent.

Thus the CONNEG protocol enforces and encourages the opponents to make progress. Because of the mutually known partial order over the space of offers; some attributes being offered are mutually known to be of more value to a participant's opponent. For instance, if a warranty is being offered by the seller, a longer warranty is always preferred by the buyer; lower prices are preferred by the buyer and higher prices by the seller; for some set-valued attributes, the recipient is known to prefer getting supersets over subsets (more is better). The CONNEG protocol uses these known orders to enforce progress; a participant's next offer cannot be worse for the other than his previous ones. CONNEG also encourages the opponents to make reasonable progress. Because a participant's given offer stays open for the opponent to accept even after the current exchange, the participant may not rescind that offer. This encourages the negotiation toward cooperative exchanges.

CONNEG was orginally developed for the negotiation of privacy information, where statements in Platform for Privacy Preferences (P3P) [7] are exchanged. Privacy information provides a fruitful area for negotiation. The full combinatorially large space of offers is available. When negotiating in domains where many combinations of attributes do not make sense, the negotiation may have just a fixed set of realizable deals to consider. For instance in the car shopping domain, if you want the power steering option, you must buy the option package that includes power windows and power seats; these cannot be negotiated separately. In privacy, however, all offers are sensible, so the PrivacyPact protocol is important for controlling the negotiation.

2.3 Existing Negotiation Strategies

In this section, we describe relevant existing research on strategies used in bilateral multi-issue negotiations where the parties involved have incomplete information about each other's preferences.

2.3.1 Similarity Maximization

As mentioned earlier, a common approach in multi-issue negotiations is to choose the next offer (from among those that are acceptable to the agent) by selecting an offer that is very similar to offers that have been made in the past by one's opponent. The assumption is that such an offer is more likely to be acceptable to the opponent because it closely matches offers that the opponent clearly believed to be suitable.

Faratin et al. [8] present a trade-off strategy, where multiple issues are traded off against one another in order to increase the social welfare of the system. They use fuzzy similarity to arrive at an estimate of the opponent's preference structure. This is then incorporated into a hill-climbing algorithm to search for a trade-off that is most likely to satisfy the opponent. When it is deemed unlikely that an agreement will be reached with the current set of potential offers induced by these trade-offs, the utility level of offers is reduced and a new set of offers is considered.

Coehoorn and Jennings [6] build on this work by describing a *kernel density estimation*, used to try to estimate an opponent's preferences by looking only at the negotiation history. This approximation is then used to develop an efficient negotiation strategy based on the principle of trade-offs and on the hill-climbing search of Faratin et al.

Buffett et al. [2] investigate three types of negotiation strategies: a "miserly" strategy, in which a negotiator always proposes the offer that has the highest utility for itself, a "cooperative" strategy, in which the next offer chosen is the one that is most similar to the opponent's previous offers, and a "hybrid" approach that combines the miserly and cooperative strategies. The hybrid approach accomplishes the best of both worlds by examining the best n legal offers (in terms of utility for the offerer), and selects the most similar to the opponent's offers from those.

2.3.2 Bayesian Classification

Buffett and Spencer [4] present a Bayesian approach to learning an opponent's preferences during a negotiation. Initially, all the candidate preference relations over the set of offers are divided into classes, where all preference relations in a given class are similar to each other. Based on the offers made by the opponent over the course of the negotiation, an agent uses a Bayesian classifier to estimate the likelihood that the opponent's true preference relation is a member of each class. In using this approach, the agent does not necessarily identify the opponent's exact preference relation; however, it determines the *class* that is most likely to contain this relation. Any relation in this class is likely to be very similar to the opponent's relation, allowing a reasonable negotiation strategy to be developed.

