Opponent Modeling with Information Adaptation (OMIA) in Automated Negotiations

Yuchen Wang(✉), Fenghui Ren, and Minjie Zhang

School of Computing and Information Technology,
University of Wollongong, Wollongong, NSW 2522, Australia
yw808@uowmail.edu.au, {fren,minjie}@uow.edu.au

Abstract. Opponent modeling is an important technique in automated negotiations. Many of the existing opponent modeling methods are focusing on predicting the opponent's private information to improve the agent's benefits. However, these modeling methods overlook an ability to improve the negotiation outcomes by adapting to different types of private information about the opponent when they are available beforehand. This availability may be provided by some prediction algorithms, or be prior knowledge of the agent. In this paper, we name the above ability as Information Adaptation, and propose a novel Opponent Modeling method with Information Adaptation (OMIA). Specifically, the future concessions of the opponent will firstly be learned based on the opponent's historical offers. Then, an expected utility calculation function is introduced to adaptively guide the agent's negotiation strategy by considering the availability and value of the opponent's private information. The experimental results show that OMIA can adapt to different types of information, helping the agent reach agreements with the opponent and achieve higher utility values comparing to those which lack the information adaptation ability.

Keywords: Automated negotiations · Opponent modeling
Information adaptation

1 Introduction

Negotiation is an important activity between people or parties who discuss issues intending to reach agreement. Negotiation often involves significant cost in human resources and time. As a result of this, automated negotiation techniques have attracted increasing attentions during the last two decades [10]. The benefits include resource-saving in manpower and time [4], avoiding social confrontation [3] and automatic bargains in e-markets [15].

One key challenge in automated negotiations is to reach beneficial negotiation results when private information about the opponent is unknown. The private information is in contrast to public information known by all negotiators.

G. Sukthankar and J. A. Rodriguez-Aguilar (Eds.): AAMAS 2017 Best Papers,
LNAI 10642, pp. 21–35, 2017.
https://doi.org/10.1007/978-3-319-71682-4_2

For instance, the public information includes the maximum negotiation time, the historical offers exchanged by negotiators, etc., while the private information contains a negotiator's personal reservation value, deadline, etc. Obviously, sharing this private information between negotiators is not applicable. To tackle this problem, researchers have put effort into opponent modeling techniques [2]. These techniques mostly give agents an ability of exploring the opponent's private information and adapting to the opponent's negotiation behavior in order to achieve satisfactory negotiation outcomes.

In current literature, a number of opponent modeling methods have been developed by employing different learning methods, such as Bayesian learning [18], Non-linear regression [14], Kernel density estimation [5], and Artificial neural networks [12]. Among these methods, four types of opponent's private information are commonly selected as their learning goals, which are reservation value [17], negotiation deadline [9], bidding strategy [16] and offer acceptance possibility [13]. However, these approaches lack an ability of information adaptation.

The information adaptation indicates an ability to make better negotiation decisions when some types of the opponent's private information are available. This availability may be provided by some prediction algorithms, or be prior knowledge of the agent. For example, when the opponent's reservation value is available, the agent should adapt its negotiation strategy and try to make agreements close to this reservation value. Lacking the ability of information adaptation means losing a potential behavior guidance toward different types of available information, thus the negotiation outcomes would be negatively affected.

In order to give agents such an information adaptation ability, in this paper, we propose a novel opponent modeling method called Opponent Modeling with Information Adaptation (OMIA). Traditionally, an opponent modeling method could guide the agent's negotiation strategy adapting to the historical offers of the opponent. In this paper, OMIA should not only adapt to the historical offers, but also adapt to different types of private information both when they are available and unavailable. The types of private information considered in this paper are reservation value, deadline, bidding strategy and acceptance possibility due to their importance in automated negotiations [1].

To establish OMIA, the major challenge is how to create one opponent model that can simultaneously adapt to the five types of information based on their availability (historical offers are always available) and values. Our idea is to establish OMIA from a probability point of view. In particular, these probability distributions are going to be estimated: (i) the probability distributions of the opponent offers' utility in future time; (ii) the probability that the opponent accepts a particular offer; (iii) the probability that the opponent quits in particular future time. The process of OMIA is as follows. First, OMIA predicts the future concessions (i) of the opponent using the historical offers. Then, an expected utility calculation function is introduced to estimate (ii) and (iii). Then, based on (i), (ii) and (iii), this function calculates the expected utility the agent will gain when it takes different concession strategies. The maximum expected utility determines the best concession strategy that the agent should take. The availability and values of the private information will influence the results of all probability distributions, affecting the value of the expected utility, and thus guiding the concession strategy of the agent.

