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