



ELSEVIER

Contents lists available at ScienceDirect

Engineering Applications of Artificial Intelligence

journal homepage: www.elsevier.com/locate/engappai

An efficient automated negotiation strategy for complex environments[☆]

Siqi Chen^{*}, Gerhard Weiss

Department of Knowledge Engineering, Maastricht University, P.O. Box 616, 6200 MD Maastricht, The Netherlands



ARTICLE INFO

Article history:

Received 21 February 2013
 Received in revised form
 21 May 2013
 Accepted 16 August 2013
 Available online 14 September 2013

Keywords:

Multi-agent systems
 Automated multi-issue negotiation
 Electronic business
 Opponent modeling
 Counter-offer prediction
 Empirical game theory

ABSTRACT

A complex and challenging bilateral negotiation environment for rational autonomous agents is where agents negotiate multi-issue contracts in unknown application domains with unknown opponents under real-time constraints. In this paper we present a negotiation strategy called EMAR for this kind of environment that relies on a combination of Empirical Mode Decomposition (EMD) and Autoregressive Moving Average (ARMA). EMAR enables a negotiating agent to acquire an opponent model and to use this model for adjusting its target utility in real-time on the basis of an adaptive concession-making mechanism. Experimental results show that EMAR outperforms best performing agents from the recent Automated Negotiating Agents Competitions (ANAC) in a wide range of application domains. Moreover, an analysis based on empirical game theory is provided that shows the robustness of EMAR in different negotiation contexts.

© 2013 Elsevier Ltd. All rights reserved.

1. Introduction

Autonomous agents that act and interact to reach their design objectives are of importance for a broad spectrum of potential applications in domains and fields such as task and service allocation (Dang and Huhns, 2006), electronic commerce and electronic markets (Lau et al., 2008; Ragone et al., 2008), distributed multi-project scheduling (Adhau et al., 2012), supply chain management (Wang et al., 2009), and pervasive computing (Park and Yang, 2008). In these agent-mediated applications, automated negotiation is central for efficiently establishing contracts about goods or services between agents that have conflicting interests. The work described in this paper focuses on automated bilateral multi-issue negotiation (e.g., Lai et al., 2004). A key feature of this negotiation form is that two agents negotiate with the intention to agree on a profitable contract for a product or service, where the contract consists of multiple issues which are of conflictive importance for the negotiators. Examples of such issues are price and quality. Specifically, this paper concentrates on realistic scenarios for bilateral multi-issue negotiations which are particularly complex for the following three reasons. First, the negotiating agents do not know each other (i.e., they have not

encountered before) and thus have no information about the preferences or strategies of their respective opponents. This makes it difficult to efficiently reach an agreement, especially because none of the agents knows what offers the respective other agent may consider as attractive. Second, we concentrate on negotiation with deadline and discount, that is, the negotiation is under real-time constraints (the agents thus should take into consideration at each time point the remaining negotiation time) and the final utility decreases over time according to some discounting factor, which means that unnecessary delays would result in negative effects on the negotiation outcome. And third, computational efficiency is important because agents may have very limited computing resources. Negotiation scenarios showing these characteristics are particularly challenging but common in reality. For example, in the case of open electronic sales platforms an agent may be engaged in bilateral multi-issue negotiations with a content service provider agent which it has never met before, and if the negotiation agent runs on a small mobile device then computational efficiency may be particularly crucial. Moreover, it may be that an agreement has to be achieved within short time because the provider may switch to other requests and because the agent may benefit from an early finish in terms of a discounted price or an extended service time. An environment showing the above features obviously places high demands upon the negotiation abilities of the agents, and we thus refer to negotiations running in such environments as complex negotiations.

Automated negotiation requires from the agents to have a high level of decision autonomy, so that they can decide on their own in real-time when and under what conditions they should perform what actions in order to reach a satisfactory agreement. This

[☆]This paper provides a substantially extended description and analysis of the EMAR negotiation strategy (Chen and Weiss, 2012b). The extension includes, among others, a detailed experimental comparison with the best agents from the 2012 International Negotiating Agents Competition (ANAC, 2012) and a game-theoretical robustness analysis.