One of the key assumptions made by Buffett and Spencer is the concession assumption. Essentially, this means that offers made by the opponent later in the negotiation are assumed to have lower utility for the opponent than earlier offers, as the opponent attempts to help direct the negotiation toward a mutually acceptable solution. In Section 3 of this paper, we describe a method whereby we attempt to detect offers made by the opponent that are not in fact concessions. In particular, this might occur when the opponent is proposing a new offer because it is similar to our earlier offers; this offer is assumed to be more desirable for us, but it might have higher utility for the opponent as well, thereby violating the concession assumption.

3. PREFERENCE LIKELIHOOD COMPUTATION

3.1 Assumptions

We assume that the agent with whom we negotiate is rational in the sense that it strives towards making deals that will satisfy itself, and will also strive towards making offers that are enticing to the opponent in order to reach an agreement quickly. Specifically, we assume each of the following:

1. Let a be an agent about to make an offer to a' in a negotiation and let s_1 and s_2 be candidate offers. Then a believes that $u_{a'}(s_1) > u_{a'}(s_2)$ and thus s_1 has a higher probability of being accepted (implying that the negotiation will have a lower expected duration

time) if and only if a believes that s_1 is more similar to a''s previous offers than s_2.

2. We are dealing with negotiation agents that 1) care about outcome preferences and 2) care about the length of the negotiation. That is, we assume that agents prefer high outcome preference to low outcome preference and less negotiation time to more negotiation time. We also assume that these two factors are utility independent, which implies that an agent's overall utility of a negotiation outcome and the time to reach that outcome can be represented by a bilinear function of the two.

3. Since we employ the PrivacyPact protocol, there is a danger in submitting offers that yield significant utility to the opponent and at the same time too little utility for oneself. This is because all offers remain "on the table" for the duration of the negotiation. So a large concession that yields high utility to the opponent will make any subsequent offers that have high utility for the sender and low utility for the opponent meaningless, since the opponent can ignore these offers and accept the more preferable offer at any time. Therefore we assume that an agent a has a reserve utility $u_a^r(t)$ on its offers that monotonically decreases as time t increases. That is, at time t, a will not make an offer with utility for itself less than $u_a^r(t)$. The agent will set this value according to how accommodating it chooses to be to the opponent (i.e. how quickly it prefers to have the negotiation settled). The lower $u_a^r(t)$, the more cooperative the agent will be. The agent then examines all legal offers with utility greater than this amount, and offers the one believed to have highest utility for the opponent.

These assumptions yield a set of strategies that includes (but is not limited to) the similarity-based strategies discussed in section 2.3.

3.2 Similarity and Distance

In this section we define two key concepts: 1) the similarity between two offers, and 2) the distance between two preference relations.

3.2.1 Similarity of Offers

Let s and s' be two offers. We adopt a similarity measurement similar to Faratin et al. Specifically, the similarity $sim(s, s')$ of s and s' is defined as the sum of the weighted similarities of the values for each attribute. Let $s(i)$ and $s'(i)$ be the value of the ith attribute for s and s' respectively, and let $h(x, y)$ be the equivalence operator for two attribute values, which returns a value between 0 and 1 depending on the level of equivalence or similarity between x and y. Then the similarity of s and s' is

$$sim(s, s') = \sum_{i=1}^{n} w_i \cdot h(s(i), s'(i)) \tag{1}$$

for the n-attribute case, where w_i are the attribute weights which sum to 1.

3.2.2 Distance between Preference Relations

We measure the likeness of two preference relations in terms of the *distance* between them. Let $\succ$ and $\succ'$ be two such preference relations over the set S of offers. Thus $\succ$ and $\succ'$ specify total orderings over the elements of S. Let $r_\succ : S \to Z$ specify the *rank* of offer in S according to $\succ$, where if s^- and s^+ are the least and most preferred offers respectively, then $r_\succ(s^-) = |S|$ and $r_\succ(s^+) = 1$, and $s \succ s' \Leftrightarrow r_\succ(s) < r_\succ(s')$. The distance $d(\succ, \succ')$ between $\succ$ and $\succ'$ is the total of the differences in rank for each offer $s \in S$:

$$d(\succ, \succ') = \sum_{s \in S} |r_\succ(s) - r'_\succ(s)| \tag{2}$$

3.3 The Likelihood of a Concession

Based on the assumptions about an opponent's negotiation strategy given in the previous section, we determine the likelihood that an opponent's most recent offer in a negotiation is less preferred to the opponent than one of its previous offers. That is, we determine the likelihood that each offer is a concession in utility compared to each of the previous offers. From these probabilities, we then determine the most likely ordering of the opponent's offers according to its personal preference relation.