The merits of OMIA are: (a) OMIA could adapt to the behavior of the opponent only using the historical offers. (b) OMIA could adapt to different types of private information based on their availability and values, and the agent can choose the types of information to adapt to. (c) OMIA makes little assumptions about the opponent (e.g. bidding strategy, utility function, etc.), making itself a highly robust model.

The remainder of this paper is organized as follows. Section 2 describes the general negotiation setting. Section 3 introduces the proposed OMIA. Section 4 describes how OMIA adapts to various types of information. Section 5 demonstrates the experimental results. Related work is presented in Sect. 6, and Sect. 7 makes a conclusion.

2 Negotiation Setting

In this paper, we study the bilateral single-issue automated negotiation, which consists of two agents negotiating over a single issue. The alternating offers protocol [11] is employed in this paper where two agents exchange offers in turns until one agent accepts an offer or reaches its deadline. The time can either be a continuous or discrete variable. In this paper, we use the discrete time setting as we are focusing on adapting to various types of input information and currently do not take computational cost into account. A monotonic concession process is assumed where there are no decommitment behaviors during the negotiation, i.e., the agent will not regret its compromise and ask for more benefits from the opponent in newly generated offers. In bilateral single-issue automated negotiation, the utility of the opponent's offers is monotonically increasing.

The bidding strategy of the opponent denotes how it provides its offers during a negotiation. A common one is called the time-dependent strategy where the offer is given based on time [7]. The utility function is given by:

$$u_t = u^{min} + (u^{max} - u^{min})(\frac{t_c}{t_{max}})^{\beta}, \tag{1}$$

where u^{min} and u^{max} denote the minimum and maximum utility of the opponent's offers respectively. t_c and t_{max} mean the current and the maximum negotiation time respectively. β is the concession parameter. $0 < \beta < 1$, $\beta = 1$ and $\beta > 1$ represent three types of concession strategies, which are called Conceder, Linear and Boulware respectively. The Conceder concedes dramatically toward u^{max} at the early stage of a negotiation, while the Boulware only makes significant concession when the time is close to t_{max}.

3 Opponent Modeling with Information Adaptation

3.1 Basic Notation

Before introducing OMIA, we first establish the notation for future use. We use t to denote the time of a negotiation, with different subscripts indicating specific time. For example, t_1 is the first time when the agent exchanges an offer with the opponent, t_c the current time, and t_{max} the maximum negotiation time. We use a discrete time setting so that the time is measured based on the negotiation rounds elapsed.

The offers received from the opponent are measured as utility values, which represent the agent's preferences over them. Higher utility values are preferred than lower ones. These utility values are the only information that the agent could obtain and make use of to build its opponent model when no other information is available. Also, the agent does not know the opponent's utility function. Thus, all offers exchanged by both sides will be measured in the agent's utility space. The utility that the agent receives from and offers to the opponent at time t could be respectively written as:

$$u_t^{oppo}, u_t^{my} \in [0, 1] \tag{2}$$

The utility is quantified between 0 and 1. Here we do not specify the types of utility function that maps offers to utility values because this is dependent on particular negotiation scenarios and should be determined accordingly.

As this paper is focusing on modeling the opponent, to make the formulas neater, we simplify the notation of utility received from the opponent by removing the superscript:

$$u_t \Leftrightarrow u_t^{oppo} \tag{3}$$

At time t_c, the agent will have received c offers from the opponent. These are called the *historical offers* and we use U_{t_c} to indicate it:

$$U_{t_c} = \{u_{t_i} | i = 1, ..., c\} \tag{4}$$

These historical offers are the only information that the agent has at the current time when no other sources of information are available.

All the notations used in this paper are listed in Table 1.

3.2 Predicting Future Concessions

The process of OMIA is first to construct the probability distributions $P(u_{t_f})$ of utility u_{t_f} for every future time t_f based on the historical offers. We will introduce this step in this Subsect. (3.2). Then, an expected utility calculation function is introduced to incorporate $P(u_{t_f})$, the probability that the opponent accepts or rejects the agent's offer $P(a_t = \{1, 0\})$, and the probability that the opponent quits or does not quit the negotiation at time t $P(q_t = \{1, 0\})$. This step will be described in the next Subsect. (3.3).