^{*} Corresponding author. Tel.: +31 43 38 83909.

E-mail address: siqi.chen09@gmail.com (S. Chen).

objective is, however, difficult to achieve in practice, mainly due to the lack of sufficient knowledge about the opponents. Although methods have been proposed for solving this problem, they are typically based on simplifying assumptions regarding the opponent models used by the individual negotiating agents (see Section 2). Against this background, this paper presents a negotiation strategy called EMAR for the type of complex scenarios described above that aims at avoiding unrealistic assumptions. EMAR integrates two aspects that are known to be essential to successful negotiation among humans: efficient opponent modeling and adaptive concession making. Opponent modeling realized by EMAR predicts the utilities of the opponent's future counter-offers through two standard mathematical techniques, namely, Empirical Mode Decomposition (EMD, e.g., Hunag et al., 1998) and Autoregressive Moving Average (ARMA, e.g., Box et al., 1994). As the above underlining shall indicate that the acronym EMAR is composed of "EM" and "AR". Adaptive concession making is achieved by dynamically adapting the concession rate (i.e., the degree at which an agent is willing to make concessions in its offers) on the basis of the utilities of future counter-offers which can be expected according to the learned opponent model. EMAR combines these techniques to achieve agreements that maximize the own benefit specifically in complex negotiation environments (as characterized above) in any kind of domain such as electronic commerce or automated trading.

The remainder of this paper is structured as follows. Section 2 overviews important related work. Section 3 describes the standard negotiation environment underlying the described research. Section 4 presents EMAR in detail. Section 5 offers a detailed experimental analysis of EMAR. Section 6 identifies some important research lines induced by the described work and concludes the paper.

2. Related work

An early influential work in the field of automated negotiation is Faratin et al. (1998). This work raised awareness of issues related to concession making and tactical negotiation that are also relevant to the approach described here. Based upon this early work and subsequent research it triggered, it had been realized that successful negotiation needs to be based in one way or another on opponent modeling. Today various approaches are available that aim at generating and utilizing opponent models in order to optimize an agent's negotiation behavior (see Hendriks, 2011 for a good overview).

Available approaches can be classified into two groups. First, approaches that aim at learning the opponent's *preference profile*, including, e.g., the opponent's reservation value (i.e., the minimum utility an agent wants to obtain) and issue weight/value ordering. An example of such an approach is Coehoorn and Jennings (2004), which exploits kernel density estimation as an approximation technique for making negotiation trade-offs, thereby reaching a profitable outcome with less concession. Another example is Lin et al. (2008), where a Bayesian learning is used to approximate the opponent preference profile. A critical drawback of preference modeling is that it tends to quickly become computationally intractable for complex domains having a complicated domain structure or a large outcome space, i.e., the number of possible agreements is huge (especially if real-time constraints apply).

Second, approaches that aim at learning the opponent's *negotiation strategy*. For instance, Hou (2004) employs non-linear regression to predict the opponent's tactic (though in the context of single-issue negotiation), supposing that the opponent uses a pure tactic as introduced in Faratin et al. (1998) and that the types of tactics are also fixed. Saha et al. (2005) make use of Chebyshev polynomials to estimate the chance that the negotiation partner

accepts an offer in the context of repeated single-issue negotiations. Brzostowski and Kowalczyk (2006) investigate the prediction of future counter-offers online on the basis of the previous negotiation encounters by using differentials, thereby assuming that the opponent strategy is based on a mix of time- and behavior-dependent one. In Carbonneau et al. (2008) an artificial neural network is constructed with three layers to predict future counter-offers in a specific domain, and the performance was fairly good competing against a human negotiator. Its training process however requires a very large database of previous encounters and thus only off-line mode can be applied. In general, these approaches to negotiation strategy learning tend to suffer from relying on simplifying assumptions (as described above) and are therefore limited in their usage for complex negotiations.

