NRC Publications Archive Archives des publications du CNRC

Learning Opponents' Preferences in Multi-Object Automated Negotiation

Buffett, Scott; Spencer, Bruce

This publication could be one of several versions: author's original, accepted manuscript or the publisher's version.
/ La version de cette publication peut être l'une des suivantes : la version prépublication de l'auteur, la version acceptée du manuscrit ou la version de l'éditeur.

Publisher's version / Version de l'éditeur:

*Seventh International Conference on Electronic Commerce (ICEC'05)
[Proceedings], 2005*

NRC Publications Archive Record / Notice des Archives des publications du CNRC :
<https://nrc-publications.canada.ca/eng/view/object/?id=c17eadb6-48f7-4b12-a20a-915a7b066a8>
<https://publications-cnrc.canada.ca/fra/voir/objet/?id=c17eadb6-48f7-4b12-a20a-915a7b066a8e>

Access and use of this website and the material on it are subject to the Terms and Conditions set forth at
<https://nrc-publications.canada.ca/eng/copyright>

READ THESE TERMS AND CONDITIONS CAREFULLY BEFORE USING THIS WEBSITE.

L'accès à ce site Web et l'utilisation de son contenu sont assujettis aux conditions présentées dans le site
<https://publications-cnrc.canada.ca/fra/droits>

LISEZ CES CONDITIONS ATTENTIVEMENT AVANT D'UTILISER CE SITE WEB.

Questions? Contact the NRC Publications Archive team at
PublicationsArchive-ArchivesPublications@nrc-cnrc.gc.ca. If you wish to email the authors directly, please see the first page of the publication for their contact information.

Vous avez des questions? Nous pouvons vous aider. Pour communiquer directement avec un auteur, consultez la première page de la revue dans laquelle son article a été publié afin de trouver ses coordonnées. Si vous n'arrivez pas à les repérer, communiquez avec nous à PublicationsArchive-ArchivesPublications@nrc-cnrc.gc.ca.

National Research
Council Canada

Institute for
Information Technology

Conseil national
de recherches Canada

Institut de technologie
de l'information

NRC - CNRC

Learning Opponents' Preferences in Multi-Object Automated Negotiation *

Buffett, S., and Spencer, B.
August 2005

* published at the Seventh International Conference on Electronic Commerce (ICEC'05). August 15-17, 2005. Xi'an, China . NRC 48260.

Copyright 2005 by
National Research Council of Canada

Permission is granted to quote short excerpts and to reproduce figures and tables from this report, provided that the source of such material is fully acknowledged.

Learning Opponents' Preferences in Multi-Object Automated Negotiation

Scott Buffett and Bruce Spencer

National Research Council Canada

Institute for Information Technology - e-Business

46 Dineen Drive, Fredericton, New Brunswick

Canada, E3B 9W4

scott.buffett@nrc.gc.ca, bruce.spencer@nrc.gc.ca

ABSTRACT

We present a classification method for learning an opponent's preferences during a bilateral multi-issue negotiation. Similar candidate preference relations are grouped into classes, and a Bayesian technique is used to determine, for each class, the likelihood that the opponent's true preference relation over the set of offers lies in that class. Evidence used for classification decision-making is obtained by observing the opponents' 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 [6, 9, 10]. Much work has been done recently in this field 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.

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.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ICEC'05, August 15–17, 2005, Xi'an, China.

Copyright 2005 ACM 1-59593-112-0/05/08...\$5.00.

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 opponent preferences typically assumes that several negotiations will take place, over which the preferences will gradually be learned.

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 the total order over the opponent's preferences as well as possible 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. Section 3 presents our classification scheme for opponents' preferences, and we show how the technique can be extended for use in the more general model of multi-issue negotiation. Section 4 then sheds some light on the effectiveness of our technique by describing experimental results, and finally sections 5 and 6 offer conclusions and discuss plans for future work.

2. NEGOTIATION FRAMEWORK

Several issues arise in the definition of an automated negotiation system. Here we assume that negotiations are done

by alternating offers - no other information is exchanged between participants. We partition the set of such negotiations according to the content of offers: one issue or several issues, and for each issue whether it is scalar-, discrete- or set-valued. Note that some issues have a natural orientation; customers want to pay less money, and information holders want to divulge fewer data. This imposes a mutually known partial order over the set of values of that attribute (which our protocol will exploit below). An example of a single-issue negotiation with a scalar-valued issue would be haggling over the price of a single item, where price is a scalar and the owner of the object seeks to increase the price. When money is exchanged, it is (almost) always a scalar value. Sometimes discrete values are used, where the discussion is over the choice of one of several possible colours, for example. Another type of discrete-valued attribute would be a binary one that represents the inclusion or exclusion of a special feature, such as an extended warranty. Often set-valued attributes are used in offers. For example, a carpenter might negotiate to perform a set of tasks for a mechanic in exchange for a set of the mechanics services on the carpenter's car. In this multi-issue case each offer is a pair of sets: the carpenter's services and the mechanic's services. For a different example, if a set of carpentry services and a set of mechanics services are exchanged for money, each offer would be composed of two sets of services and a scalar price.