Table 1. Notation

Notation	Meaning
t_i	Negotiation time[a]
t_c	Current time
t_f	Future time
u_t	Received offer's utility at time t
u_t^{my}	Utility of the agent's counter-offer at time t
U_{t_c}	Historical offers at current time
$Learn(\cdot)$	Learning algorithm
$GP(\cdot)$	Gaussian process
$P(u_t)$	Probability distribution for u_t
$f_n(u_t; \mu_t, \sigma_t, 0, 1)$	Normalized and truncated probability density function for u_t
$F_n(u_t; \mu_t, \sigma_t, 0, 1)$	Normalized and truncated cumulative distribution function for u_t
E_{t_f}	Expected utility the agent will gain when trying to reach agreement at time t_f
t_m	Time between t_c and t_f
$u_{t_m}^{my}(t_f)$	Counter-offer corresponding to E_{t_f} provided at time t_m
$E_{t_f}(u_{t_m}^{my}(t_f))$	Expected utility from the agent's offers $u_{t_m}^{my}$ accepted by the opponent between time t_c and t_f
$E_{t_f}(u_{t_f})$	Expected utility from accepting the opponent's offer u_{t_f}
E^*	Maximum expected utility among E_{t_f}
t^*	Time corresponding to E^*
$P(a_t = \{1, 0\})$	Probability that the opponent accepts or rejects the agent's offer u_{t-1}^{my}
$P(q_t = \{1, 0\})$	Probability that the opponent quits or does not quit the negotiation at time t

[a] We use a positive integer subscript i to denote the negotiation time in accordance with the discrete time setting

First, we are going to predict the probability distributions $P(u_{t_f})$ using the historical offers. These distributions can be regarded as an estimation of the opponent's future concessions. Thus, the agent can exploit the concessions of the opponent by adaptively providing counter-offers. Predicting $P(u_{t_f})$ naturally requests a learning algorithm to analyze the historical offers and provide predictions for any future time:

$$Learn(U_{t_c}) \Rightarrow < P(u_{t_c+1}), ..., P(u_{t_{max}}) >, \tag{5}$$

where $Learn(\cdot)$ is a learning function that can satisfy this requirement. Any applicable learning approaches can be applied here. This allows the users of OMIA to flexibly choose the learning approaches.

In this paper, we choose the Gaussian process technique based on three considerations: (1) Gaussian process is a powerful non-linear interpolation tool and has been applied to address various learning tasks [6,8]. (2) Gaussian process provides both an estimation and its uncertainty (essentially a Gaussian distribution) for an unseen utility at any time, which is highly in accordance with our aims. (3) Gaussian process can work with variables with a continuous domain, which gives potential to expand OMIA to real-time negotiation. For the parameters, we use Matérn covariance function and linear mean function due to their robustness [16].

The Gaussian process will predict new probability distributions $P(u_{t_f})$ using the historical offers U_{t_c} at time t_c:

$$GP(U_{t_c}) \Rightarrow < P(u_{t_{c+1}}), ..., P(u_{t_{max}}) > \tag{6}$$

The output of the Gaussian process for any time t is a Gaussian distribution given by:

$$P(u_t) = f(u_t; \mu_t, \sigma_t) = \frac{1}{\sigma_t\sqrt{2\pi}} exp(\frac{-(u_t - \mu_t)^2}{2\sigma_t^2}), \tag{7}$$

where μ_t is the mean, i.e. the most likely value of u_t, and σ_t is the standard deviation.

The utility space is bounded in the range of $[0, 1]$, and the expected mean μ_t may exceed this bound. Thus, a normalization is needed to make the mean between 0 and 1:

$$f_n(u_t; \mu_t, \sigma_t) = normalize(f(u_t; \mu_t, \sigma_t)) \tag{8}$$

Furthermore, a truncated normal distribution should also be used to make the distribution in the range of $[0, 1]$. Its probability density function is given by:

$$\begin{aligned} P(u_t) &= f_n(u_t; \mu_t, \sigma_t, 0, 1) \\ &= \frac{f_n(u_t; \mu_t, \sigma_t)}{F_n(1; \mu_t, \sigma_t) - F_n(0; \mu_t, \sigma_t)} \end{aligned} \tag{9}$$

where $f_n(u_t; \mu_t, \sigma_t)$ is the normalized probability density function of opponent's utility at time t and $F_n(u_t; \mu_t, \sigma_t)$ is its cumulative distribution function.

When a negotiation goes to time t_c, Eqs. (6), (7), (8) and (9) will be performed to get probability distributions $P(u_{t_{c+1}}), ..., P(u_{t_{max}})$.

3.3 Making Concessions by Calculating Expected Utility

We assume that the opponent provides its offer first, so at time t_c, the opponent's latest offer is u_{t_c} and the agent's latest offer is $u^{my}_{t_{c-1}}$. After constructing the probability distributions for every future time after t_c, the agent needs to make a decision about whether to accept u_{t_c} or provide a counter-offer $u^{my}_{t_c}$.