An example of recent research addressing this complexity issue is Williams et al. (2011): Gaussian processes are applied to predict the future opponent concession before the deadline of negotiation is reached and to set the agent's "optimum" concession rate accordingly. This approach performed better than the best negotiating agents of ANAC 2010 and made the third place in ANAC 2011 (ANAC stands for International Automated Negotiating Agents Competition). Chen and Weiss (2012a) proposed the negotiation approach called OMAC that learns an opponent's strategy in order to predict future utilities of counter-offers by means of discrete wavelet decomposition and cubic smoothing splines. Another example is Hao and Leung (2012), where a negotiation strategy named ABiNeS was introduced for negotiations in complex environments. ABiNeS adjusts the time to stop exploiting the negotiating partner and also employs a reinforcement-learning approach to improve the acceptance probability of its proposals. An agent based on ABiNeS, called CUHKA-agent, was the winner of ANAC 2012.

The proposed approach, EMAR, which belongs to the "negotiation strategy learning" class, is designed for complex negotiations. It is able to predict the opponent's future moves in an ongoing negotiation and then to optimize its expected payoff, thereby relying on adaptive decision-making without requiring any prior knowledge about the opponents and the negotiation domains.

3. Negotiation environment

We adopt a basic bilateral multi-issue negotiation setting which is widely used in the agents field (e.g., Coehoorn and Jennings, 2004; Faratin et al., 1998) and the negotiation protocol we use is based on a variant of the alternating offers protocol proposed in Rubinstein (1982). Let $I = \{a, b\}$ be a pair of negotiating agents, i represents a specific agent ($i \in I$), J be the set of issues under negotiation, and j be a particular issue ($j \in \{1, \dots, n\}$, where n is the number of issues). The goal of a and b is to establish a contract for a product or service. Thereby a contract consists of a package of issues such as price, quality and quantity. Each agent has a lowest expectation for the outcome of a negotiation; this expectation is called reservation value ϑ . w_j^i ($j \in \{1, \dots, n\}$) denotes the weighting preference in which agent i assigns to issue j , where the issue weights of an agent are normalized (i.e., $\sum_{j=1}^n w_j^i = 1$ for each agent i). During negotiation agents (a and b) act in conflictive roles which are specified by their preference profiles. In order to reach an agreement they exchange offers O in each round to express their demands. An offer is a vector of values, with one value for each issue. The utility of an offer for agent i is calculated by the utility function defined as

