Using Opponent Models for Efficient Negotiation (Extended Abstract)

Koen Hindriks
Man-Machine Interaction Group
Delft University of Technology
Mekelweg 4, Delft, The Netherlands
+31.15.2781315
k.v.hindriks@tudelft.nl

Catholijn Jonker
Man-Machine Interaction Group
Delft University of Technology
Mekelweg 4, Delft, The Netherlands
+31.15.2781315
c.m.jonker@tudelft.nl

Dmytro Tykhonov
Man-Machine Interaction Group
Delft University of Technology
Mekelweg 4, Delft, The Netherlands
+31.15.2783737
d.tykhonov@tudelft.nl

ABSTRACT

Information about the opponent is essential to improve automated negotiation strategies for bilateral multi-issue negotiation. In this paper we propose a negotiation strategy that combines a Bayesian technique to learn the preferences of an opponent during bidding and a Tit-for-Tat-like strategy to avoid exploitation by the opponent. The learned opponent model is used to achieve two important goals in negotiation. It may be used to increase the efficiency of negotiation by searching for Pareto optimal bids and to avoid exploitation by making moves that mirror the move of the other party. The performance of the proposed negotiation strategy is analyzed in a tournament setup.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence – *intelligent agents, multi-agent systems,*

General Terms

Algorithms, Performance, Economics, Experimentation, Theory.

Keywords

Automated Multi-Issue Negotiation, Opponent Modelling, Bayesian Learning, Negotiation Strategy, Tit-for-Tat.

1. INTRODUCTION

In bilateral negotiation, two parties aim at reaching a joint agreement. They do so by exchanging various offers or bids using e.g. an alternating offers protocol called the "negotiation dance" in [4]. In reaching such an agreement both parties usually aim to satisfy their own interests as best as possible, but have to take their opponent's preferences into account as well to reach an agreement at all. This is complicated, by the fact that negotiating parties are generally not willing to reveal their preferences in order to avoid exploitation. As a result, both parties have incomplete information which makes it hard to decide on a good negotiation move and hard to reach an optimal agreement.

One way to approach the problem of incomplete information is to learn an opponent's preferences given the negotiation moves that an opponent makes during the negotiation. For example, negotiating agents can obtain a good approximation of an opponent's preferences in a single-session negotiation by studying the offers made by the opponent [3]. The approach of [3] focuses on obtaining a correct model of the opponent's preferences in the shortest time possible from as few bids as possible.

In this paper, we propose a negotiation strategy that uses a model of the opponent's preferences not only to increase the efficiency of the negotiated agreement but also to avoid exploitation by the other party in a sophisticated manner. This paper shows that two important goals in any negotiation can be realized when a reasonable estimate of the preferences of an opponent is available. For that purpose we combine the Bayesian learning technique as proposed in [3] with a Tit-for-Tat-like tactic, see e.g., [1], and the classification of negotiation moves as described in [2]. The opponent profile together with the classification scheme is used to develop a Tit-for-Tat Bayesian negotiation strategy, that is sophisticated in three ways.

First, its bidding can be understood by the opponent as signalling whether a move is appreciated or not (which is not as easy as it seems). Second, it is a friendly strategy that does not punish the opponent for making a move that can be understood as an honest mistake. Finally, the proposed strategy is based on a rationality assumption, i.e., that an opponent would tend to accept more preferred offers over less preferred. In line with this assumption, the strategy searches for Pareto efficient offers, i.e., offers that cannot be improved for both parties simultaneously.

2. NEGOTIATION STRATEGY

The ultimate goal of a negotiation strategy is to reach the most beneficial deal possible. Negotiation is, however, a joint decision requiring acceptance of an offer by the other party. Any negotiator thus faces a dilemma involving two conflicting goals in a negotiation: how to maximize the utility of an agreement for the agent itself and how to maximize the chance of acceptance by the opponent at the same time. To achieve the first goal the strategy proposed is based on the Tit-for-Tat-like tactic that has proved to be efficient in co-operative problem-solving negotiation settings [1]. To achieve the second goal a model of the opponent's preferences that is learned during a negotiation session is used.

The negotiation strategy assumes that a preference profile can be defined as a function of the evaluation functions associated with the individual issues. More specifically, the utility function has a linear additive structure. We assume that a learning technique used to model the opponent preferences tries to reconstruct original utility function representing the opponent preferences with an approximate function. The authors of [3] propose an algorithm to model opponent preferences based on the Bayesian learning technique. The algorithm learns the probability distribution over a set of hypotheses about evaluation functions and weights of the issues. The probability distribution defined over the set of hypothesis represents the agent's belief about the opponent's preferences. Structural assumptions about the

evaluation functions and weights decrease the number of parameters to be learned and simplify the learning task. During a negotiation when a new bid is received from the opponent the probability of each hypothesis is updated using Bayes' rule.