Our idea is to exploit the future concessions of the opponent by calculating the expected utility E_{t_f} when the agent tries to reach agreement with the opponent at particular future time t_f. Here we regard reaching agreement at future time t_f as conceding toward μ_{t_f} and accepting the opponent's offer u_{t_f} if the agent's offers before time t_f are not accepted by the opponent. The μ_{t_f} is the estimated mean value of u_{t_f}. Let $u^{my}_{t_c}(t_f)$ be the agent's counter-offer at time t_c when conceding toward μ_{t_f}. Let E^* be the maximum expected utility among all E_{t_f}, and the corresponding time is t^*. The $u^{my}_{t_c}(t^*)$ will be set as the agent's final counter-offer $u^{my}_{t_c}$ ready to be provided to the opponent.

We choose μ_{t_f} as concession targets because of these two considerations:

(1) The μ_{t_f} is a moderate value which is not too high. As the agent does not know how the opponent will react to its concessions, avoiding being extreme is a reasonable way. Also, it is not too low so the agent would gain sufficient utility when conceding toward it.
(2) The μ_{t_f} indicates the most possible value that the opponent will offer at time t_f. Thus, conceding toward it would have a high chance of reaching agreement with the opponent.

In terms of determining how the agent concedes toward μ_{t_f}. Again, avoiding being extreme is a reasonable way. To be more exact, the agent should not keep its offers unchanged until time t_f nor concede immediately to μ_{t_f}. So we choose to concede linearly toward every μ_{t_f}.

When trying to reach agreement at future time t_f, the expected utility E_{t_f} consists of following two parts.

First, the opponent has a chance to accept one of the counter-offers $u^{my}_{t_m}(t_f)$ provided by the agent from time t_c to t_{f-1}. We use $E_{t_f}(u^{my}_{t_m}(t_f))$ to represent this part of expected utility. The $u^{my}_{t_m}(t_f)$ is calculated by performing linear interpolation:

$$u^{my}_{t_m}(t_f) = u^{my}_{t_{c-1}} + (\mu_{t_f} - u^{my}_{t_{c-1}})\frac{t_m - t_{c-1}}{t_f - t_{c-1}} \tag{10}$$

Each $u^{my}_{t_m}(t_f)$ is accepted when the opponent rejects all previous offers, does not quit, and accepts this one. The $E_{t_f}(u^{my}_{t_m}(t_f))$ is defined as:

$$\begin{aligned} E_{t_f}(u^{my}_{t_m}(t_f)) = (&\prod_{n=c+1}^{m} P(a_{t_n}=0)P(q_{t_n}=0)) \\ &\times P(a_{t_{m+1}}=1)u^{my}_{t_m}(t_f) \end{aligned} \tag{11}$$

When no extra private information is available, the $P(q_{t_n} = 0)$ could be specified by some prior assumptions, e.g. the opponent will not quit.

Second, the opponent has a chance to reject all offers from time t_c to t_{f-1}, does not quit, and gives u_{t_f} at time t_f. Then the agent will accept u_{t_f}, which constitutes the second part of the expected utility. We use $E_{t_f}(u_{t_f})$ to denote this part and it can be calculated as:

$$\begin{aligned} E_{t_f}(u_{t_f}) = (&\prod_{n=c+1}^{f} P(a_{t_n}=0)P(q_{t_n}=0)) \\ &\times \int_0^{u^{my}_{t_{f-1}}(t_f)} u_{t_f}P(u_{t_f})du_{t_f}, \end{aligned} \tag{12}$$

where $\int_0^{u^{my}_{t_{f-1}}(t_f)} u_{t_f} P(u_{t_f}) du_{t_f}$ is the expected utility that the opponent will give when it rejects offer $u^{my}_{t_{f-1}}(t_f)$.

In terms of determining the acceptance possibility $P(a_t = 1)$ when no extra private information is available, a common assumption for the opponent to accept an offer is that the offer will give extra benefit to the opponent, i.e., the agent provides less utility than what the opponent is going to give:

$$\begin{aligned} P(a_t = 1) &= P(u^{my}_{t-1} \leq u_t) \\ &= 1 - P(u_t < u^{my}_{t-1}) \\ &= 1 - F_n(u^{my}_{t-1}; \mu_t, \sigma_t, 0, 1), \end{aligned} \tag{13}$$

where $F_n(u^{my}_t; \mu_t, \sigma_t, 0, 1)$ is the normalized and truncated cumulative probability distribution of the opponent offers' utility.