$$U^i(O) = \sum_{j=1}^n (w_j^i \cdot V_j^i(O_j)) \quad (1)$$

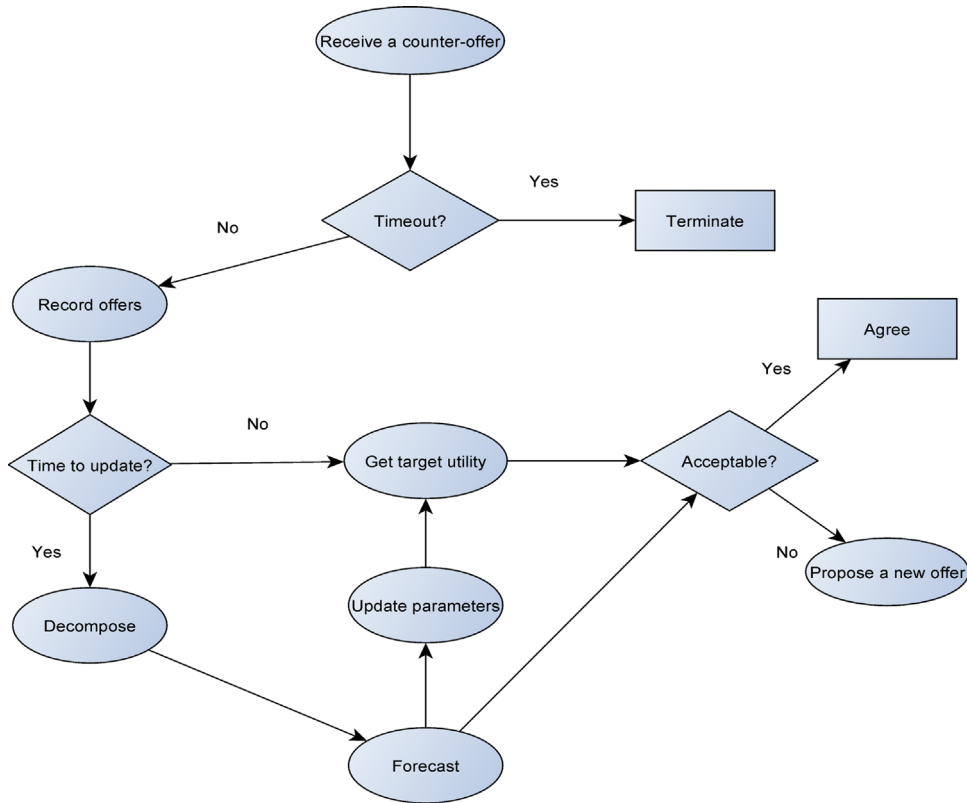


Fig. 1. Flowchart of EMAR.




Fig. 2. Illustrating the resulting components achieved by the application of EMD, where R_0 is the original signal, the i -th IMF is given as C_i and the residue is shown by R_3 . R_0 is offered by the agent IAMhaggler2011 in domain *England vs Zimbabwe*. (Details of the agent and the domain are given in Section 5.1.)

where w_j^i and O are as defined above and V_j^i is the evaluation function for i , mapping every possible value of issue j (i.e., O_j) to a real number.

Following Rubinstein's alternating bargaining model (Rubinstein, 1982), each agent makes, in turn, an offer in the form of a contract proposal. Negotiation is time-limited instead of being restricted by a fixed number of exchanged offers; specifically, each negotiator has a hard deadline by when it must have completed or withdraw the negotiation. The negotiation deadline of agents is denoted by t_{max} . In this form of real-time constraints, the number of remaining rounds is not known and the outcome of a negotiation depends crucially on the time sensitivity of the agents' negotiation strategies.

This holds, in particular, for discounting domains, that is, domains in which the utility is discounted with time. As usual for discounting domains, we define a so-called discounting factor δ ($\delta \in [0, 1]$) and use this factor to calculate the discounted utility as follows:

$$D(U, t) = U \cdot \delta^t \tag{2}$$

where U is the (original) utility and t is the standardized time. As an effect, the longer it takes for agents to come to an agreement the lower is the utility they can obtain.

After receiving an offer from the opponent, O_{opp} , an agent decides on acceptance and rejection according to its interpretation $I(t, O_{opp})$ of the current negotiation situation. For instance, this

decision can be made in dependence on a certain threshold $Thres^i$: agent i accepts if $U^i(O_{opp}) \geq Thres^i$, and rejects (and proposes a counter-offer) otherwise. As another example, the decision can be based on utility differences. Negotiation continues until one of the negotiating agents accepts or withdraws due to timeout.

4. EMAR

EMAR includes two core stages – opponent modeling and adaptive concession making – as described in detail in Sections 4.1 and 4.2, respectively. A third important stage, namely, its response mechanism to counter-offers, is described in Section 4.3. An overview of EMAR is given in Algorithm 1 and its flowchart is shown in Fig. 1. The individual steps of Algorithm 1 are explained in the following.