The availability of information about an opponent's preferences enables an agent to make judgments about the opponent's moves. A classification of negotiation moves proposed in [2] gives an interesting insight into the dynamic properties of agent negotiation behaviour. The classification of the negotiation moves allowed us to develop a sophisticated Tit-for-Tat tactic. The basic idea of Tit-for-Tat in multi-issue negotiation is to respond to an opponent move with a symmetrical one, as depicted in Figure 1. All rational negotiation strategies known to us try to make concession moves at some points during the negotiation. The most reasonable response to a concession move would be a concession move of approximately the same concession size. This is called "mirroring" the move of the opponent.

Mirroring simply in this manner would imply that an unfortunate move of the opponent would be answered with an unfortunate step. However, it is not rational to consciously make unfortunate steps. Therefore, we conclude that the pure tactic by mirroring the opponent moves is too simplistic. Instead we use an approximation of the Pareto frontier computed using the learned opponent model and the agent's own preference profile to add an additional step. The strategy thus developed is called Bayesian Tit-for-Tat.

In this step after matching the opponent move according to the Tit-for-Tat tactic the Bayesian Tit-for-Tat strategy searches for a bid on the approximated Pareto frontier that is on the same iso-curve with the matched bid, see Figure 1.


Figure 1 -Bayesian Tit-for-Tat Strategy

In short, we constructed the Bayesian Tit-for-Tat strategy on the basis of the assumption that by maximizing the opponent's utility in every offer, we increase the chance of acceptance. Therefore, if after mirroring the opponent's move the efficiency of our next move can be increased by selecting an equivalent offer (with respect to the agent's preference profile) on the Pareto frontier our strategy will choose to make that offer. The agent accepts a bid from its opponent when the utility of that bid is higher than the utility of its own last bid or the utility of the bid it would otherwise propose next.

Note, that learning the opponent preference is not perfect and therefore, classification of the negotiation moves can be imprecise. As a result, the adequacy of our strategy can only be established experimentally in a tournament as proposed in [2]. As discussed, the main objective associated with a negotiation strategy is to gain the best agreement possible in a negotiation. Utility of an agreement, therefore, measures the efficiency of a strategy. For every negotiation domain and preference profile the

utility of agreements achieved by a strategy were averaged over all opponent strategies in the tournament. We assume that an efficient negotiation strategy should perform better than the Zero Intelligence (ZI) strategy [2]. Therefore, we calculate the percentage of the utility increase compared to the utility of the ZI strategy (see Table 1).

Table 1 - Increase in utility relative to the ZI strategy

Negotiation Domain	Negotiation Strategy			
	ABMP	Trade-Off	Bayesian Smart	Bayesian Tit-for-Tat
Car	16%	12%	13%	14%
Party domain	13%	9%	13%	14%
Service Oriented	14%	17%	25%	38%
AMPO vs City	10%	13%	14%	20%
Employment contr.	11%	40%	44%	47%

The results show that on all domains the Bayesian Tit-for-Tat strategy performs better than all other strategies currently available in the negotiation repository, except the 2nd hand car negotiation domain where the ABMP strategy is very efficient.

3. CONCLUSIONS

The Bayesian Tit-for-Tat strategy for closed multi-issue negotiation as proposed in this paper shows that two important goals in any negotiation can be realized when a reasonable estimate of the preferences of an opponent is available: to increase the efficiency of the negotiated agreement and to avoid exploitation by the other party. Our strategy combines three fundamental techniques: Bayesian learning, Tit-for-Tat, and the classification of negotiation moves as developed for the analysis of negotiation strategies [2]. The Bayesian technique as proposed in [3] is used to learn the opponent's preferences during the negotiation. Tit-for-Tat, e.g., [1], is applied to avoid exploitation by a form of mirroring of the bids of the opponent. The opponent profile together with the classification scheme is used to develop a negotiation strategy that is sophisticated in three ways.

Firstly, its bidding can be understood by the opponent as signalling whether a move is appreciated or not (Tit-for-Tat). It forms the basis of our aim to be *transparent* to the opponent. Secondly, it is a *friendly* strategy in that it does not punish the opponent for making a move that can be understood as an honest mistake. Third, using the learned opponent model our strategy allows an agent to propose nearly Pareto optimal offers using the learned opponent preferences to approximate the Pareto frontier.

The performance of the proposed Bayesian Tit-for-Tat negotiation strategy has been analyzed in a tournament setup, using domains of different characteristics and a number of existing negotiation agents. The results show the Bayesian Tit-for-Tat strategy realizes a significant relative increase in utility with respect to the other strategies that were part of the tournament setup, including the Zero Intelligence, ABMP, Trade-Off, and Bayesian agents.

4. REFERENCES

[1] Axelrod, R. 1984. *The Evolution of Cooperation*. Basic Books, Inc., Publishers, New York, USA.

[2] Hindriks, K., Jonker, C. M., Tykhonov, D. 2007. Analysis of Negotiation Dynamics In: *Proc. of CIA 2007*, Delft, The Netherlands, Springer-Verlag, LNAI4676, pp. 27-35.

[3] Hindriks, K., and Tykhonov, D. 2008. "Opponent Modelling in Automated Multi-Issue Negotiation", AAMAS'08.

[4] Raiffa, H. 1982. *The Art and Science of Negotiation*, Harvard University Press.