The final expected utility E_{t_f}, which the agent is expected to reach at time t_f, is given by:

$$\begin{aligned} E_{t_f} &= \sum_{m=c}^{f-1} E_{t_f}(u^{my}_{t_m}) + E_{t_f}(u_{t_f}) \\ &= \sum_{m=c}^{f-1} (\prod_{n=c+1}^{m} P(a_{t_n} = 0) P(q_{t_n} = 0)) P(a_{t_{m+1}} = 1) u^{my}_{t_m} \\ &+ (\prod_{n=c+1}^{f} P(a_{t_n} = 0) P(q_{t_n} = 0)) \int_0^{u^{my}_{t_{f-1}}} u_{t_f} P(u_{t_f}) du_{t_f} \end{aligned} \tag{14}$$

For each future time t_f, we follow the same process to calculate the expected utility. The maximum expected utility E^* and the corresponding time t^* can be obtained by:

$$E^* = max(E_{t_f}) \quad t_f \in [t_{c+1}, t_{max}] \tag{15}$$

$$t^* = \text{argmax}_{t_f \in [t_{c+1}, t_{max}]} E_{t_f} \tag{16}$$

Finally, we compare E^* with u_{t_c}. As u_t is monotonically increasing, u_{t_c} is the maximum utility that the agent has received at current time. If $E^* > u_{t_c}$, it means that conceding toward μ_{t_f} may be more valuable so that the agent will reject u_{t_c} and choose to concede toward μ_{t^*} by setting $u^{my}_{t_c}$ as $u^{my}_{t_c}(t^*)$. Otherwise, if $E^* \leq u_{t_c}$, the agent will accept u_{t_c}. This decision procedure at time t_c can be formed as:

$$Decision(t_c) = \begin{cases} \text{accept } u_{t_c}, & E^* \leq u_{t_c} \\ \text{reject } u_{t_c} \text{ and provide } u^{my}_{t_c}(t^*), & E^* > u_{t_c} \end{cases} \tag{17}$$

4 Adaptation to Four Types of Private Information

In this section, we demonstrate how to use the proposed OMIA to adapt to four types of private information, which are (1) reservation value, (2) deadline, (3) bidding strategy and (4) acceptance probability.

Reservation value. The opponent will never provide an offer exceeding its reservation value. To adapt to a given reservation value u^r, the $P(q_{t_n} = 0)$ in Eq. 14 will be influenced by the opponent offers' utility. Given an u^r, the probability of not quitting $P(q_t = 0)$ could be specified by:

$$P(q_t = 0) = \begin{cases} 0, & u_t > u^r \\ 1, & u_t \leq u^r \end{cases} \tag{18}$$

Deadline. Similar to the reservation value, we also adapt to the deadline by specifying the $P(q_t = 0)$. Given an deadline t^d, $P(q_t = 0)$ is set by:

$$P(q_t = 0) = \begin{cases} 0, & t > t^d \\ 1, & t \leq t^d \end{cases} \tag{19}$$

Bidding strategy. Knowing the information of the opponent's bidding strategy means the agent knows what the opponent will offer at specific time. That is, pairs of time and utility (t_i, u_{t_i}) are given beforehand. OMIA utilizes (t_i, u_{t_i}) together with the historical offers U_{t_c} to train the Gaussian process models and predict u_{t_f}.

$$GP(U_{t_c}, (t_i, u_{t_i})) \Rightarrow < P(u_{t_{c+1}}), ..., P(u_{t_{max}}) > \tag{20}$$

Acceptance possibility. OMIA models the acceptance possibility $P(a_t = 1)$ from a probability point of view only with historical offers of the opponent. If there are other kinds of methods giving more precise estimations, OMIA can directly use them in Eq. 14 so that the expected utility is computed with adaptation to these acceptance possibility calculation methods.

5 Experimental Results

5.1 Experimental Setting

In this experiment, an agent and an opponent negotiate over a single issue. The utility of the opponent's offers follows a time-dependent strategy, i.e. Conceder, Linear or Boulware. The concession parameter is denoted as β. The utility of both sides is quantified in the agent's utility space in $[0, 1]$. The negotiation time is set from 1 to 50.

We design a series of experiments to show the information adaptation ability of OMIA. Table 2 shows the detailed experimental setting, including the experiment type, the information adapting to, opponent strategy, assumed estimated private information of the opponent, and assumptions for the opponent's decision for quit. The estimations of these types of private information may be provided by some prediction algorithms, or be prior knowledge of the agent. In this experiment, these estimations are set as prior knowledge and are assumed to be true.