The Privacy Pact protocol [1] 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. In this section, we formalize the framework of negotiation, and briefly describe the protocol as used in this setting.

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. 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 $s \mu S$. 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 Privacy Pact protocol. However, we assume that these other issues are mutually utility independent with respect to the set of objects, 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, for the remainder of the paper we 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

$u_a : 2^S \rightarrow \mathbb{R}$ for each $a \in \{p, c\}$ is such that $\forall s, s' \subseteq S, s \subseteq s' \Rightarrow u_p(s) \geq u_p(s')$ and $u_c(s) \leq u_c(s')$.

We also assume that each actor has a *break-even point* a_a , where no offer s such that $u_a(s) < \alpha_a u_a(s) < a_a$ is acceptable, and a deadline d_a by which a deal must be made. Utility may also be a function of time, perhaps decreasing to the point where no deal made past time d_a has utility greater than a_a . We maintain the utility independence assumption and assume that an actor's preference relation over S stays constant regardless of time.

Privacy Pact gives the rules of the conversation between the two negotiators. 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 S of negotiation, using the first two messages. 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 each participant to the other, starting first with an offer from the consumer. Each offer consists of a single some $s \mu S$. 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 negotiation ends with this final offer being the agreement.

3. CLASSIFICATION OF PREFERENCE RELATIONS

3.1 Identifying Classes

The technique for learning opponents' preference relations presented in this paper is based on Bayesian classification. Given a mutually exclusive and exhaustive set of classes, the idea is to determine the likelihood of the opponent's preference relation being a member of each class. 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.

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 the set of offers received from the opponent. Equally important in the task of examining the evidence is also the order in which those offers were received. Typically during a negotiation session, an actor has a tendency to make concessions as time passes. So if an actor a follows offer s_i with s_{i+1} , then we can assume with a high probability that $u_a(s_i) > u_a(s_{i+1})$. While this concession assumption might seem too strong, consider that the Privacy Pact protocol dictates that all previous offers remain on the table. Thus 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.

If s_i and s_{i+1} intersect, for example $s_i = \{x, y\}$ and $s_{i+1} = \{x, z\}$, then one may also conclude with high probability that $u_a(\{y\}) > u_a(\{z\})$. Note that in general it is unreasonable to assume utility independence of the elements of S since some may combine well with others for various reasons. Thus $u_a(\{x, y\}) > u_a(\{x, z\})$ does not necessarily imply that $u_a(\{y\}) > u_a(\{z\})$, since the presence of x may affect the preferences of y and z . However, if these dependencies are relatively few, then one can at least say with high probability that $u_a(\{y\}) > u_a(\{z\})$.

The probability of an opponent's preference relation being a member of a class C given a negotiation N is computed as follows. Let C be a partition over the set of preference relations, each $C \in \mathcal{C}$ with 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 the negotiation history (i.e. the sequence of offers given by each party) thus far in N . Then the probability $P(C|E)$ that the opponent's preference relation is in C given the evidence E is given by Bayes' Rule:

$$P(C|E) = \frac{P(E|C) \cdot P(C)}{P(E)} \quad (1)$$

Thus for each $C \in \mathcal{C}$, one must be able to determine the probability $P(E|C)$ of a sequence of offers given that the opponent's preference relation is a member of C . These probabilities can be estimated by expert analysis or by examination of historic data from previous negotiations. The key is that it does not require that we have experience with this particular opponent. We only need population data. The offers we receive from the opponent then give us insight on the part of the population with which we are dealing.

Determining the initial classes, as well as the prior probabilities of each, is a separate problem. Techniques from the area of data mining, such as collaborative filtering, may work well here. This is however outside of the scope of the current paper, and we defer such research to future work. For this paper we use a simple definition of the initial set of classes, which we detail in section 4.