The central criterion for determining the likelihood of the opponent's preferences lies in the similarity measures of the opponent's offers. By using the assumptions in the previous section, some inferences can be made on the probabilities of opponent concessions by analyzing these similarities. Let the players in the negotiation be the business b and customer c. We take the point of view of the customer. Let a_i be the ith offer from agent $a \in \{b, c\}$ and let $sim_{all}(b_i, j)$ be a function that measures the similarity between b's ith offer and c's first j offers. That is, $sim_{all}(b_i, j)$ is a function of $sim(b_i, c_k)$ for all $k \le j$, such as the average over the set, for example. Let $p_{i,j}$ denote the probability that the business prefers offer b_i over offer b_j, where $i < j$. Theorem 1 first defines when and where concessions are made with certainty:

THEOREM 1. *Let b_i and b_j be two offers from the business where $i < j$. If $sim_{all}(b_i, i-1) < sim_{all}(b_j, i-1)$ then $u_b(b_i) > u_b(b_j)$. Moreover, $u_b(b_t) > u_b(b_j)$ for all $t \le i$.*

Proof: Let $u_b^r(i)$ be the business' reserve utility when b_i was offered. Then, by assumption 3, b_i was believed to be the best for the customer over all offers with utility of at least $u_b^r(i)$. By assumption 1, this implies that b_i maximizes $sim_{all}(b_i, i-1)$. Since $sim_{all}(b_i, i-1) < sim_{all}(b_j, i-1)$, it must be the case that that b_j was not considered and thus $u_b^r(i) > u_b(b_j)$. As for the second part of the theorem, since $u_b^r(t)$ never increases with time, $u_b^r(t) > u_b(b_j)$ for all $t < i$, and thus $u_b(b_t) > u_b(b_j)$ for all $t < i$.

For the general case, when comparing two offers b_i and b_j to determine the likelihood that $u_b(b_i) > u_b(b_j)$, we compare

i	$sim_{all}(b_i, i-1)$	$sim_{all}(b_7, i-1)$
1	-	-
2	0.40	0.46
3	0.50	0.38
4	0.45	0.50
5	0.55	0.48
6	0.58	0.50
7	0.44	0.44

Table 1: Similarity data for Example 1

the *earliest possible entry time* of each offer. The earliest possible entry time for an offer b_i is the earliest time (i.e. offer round) that b_i could have been considered. That is, it is the earliest possible time t such that $u_b(b_i)$ could be greater than or equal to $u_b^r(i)$. This concept is formally defined as follows:

DEFINITION 1 (EARLIEST POSSIBLE ENTRY TIME). *Let b_j be an offer. The earliest possible entry time e_j for b_j is $i+1$ if $sim_{all}(b_i, i-1) < sim_{all}(b_j, i-1)$ (where $i < j$) and $sim_{all}(b_k, k-1) \geq sim_{all}(b_j, k-1)$ for all $i < k \leq j$.*

DEFINITION 2 (ACTUAL ENTRY TIME). *Let b_j be an offer. The actual entry time a_j for b_j is the earliest t such that $u_b^r(t) \leq u_b(b_j)$.*

EXAMPLE 1. *Consider the similarities given in Table 1. The similarity of each offer compared to all of the customer's previous offers at the offer time is given, as well as the similarity of b_7 to the customer's previous offers at those times. At time 2, b_7's similarity was greater than that of b_2, but it was not offered. This means that b_7 had not been considered yet by that time. This was the case at time 4 as well, and is the last time that this occurred. Thus the earliest possible entry time for b_7 is 5. The actual entry time of b_7 could be either 5, 6 or 7.*