We conduct 2 types of experiments. One is a case study in which the agent negotiates with a Boulware opponent with β being 4. A non-adaptive behavior is

Table 2. Experimental setting

Experiment type	Adapting to	Opponent strategy	Assumed estimated private information	Quit of opponent
Case study	No adaptation	Boulware, $\beta = 4$	/	Assume not quit
	Historical offers		/	
	Historical offers and bidding strategy		An offer worth utility 0.1296 will be given by the opponent at time 30	
	Historical offers and acceptance possibility		The acceptance possibility will become 0% after time 30	
Empirical analysis	Historical offers and reservation value	Conceder, Linear and Boulware	Three experiments with reservation value being 0.6, 0.7 and 0.8	Quit after its offer's utility > the estimated reservation value
	Historical offers and deadline	Uniformly select 20 β for each strategy type in the range of $[0.3, 0.8]$, $[0.9, 1.1]$, and $[2, 4]$ respectively	Three experiments with deadline being 30, 35 and 40	Quit after time > the estimated deadline

firstly studied. Then, we study three types of information to adapt to, which are the historical offers, bidding strategy and acceptance possibility. The historical offers are public information and can be used directly. For bidding strategy, we select a pair of time and utility $(40, 0.4096)$ on the curve of opponent's offers as its assumed estimated value, and use it along with the historical offers as the input of Gaussian process. This means that the agent knows the opponent will offer 0.4096 at time 40. Thus, we can compare the behavior of the agent with and without this estimated pair of time and utility. For acceptance possibility, we assume that the acceptance possibility will become 0 after time 30. We can then compare how will the agent behave with and without the estimation for acceptance possibility.

The other type of experiment is empirical analysis for the private information of reservation value and deadline. For each one, we have three experiments using three assumed estimated values, which are 0.6, 0.7 and 0.8 for reservation values, and 30, 35 and 40 for deadlines. In every experiment, different numbers of β are selected to cover three types of opponent (Conceder, Linear and Boulware) in order to get average results. We uniformly select 20 β for each strategy type in the range of $[0.3, 0.8]$, $[0.9, 1.1]$, and $[2, 4]$, respectively. Totally, we get 60 results for an experiment with particular assumed estimated value.

5.2 Results Analysis

The experimental results for the case studies are presented in Figs. 1 and 2. The results for the empirical analysis are showed in Tables 3 and 4.

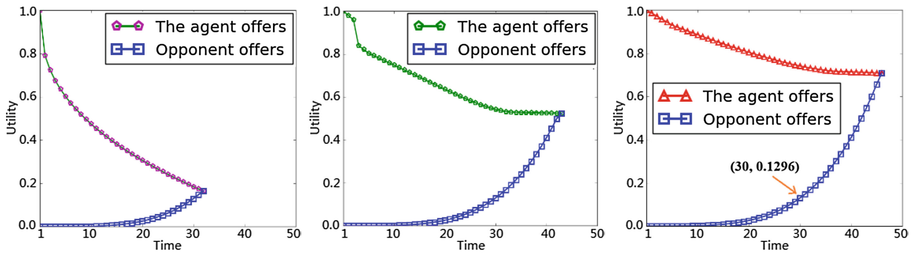

(a) Non-adaptive behavior (b) Adapting to historical offers only (c) Adapting to both historical offers and bidding strategy

Fig. 1. Negotiation process of the agent with non-adaptive behavior, adaptation to historical offers, and adaptation to both historical offers and bidding strategy

Figure 1(a) shows the negotiation process of the agent with a non-adaptive Conceder negotiation strategy. Figure 1(b) shows the process when the agent applies OMIA, adapting to the opponent's historical offers U_{t_c}. Figure 1(c) shows the process when the agent adapts to both U_{t_c} and the bidding strategy of the opponent. It can be seen that the agent with non-adaptive strategy achieves a low utility value, while the agent with adaptation to U_{t_c} achieves a higher utility value than the non-adaptive agent. When the opponent's bidding strategy (represented by the assumed estimated time-utility pair $(40, 0.4096)$) is given, we can see that the agent has more confidence about the future concessions of the opponent, makes fewer concessions during the early stage of the negotiation process, and finally achieves a higher utility value than the agent with adaptation only to U_{t_c}.

Figure 2(a) and (b) illustrate the negotiation process without and with adaptation to the assumed estimated acceptance possibility. The acceptance possibility will become 0 after time 30. The agent, receiving this information but not adapting to it, fails to reach agreement with the opponent. By contrast, the agent, adapting to this information, makes more concessions and finally reaches agreement at time 30. This is because the given acceptance possibility makes the expected utility when conceding toward time 30 maximum.