3.2 Classification Mechanism

The key to correct classification is to accurately compute $P(E|C)$ for any evidence E obtained from the negotiation history, given any class C . These probabilities will be influenced by any assumptions that might be made about the opponent (i.e. level of cooperativeness, concession rate, etc), and also on the class-building scheme employed. One would also believe that the effectiveness of a classification mechanism greatly depends on the class structure. Thus the possibilities for classification techniques are seemingly unlimited. We develop one such technique here, and demonstrate its effectiveness in the following section.

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. Moreover, we make the added assumption of utility independence between items. That is, for all s_1, s_2, s', s'' subsets of S and opponent a ,

$$u_a(s' \cup s_1) < u_a(s' \cup s_2) \Rightarrow u_a(s'' \cup s_1) < u_a(s'' \cup s_2) \quad (2)$$

Note that neither of these assumptions need actually be true of the opponent in practice. They are only used for classification purposes. As long as they are usually true, the classification mechanism will work reasonably well. If one could utilize other knowledge about the opponent, such as concession rate/frequency or special cases where utility independence does not hold between objects, then the mechanism can be made to work even better.

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. Each time a new offer is received, the partial order has the potential to be more specifically defined. This reduces the set of consistent preference relations. If sufficient offers are received, the total order could be defined, thus leaving only one possible preference relation. This is not likely to happen, so we turn to the problem of determining the most likely class.

We define *evidence* E obtained from a negotiation history as the number of initial candidate preference relations that have been eliminated as a result of the negotiation. That is, E gives the number of initial relations that are inconsistent with the orderings learned from the sequence of offers received from the opponent. More formally, $E = (X, N)$ is a tuple denoting the number X of initial candidate preference relations eliminated as a result of the sequence of N offers received. Thus the probability $P(E|C) = P((X, N)|C)$ is the probability of having X relations eliminated after N offers given that the opponent's true preference relation is a member of class C .

Probability distributions over various types of evidence for a given class can be determined by either examining population data on past negotiations, by analytic methods, or even by constructing realistic models with reasonable negotiation strategies and simulating data. Once these distributions are constructed (or can at least be computed on demand), classification of preference relations can commence.

Let $P(C)$ be the prior probability that each $C \in \mathcal{C}$ 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 $E = (X, N)$ being observed given that the opponent's relation is a member of C where X is the number of eliminations after N offers from the opponent, and let $P(R(C, N))$ be the set of preference relations from class C that are consistent with the evidence received through N rounds. Initially, $P(R(C, 0))$ is the set of relations from C that are consistent with the initial partial order as defined in section 2. The evidence after N rounds of negotiations is then $E = (X, N)$ where $X = |P(R(C, 0))| - |P(R(C, N))|$. The probability $P(E|C)$ is then computed for each class, and the posterior probability $P(C|E)$ that the opponent's preference relation lies in class C given the evidence E is then computed using equation 1. This process may continue perhaps until one class has sufficiently high enough probability to make some

reasonable assumptions about the opponents preferences, and effective strategy development can follow as a result.

3.3 Multi-object Negotiation as Multi-issue Negotiation

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

With the assumption of conceding offers, 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^0 . Then, with the assumption of utility independence, several preferences can be induced with the analog of equation 2:

$$u(\langle v^1_1, \dots, v^1_{|I} \rangle) < u(\langle v^2_1, \dots, v^2_{|I} \rangle) \Rightarrow u(\langle v^3_1, \dots, v^3_{|I} \rangle) < u(\langle v^4_1, \dots, v^4_{|I} \rangle)$$

$$\text{Where } \forall i \quad v^1_i = v^2_i = v^3_i = v^4_i, \\ v^1_i = v^3_i = v^2_i = v^4_i$$

(3)

4. EXPERIMENTATION

To test the ideas put forth in this paper, we considered the

case where S contained 5 objects v, w, x, y, z and took the point of view of the producer. Thus the goal was to determine the consumer's preference relation. The candidate preference relations were divided into four simple classes:

1. The class of all preference relations where any two objects are preferred over any one.
2. The class of all preference relations where one particular object is preferred over any two other objects. Any pair is preferred over any one other than this.
3. The class of all preference relations where one particular object is preferred over any three other objects. Any pair is preferred over any one other than this.
4. The class of all preference relations where one particular object is preferred over any three other objects, and another object is preferred over any pair not including the first one mentioned. Any pair is preferred over any one other than these two.