Thus if the earliest possible entry time for b_j is t, then $u_b^r(t-1) > u_b(b_j)$ but possibly $u_b^r(t) < u_b(b_j)$. The actual entry time is the unknown time t where the utility for b_j first meets or exceeds $u_b^r(t)$. We determine the likelihood of preference over two offers by comparing the earliest possible entry times. The earlier the earliest possible entry time is, the higher the utility is likely to be. In the case of Theorem 1, not only is the earliest possible entry time for b_i earlier than the earliest possible entry time for b_j, we know that the *actual* entry time for b_i is earlier than the earliest possible entry time for b_j. Because of the monotonically decreasing $u_b^r(t)$ assumption, this implies that b_i is preferred over b_j with certainty. We compute the probability $p_{i,j}$ that $u_b(b_i) < u_b(b_j)$ for the general case as follows.

Let b_i and b_j be offers where $i < j$ with earliest possible entry times e_i and e_j. We assume that the unknown actual entry time for an offer b_k is chosen from a uniform probability distribution over the times e_k to k (inclusive). All joint outcomes for the actual entry times of b_i and b_j are considered, each with probability $prob((a_i, a_j)) = 1/(i - e_i + 1) \times 1/(j - e_j + 1)$ of occurring. For each joint outcome (a_i, a_j) of hypothetical actual entry times the probability $prob(b_i \prec b_j|(a_i, a_j))$ that b_i is preferred over b_j is taken as follows:

$$\begin{array}{ll} prob(b_i \succ b_j|(a_i, a_j)) = 1 & \text{if } a_i < a_j \\ prob(b_i \succ b_j|(a_i, a_j)) = 0.5 & \text{if } a_i = a_j \\ prob(b_i \succ b_j|(a_i, a_j)) = 0 & \text{if } a_i > a_j \end{array}$$

Thus an earlier entry time implies higher utility. When the entry times are the same, no conclusions can be drawn on preference, so the probability that $b_i \prec b_j$ is simply left at 0.5. The probability $p_{i,j}$ is then

$$p_{i,j} = \sum_{(a_i, a_j) \in A_i \times A_j} prob((a_i, a_j)) prob(b_i \succ b_j|(a_i, a_j)) \tag{3}$$

where A_i and A_j are the sets of possible actual entry times for b_i and b_j, respectively.

EXAMPLE 2. *Let b_6 and b_7 be offers with earliest entry times $e_6 = 5$ and $e_7 = 3$. There are two possible actual entry times a_6 for b_6 (5 and 6), and five possible actual entry times a_7 for b_7 (3, 4, 5, 6 and 7). Thus there are ten possible joint outcomes for the two actual entry times, each with 1/10 probability. The ten joint outcomes, along with the probability that $b_6 \succ b_7$ for each case, are given in Table 2.*

(a_6, a_7)	$prob((a_6, a_7))$	$prob(b_6 \succ b_7 \mid (a_6, a_7))$	$prob(b_6 \succ b_7)$
$(5, 3)$	0.1	0	0
$(5, 4)$	0.1	0	0
$(5, 5)$	0.1	0.5	0.05
$(5, 6)$	0.1	1	0.1
$(5, 7)$	0.1	1	0.1
$(6, 3)$	0.1	0	0
$(6, 4)$	0.1	0	0
$(6, 5)$	0.1	0	0
$(6, 6)$	0.1	0.5	0.05
$(6, 7)$	0.1	1	0.1
$p_{i,j}$			0.4

Table 2: Determining the probability $p_{i,j}$ in Example 2

Let m be the number of possible actual entry times for b_i (i.e. $i - e_i + 1$), and let n be the number of possible actual entry times for b_j (i.e. $j - e_j + 1$). Then then probability $p_{i,j}$ can be computed by

$$p_{i,j} = \begin{cases} \dfrac{n^2 + 2(e_j - e_i)n - 1}{2mn} & \text{if } e_j \leq e_i \\ \dfrac{(j - e_i)^2 + j - e_i + m}{2mn} & \text{otherwise} \end{cases} \tag{4}$$

Using these probabilities, the most likely ordering of the offers submitted to the business so far is determined.