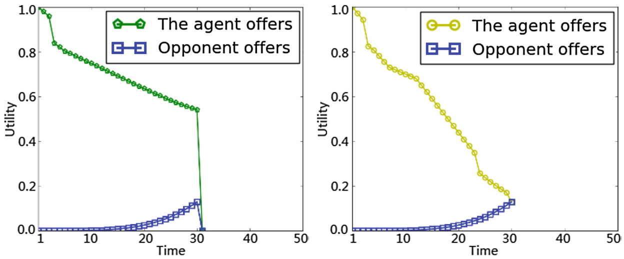

(a) Adapting to historical offers only

(b) Adapting to both historical offers and acceptance possibility

Fig. 2. Negotiation process of the agent without and with adaptation to the assumed estimated acceptance possibility

Table 3. Average utility achieved without and with adaptation to estimated reservation value

Estimated reservation value u^r	Average utility based on historical offers only	Average utility based on historical offers and estimated reservation value
0.6	0.22 ($\pm$0.008)	0.54 ($\pm$0.002)
0.7	0.36 ($\pm$0.006)	0.63 ($\pm$0.005)
0.8	0.54 ($\pm$0.009)	0.72 ($\pm$0.006)

Table 3 shows the average utility achieved without and with adaptation to estimated reservation value u^r. 95% confidence intervals are listed in the parentheses. We can see that the agent without adaptation to u^r achieves a lower average utility value than that with adaptation. This is caused by a high chance of failing to reach agreement with the opponent for the agent with a Boulware strategy and without adaptation to u^r. In addition, for the agent with adaption to u^r, the average utility achieved is always slightly lower than u^r. This shows that the agent makes agreement when the opponent offers' utility is close to u^r in order to gain as much utility as possible.

Table 4 shows the average utility achieved without and with adaptation to estimated deadline t^d. Similar to the result of the reservation value, the agent without adaptation to t^d achieves significantly lower average utility than that with adaptation. As there is a high chance of failing to reach agreement when the opponent approaches its deadline, but the agent does not adapt to it.

Table 4. Average utility achieved without and with adaptation to estimated deadline

Estimated deadline t^d	Average utility based on historical offers only	Average utility based on historical offers and estimated deadline
30	0.11 ($\pm$0.012)	0.51 ($\pm$0.014)
35	0.19 ($\pm$0.008)	0.60 ($\pm$0.004)
40	0.32 ($\pm$0.009)	0.70 ($\pm$0.009)

In summary, the experimental results show that OMIA could adapt to the historical offers by exploiting the future concessions of the opponent. Also, OMIA can adaptively guide the agent's negotiation behaviors by utilizing the availability and values of the opponent's private information in an expected utility measurement. As a result, The agent is able to successfully reach agreement with the opponent and achieve higher utility values comparing to those which lack the information adaptation ability.

6 Related Work

A lot of opponent modeling methods with different learning goals have been developed. Yu *et al.* [17] apply non-linear regression and Bayesian learning to estimate the opponent's reservation value and deadline. They introduce a concept named detecting region to estimated the lower and upper boundary of the reservation value and deadline. Historical offers are used to make the prediction more accurate during the process of the negotiation. Williams *et al.* [16] use the Gaussian process to estimate the future behavior of the opponent. The concession rate of the agent is then adaptively set based on the predictions during a single negotiation session. Time-based discounts and a risk function are applied in their model to handle the uncertainty. Oshrat *et al.* [13] create a negotiator called *KBAgent*. It can use the past negotiation results as a knowledge base to compute the acceptance probability for unseen offers. This approach allows agents to negotiate with people and can gain more utility than humans.

However, these opponent modeling methods are focusing on the use of the information they have. These methods overlook the fact that there may exist potentially available information. Having the ability of utilizing the potential information will help negotiation agents to reach agreement with the opponent and gain more utility values.

7 Conclusion

In this paper, we proposed a novel opponent modeling method called OMIA. The proposed method can not only adapt to the historical offers of the opponent, but also simultaneously adapt to four types of commonly used information in automated negotiations. Agents using OMIA can flexibly choose which types of information to adapt to. Also, OMIA makes little assumptions about the opponent, making itself a highly robust model. The experimental results showed that OMIA could exploit the future concessions of the opponent only based on the historical offers. When extra estimated information is given, OMIA could adaptively give guidance to the agent's behaviors, and the agent is able to gain as much utility from the opponent as possible under the estimated information.

Acknowledgments. This research is supported by a DECRA Project (DP140100007) from Australia Research Council (ARC), a UPA and an IPTA scholarships from University of Wollongong, Australia.