Natural consequences follow from each class description. For example, in all class 1 preference relations any three objects are preferred over any one, as well as any four, etc. The idea behind this particular scheme is that it facilitates classification when preferences for different-sized offers are known. For example, if we know that vw is preferred to wxy , then we can immediately rule out class 1. Also with this class structure, given that the preference relation is (or is likely to be) in a particular candidate class, once all preferences over objects in S are known, the entire preference relation over all subsets of S can be inferred. While it is clear that this is not the most effective classification scheme, the goal is to show that if the technique works for our experiments, then it should perform very well when more sophisticated classification techniques are employed.

In these experiments, the opponent's strategy consisted of a random concession technique. Offers submitted by the opponent had no dependence on the offers it received. One of the four classes was chosen randomly, followed by a random selection of a preference relation from the class. All

Figure 1: Average likelihood of selecting the correct class

$2^5 - 1 = 31$ offers (not including the empty set) were then ranked from best to worst according to the preference relation. Starting with the most favorable offer, which is always S since the consumer always prefers to have all objects, concessions were made by causing the opponent to move randomly 1, 2 or 3 ranks down the list to find the next offer. All negotiations lasted 9 rounds, with various statistics noted as to the accuracy of our method in relation to the true preference relation after each round. Probabilities $P(E|C)$ were estimated by initially running experiments and recording the distributions for the number of eliminations after each round, given each class.

The first experiment was designed with the goal of showing how accurately the opponent was classified after each round. Figure 1 shows the average likelihood of the correct class predicted by our technique. For the purpose of comparison, a guessing technique that makes use of no information would simply predict the correct class with an average of 25% likelihood, since there are 4 classes.

The purpose of the second experiment was to show how often our technique could find the true preference relation compared to a simple method. Given the set of remaining relations that are consistent with the sequence of offers received thus far, the simple method simply chose one at random. Our method, on the other hand, first chose a class randomly according to the likelihood of each class, and then chose a member of that class at random. Figure 2 summarizes the results. Note that the difference in the average achieved by our method compared with that for the simple technique was statistically significant when the number of offers was ≥ 3 .

Finally, the third experiment measures the performance of the technique with respect to the main goal of our efforts. In practice, we do not expect to learn the opponent's preference

measured the weighted average of the similarity of all remaining relations as compared to the true relation. The average is weighted by the likelihood of each class. Similarity between to preference relations was measured as the number of preference pairs that the two relations had in common. Since there are 31 possible offers, similarity was measured as a score out of ${}_{31}C_2 = 465$ (i.e. there are 465 pairs). For comparison, the unweighted average similarity over all remaining relations was noted as well, which measures the performance of the simple technique. Figure 3 shows the results. Note that the difference in the average achieved by our method compared with that for the simple technique was statistically significant when the number of offers was ≥ 6 .

5. 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

Figure 2: Average likelihood of selecting the correct preference relation

relation exactly. Instead, we hope to narrow it down to a class of preference relations where any in the class should be quite similar to the true relation. In this experiment we

Figure 3: Weighted average similarity to true preference relation

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.

Interest in multi-issue automated negotiation has been growing quickly in recent years. Much work in utility elicitation [2, 3, 8] 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. [6, 7] 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. Fatima et al. [5] and Coehoorn and Jennings [4] 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. Restificar and Haddawy [11], 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.

6. FUTURE WORK

One main area for future work is the investigation of technologies for determining effective classes, given a set of pre-existing data. This will likely involve the use of machine learning techniques such as 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 determine a class structure where preferences relations within a structure have maximum similarity, but can also be incorporated into the classification mechanism itself.

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.

7. REFERENCES

- [1] 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.
- [2] 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.
- [3] U. Chajewska, D. Koller, and R. Parr. Making rational decisions using adaptive utility elicitation. In *AAAI-00*, pages 363–369, Austin, Texas, USA, 2000.
- [4] 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.
- [5] 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.
- [6] S.S. Fatima, M. Wooldridge, and N. R. Jennings. An agenda-based framework for multi-issue negotiation. *Artificial Intelligence*, 152(1):1–45, 2004.
- [7] 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.
- [8] 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.
- [9] N. R. Jennings, P. Faratin, A. R. Lomuscio, S. Parsons, C. Sierra, and M. Wooldridge. Automated negotiation: prospects, methods and challenges. *Int. J. of Group Decision and Negotiation*, 10(2):199–215, 2001.
- [10] N. R. Jennings, S. Parsons, C. Sierra, and P. Faratin. Automated negotiation. In *5th International Conference on the Practical Application of Intelligent Agents and Multiagent Systems (PAAM-2000)*, pages 23–30, Manchester, UK, 2000.
- [11] 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.