One problem with this approach lies in the fact that we might not be able to duplicate the opponent's similarity calculation. Each party has full information on offer histories, but the computation for similarity (e.g. the weights of attribute value similarities) is private. Previous work [6] has had some success with this problem, so a close estimation is attainable. Also, there is a relation between b's similarity assessment of an offer and c's true utility for that offer (since b attempts to compute similarity so that $u_c(b_i) > u_c(b_i') \Leftrightarrow sim_{all}(b_i, i-1) > sim_{all}(b_i', i-1)$. So another idea is to assume that b's similarity computation is exact and use u_c instead.

4. EXPERIMENTATION

To test the technique presented in this paper for determining concession likelihoods, several negotiations over a set of arbitrary objects were simulated. The two players were the *producer*, who will be the giver of the agreed-upon set of objects, and the *consumer*, who will be the receiver of these objects. In each test run, seven objects were up for negotiation, and thus there were $2^7 = 128$ possible offers. Objects were considered desirable, and thus the consumer preferred the presence of an object over the absence of the object, and the opposite was true of the producer. We took the point of view of the producer and attempted to determine the consumer's preferences over its previous offers. That is, the producer used the technique from this paper to determine which of the consumer's offers were concessions, and which were not.

In each run, the opponent (consumer) used a similarity maximization strategy as described above. Thus each time it was the consumer's turn to make an offer, it would look at all offers with utility above some reserve utility, and choose the one that was most similar to all of the producer's previous offers. The reserve utility would then monotonically decrease over the duration of the negotiation. If this reserve utility was always high and thus allowed very few offers to be examined, then as a result the consumer would generally make steady concessions throughout the negotiation. On the other hand, if the reserve utility was relatively low and thus allowed the consumer to consider several possible offers in search of the most similar, then it becomes more likely that the consumer will often move up and down its preference ranking and therefore will make several non-concessions. This situation is the true test of our technique's performance. A utility threshold was selected for each agent, and the negotiation stopped when an agent received an offer whose utility met or exceeded its threshold. In the experiments, thresholds were chosen so that there would be exactly one mutually acceptable deal.

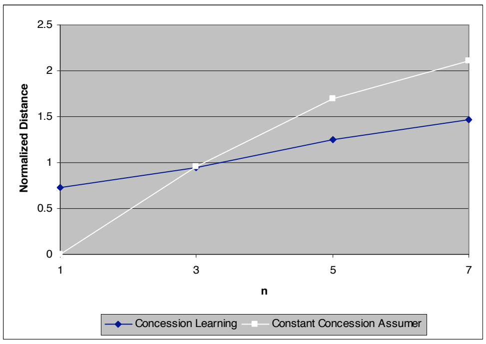

Figure 1: Average distance between the most likely ordering and the true preference ordering for each value for n in the first experiment

n	Concession Assumer	Learning Technique	Win Pct.
3	220	250	53%
5	101	372	79%
7	54	408	88%

Table 3: Number of wins for each technique (ties excluded) in the first experiment

At the end of each run, the total distance between the most likely ordering as determined by our technique and the consumer's true preference ordering over its history of offers was noted. For the sake of comparison, the distance between the ordering as determined by the concession-assuming strategy, where early offers are assumed to be preferred over later offers, and the true preference ordering was noted for each case as well. A total of 500 runs were simulated for each experiment.