References

1. Baarslag, T., Hendrikx, M.J., Hindriks, K.V., Jonker, C.M.: Learning about the opponent in automated bilateral negotiation: a comprehensive survey of opponent modeling techniques. Auton. Agents Multi-Agent Syst. **30**(5), 849–898 (2016)
2. Baarslag, T., Hendrikx, M.J., Hindriks, K.V., Jonker, C.M.: A survey of opponent modeling techniques in automated negotiation. In: Proceedings of the 2016 International Conference on Autonomous Agents and Multiagent Systems, pp. 575–576. International Foundation for Autonomous Agents and Multiagent Systems (2016)
3. Broekens, J., Jonker, C.M., Meyer, J.J.C.: Affective negotiation support systems. J. Ambient Intell. Smart Environ. **2**(2), 121–144 (2010)
4. Carbonneau, R., Kersten, G.E., Vahidov, R.: Predicting opponent's moves in electronic negotiations using neural networks. Expert Syst. Appl. **34**(2), 1266–1273 (2008)
5. Coehoorn, R.M., Jennings, N.R.: Learning on opponent's preferences to make effective multi-issue negotiation trade-offs. In: Proceedings of the 6th International Conference on Electronic Commerce, pp. 59–68. ACM (2004)
6. Deisenroth, M.P., Fox, D., Rasmussen, C.E.: Gaussian processes for data-efficient learning in robotics and control. IEEE Trans. Pattern Anal. Mach. Intell. **37**(2), 408–423 (2015)
7. Fatima, S.S., Wooldridge, M., Jennings, N.R.: Multi-issue negotiation under time constraints. In: Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems: Part 1, pp. 143–150. ACM (2002)
8. Gal, Y., van der Wilk, M., Rasmussen, C.E.: Distributed variational inference in sparse Gaussian process regression and latent variable models. In: Advances in Neural Information Processing Systems, pp. 3257–3265 (2014)
9. Ji, S.J., Zhang, C.J., Sim, K.M., Leung, H.F.: A one-shot bargaining strategy for dealing with multifarious opponents. Appl. Intell. **40**(4), 557–574 (2014)
10. Kersten, G.E., Lai, H.: Negotiation support and E-negotiation systems: an overview. Group Decis. Negot. **16**(6), 553–586 (2007)

11. Lin, R., Kraus, S., Baarslag, T., Tykhonov, D., Hindriks, K., Jonker, C.M.: Genius: an integrated environment for supporting the design of generic automated negotiators. Computat. Intell. **30**(1), 48–70 (2014)
12. Moosmayer, D.C., Chong, A.Y.L., Liu, M.J., Schuppar, B.: A neural network approach to predicting price negotiation outcomes in business-to-business contexts. Expert Syst. Appl. **40**(8), 3028–3035 (2013)
13. Oshrat, Y., Lin, R., Kraus, S.: Facing the challenge of human-agent negotiations via effective general opponent modeling. In: Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems-Volume 1, pp. 377–384. International Foundation for Autonomous Agents and Multiagent Systems (2009)
14. Ren, F., Zhang, M.: Predicting partners' behaviors in negotiation by using regression analysis. In: Zhang, Z., Siekmann, J. (eds.) KSEM 2007. LNCS (LNAI), vol. 4798, pp. 165–176. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-76719-0_19
15. Ren, F., Zhang, M.: A single issue negotiation model for agents bargaining in dynamic electronic markets. Decis. Support Syst. **60**, 55–67 (2014)
16. Williams, C.R., Robu, V., Gerding, E.H., Jennings, N.R.: Using Gaussian processes to optimise concession in complex negotiations against unknown opponents (2011)
17. Yu, C., Ren, F., Zhang, M.: An adaptive bilateral negotiation model based on Bayesian learning. In: Ito, T., Zhang, M., Robu, V., Matsuo, T. (eds.) Complex Automated Negotiations: Theories, Models, and Software Competitions. SCI, vol. 435, pp. 75–93. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-30737-9_5
18. Zhang, J., Ren, F., Zhang, M.: Bayesian-based preference prediction in bilateral multi-issue negotiation between intelligent agents. Knowl. Based Syst. **84**, 108–120 (2015)

http://www.springer.com/978-3-319-71681-7

Autonomous Agents and Multiagent Systems
AAMAS 2017 Workshops, Best Papers, São Paulo, Brazil, May 8-12, 2017, Revised Selected Papers
Sukthankar, G.; Rodriguez-Aguilar, J.A. (Eds.)
2017, XII, 297 p. 73 illus., Softcover
ISBN: 978-3-319-71681-7