In the first experiment, we tested the technique against a consumer that used strategies for reducing its reserve utility similar to that used by Buffett et al [2]. Here, the reserve utility was chosen so that the best n valid offers (according to the consumer's own utility function) were considered and the offer with highest similarity was chosen as the next offer. This value for n was then kept constant throughout the entire negotiation. Tests were performed for $n = 1, 3, 5, 7$. The producer used a similar strategy with $n = 7$ throughout. The average distance between the most likely ordering and the true preference ordering for each value for n is given in Figure 1. The number of times that each method recorded the lower distance out of the 500 cases (ties excluded) are given in Table 3.

When $n = 1$ the concession assuming method performs perfectly, since the negotiation agent will always concede. However, as n increases, the inaccuracy of the concession assuming method increases greatly. This is because the agent has a larger window of offers to examine in search of the most similar offer. As this window gets larger, the likelihood of an offer being a concession in comparison with the previous offer approches 50%. The concession learning technique on the other hand performs very well for higher n, giving a more

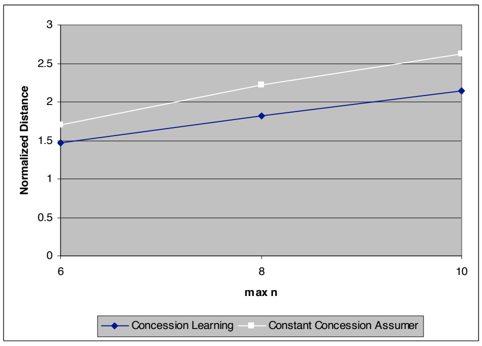

Figure 2: Average distance between the most likely ordering and the true preference ordering for each value for n in the second experiment

n	Concession Assumer	Learning Technique	Win Pct.
6	105	335	76%
8	58	381	90%
10	26	404	94%

Table 4: Number of wins for each technique (ties excluded) in the second experiment

accurate account of the opponent's preferences 88% of the time when the opponent examines seven offers at a time. A t-test shows that that our technique performs better when $n > 3$ at significance level $p < 0.01$.

In the second experiment, we tested the technique against an agent using strategies similar to those of Faratin et al. [8] Here, the agent looks at several desirable offers, each with the same (or close to the same) utility, and selects the one with the highest similarity to its opponent's offers. If this offer is rejected, then the next most similar offer is selected, and so on. If the agent feels like no progress is going to be made, a concession is given and the agent examines a set of offers with slightly lower utility. We emulate this by choosing a value for n, and then decrementing n each time an offer is made. So the agent examines the best n offers, selecting the most similar. If this is rejected, then the agent looks at the best $n - 1$ offers, which is the same as the set from the previous set, minus the offer given. This continues until n reaches one half of its original value. At this time the concession is made and n is restored to its original value. Figure 2 gives the results from experiments using $n = 6, 8, 10$. The number of times that each method recorded the lower distance out of the 500 cases (ties excluded) are given in Table 4.

Again, our technique outperforms the concession-assuming technique, and a t-test shows that that the distance is lower for our technique at significance level $p < 0.01$ for all n tested. Observing the data in Table 4, we see that our technique wins head-to-head almost all of the time, totaling a winning percentage of 94% when $n = 10$. The concession assuming technique failed in this experiment since the opponent's reserve utility is decreased only occasionally (in our tests, every 3, 4 or 5 offers). The rest of the time, offers are chosen from a static set, and thus concessions are made only 50% of the time.

5. CONCLUSIONS AND RELATED WORK

This paper presents a technique for learning an opponent's preferences over the set of its previously made offers. The technique should be most effective when used to learn about a similarity-maximizing negotiation agent. Such an agent is one that attempts to find agreements relatively quickly by making offers that are similar to its opponent's previous offers. To determine these likely preferences, patterns in the similarity between the agent's offers and our own agent's previous offers are examined. Based on these patterns, the likelihood of concession is determined, which is the probability that the agent has higher utility for an early offer than that for a later offer. These probabilities are then used to find the most likely preference relation over its offer history. Once an agent has a good model of these preferences, the agent can work to help the negotiation converge to agreement more quickly by making offers that are similar to the opponent's more highly preferred past offers. Better yet, an agent could also use this information to infer patterns in the negotiator's behaviour, and exploit this knowledge to negotiate better deals for itself.

Results show that our technique significantly outperforms a previous technique that assumes a negotiation agent always makes concessions during the course of a negotiation. When used against a similarity-maximizing agent that examines 7 offers at a time (out of 128 possible offers), our technique finds preference relations that more closely resemble the opponent's true preference relation 88% of the time. When used against an agent that only concedes in the utility of the set of possible offers occasionally, the results once again favour our technique. Here our technique outperforms the concession assumer 94% of the time when the opponent examines up to 10 offers at a time and lowers its minimal offer utility after every 5 offers.

Much work in utility elicitation [1, 3, 5, 11] 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. [9, 10] 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. The above-mentioned work by Faratin et al. [8] and Coehoorn and Jennings [6] 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 determine which of the opponent's offers should be matched and which should be disregarded. Restificar and Haddawy [13], 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 straight-forward (receivers of money always prefer more to less, while the givers prefer less to more). Similarly, Mudgal and Vassileva [12] examine the idea of learning opponent preferences during a negotiation 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.

6. FUTURE WORK

A focus for future work lies in the area of constructing effective negotiation strategies that can exploit these learned preferences. It is one thing to use this information to determine offers that the opponent is more likely to accept, as this can speed up the negotiation process. If an agent's utilities depend greatly on the time taken to ultimately find a deal, then this is worthwhile. For agents that care little about the overall negotiation time and place more weight on the utility of the final outcome, on the other hand, exploiting information on opponent preferences to negotiate better deals is not so straightforward. One idea involves 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. Tools from decision analysis, such as decision trees and Markov decision processes may play a role here as well.

7. REFERENCES

[1] C. Boutilier, Relu Patrascu, Pascal Poupart, and Dale Schuurmans. Regret-based utility elicitation in constraint-based decision problems. In *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence (IJCAI-05)*, pages 929–934, Edinburgh, Scotland, 2005.

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

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

[4] S. Buffett and B. Spencer. Learning opponents' preferences in multi-object automated negotiation. In *Proc. of the 7th International Conference on Electronic Commerce (ICEC2005)*, pages 300–305, Xi'an, China, 2005.

[5] U. Chajewska, D. Koller, and R. Parr. Making rational decisions using adaptive utility elicitation. In *AAAI-00*, pages 363–369, Austin, Texas, USA, 2000.

[6] R. M. Coehoorn and N. R. Jennings. Learning an opponent's preferences to make effective multi-issue negotiation tradeoffs. In *Proc. of the 6th International Conference on Electronic Commerce (ICEC2004)*, pages 113–120, Delft, The Netherlands, 2004.

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

[8] P. Faratin, C. Sierra, and N.R. Jennings. Using similarity criteria to make issue trade-offs in automated negotiations. *Artificial Intelligence*, 142:205–237, 2002.

[9] S.S. Fatima, M. Wooldridge, and N. R. Jennings. An agenda-based framework for multi-issue negotiation. *Artificial Intelligence*, 152(1):1–45, 2004.

[10] S.S. Fatima, M. Wooldridge, and N. R. Jennings. Optimal negotiation of multiple issues in incomplete information settings. In *Proc. 3rd Int. Conf. on Autonomous Agents and Multi-Agent Systems*, pages 1080–1087, New York, NY, 2004.

[11] P. Haddawy, V. Ha, A. Restificar, B. Geisler, and J. Miyamoto. Preference elicitation via theory refinement. *Journal of Machine Learning Research*, 4:313–337, 2003.

[12] C. Mudgal and J. Vassileva. Bilateral negotiation with incomplete and uncertain information: A decision-theoretic approach using a model of the opponent. In *Cooperative Information Agents*, pages 107–118, 2000.

[13] A. Restificar and P. Haddawy. Inferring implicit preferences from negotiation actions. In *Proc. Int'l Symposium on Artificial Intelligence and Mathematics*, Fort Lauderdale, USA, 2004.