Algorithm 1. The EMAR approach. Let t_c be the current time, δ the time discounting factor, t_{max} the deadline of negotiation and ϑ the reservation value. O_{opp} is the latest offer of the opponent and O_{own} is a new offer to be proposed by EMAR. χ is the time series composed of the maximum utilities over intervals. ξ is the lead time for prediction and ω is the estimated central tendency of χ . c_i is the i -th IMF component and r_n is the final residue (more details will be given in the following text). E is the predicted received utility series. u_{res} is the dynamic reservation utility, specifying the lowest expectation to negotiation payoff, and e_{min} is the conservative estimation of opponent concession. R is the dynamic conservative expectation function. u' is the target utility at time t_c .

```

1: Require:  $R, \delta, \xi, \vartheta, t_{max}$ 
2: while  $t_c <= t_{max}$  do
3:    $O_{opp} \leftarrow receiveMessage$ ;
4:    $recordBids(t_c, O_{opp})$ ;
5:   if  $TimeToUpdate(t_c)$  then
6:      $\chi \leftarrow preprocessData(t_c)$ ;
7:      $(c_0, \dots, c_n, r_n) \leftarrow decompose(\chi)$ ;
8:      $(\omega, E) \leftarrow forecast(c_0, \dots, c_n, r_n, \xi)$ ;
9:      $(u_{res}, e_{min}) \leftarrow updateRes(\omega, \chi, \vartheta, t_c)$ ;
10:     $R \leftarrow (u_{res}, e_{min})$ ;
11:   end if
12:    $u' \leftarrow getTarget(t_c, E, \delta, R)$ ;
13:   if  $isAcceptable(u', O_{opp}, t_c, \delta)$  then
14:      $accept(O_{opp})$ ;
15:   else
16:      $O_{own} \leftarrow constructOffer(u')$ ;
17:      $proposeNewBid(O_{own})$ ;
18:   end if
19: end while

```

4.1. Opponent modeling

Opponent modeling realized by EMAR aims at predicting the future moves of the negotiating opponents. The process of opponent modeling corresponds to the lines 2–11 in Algorithm 1. When receiving a new bid from the negotiation opponent at the time t_c , the agent records the time stamp t_c and the utility $U(O_{opp})$ this bid offers according to its utility function. The maximum utilities in consecutive equal time intervals and the corresponding time stamps are used periodically as input for predicting the opponent's behavior (lines 5 and 6). The reasons for *periodical* updating are twofold. First, this reduces the computation complexity of EMAR so that the response speed is improved. Assume that *all* observed counter-offers were taken as input, then it would be necessary to deal with perhaps many thousands of data points at once. This computational load would have a clear negative impact on the

quality of negotiation in a real-time setting. Second, the effect of noise can be reduced. This is important because in multi-issue negotiations a small change in utility of the opponent can result in a large utility change for the other agent – and this can easily result in a fatal misinterpretation of the opponent's behavior.

The general idea behind opponent modeling realized by EMAR is to apply the “divide-and-conquer” principle to construct a reasonable forecasting methodology. Opponent modeling is mainly based on a combination of Empirical Mode Decomposition (EMD, e.g., Hunag et al., 1998; Huang and Shen, 2005; Flandrin et al., 2004) and Autoregressive Moving Average (ARMA, e.g., Box et al., 1994). In more detail, EMD is first employed to decompose the time series given by the utilities of past counter-offers into a finite number of components in order to make the prediction task simplified, and then ARMA is applied to predict future values of these sub-components.

EMD, which is based on the Hilbert–Huang transform (HHT), is a decomposition technique which relies on time-local characteristics of data and can deal with nonlinear and non-stationary time series in an adaptive manner. It has been widely applied as a powerful data analysis tool in a broad scope of fields such as finance, image processing, ocean engineering and solar studies. A main advantage of EMD as a decomposition method is that it is very suitable for analyzing complicated data and is fully data driven (thus requiring no additional decomposition information) – this makes EMD an effective and efficient decomposition method. Compared to traditional Fourier and wavelet decompositions, EMD has several distinct advantages (Huang et al., 2003; Yu et al., 2008). First of all, fluctuations within a time series are automatically selected from the time series. Second, EMD can adaptively decompose a time series into several independent components called Intrinsic Mode Functions (IMFs). With the usage of the IMFs a residue can be calculated which captures the main trend of the time series. Lastly, unlike wavelet decomposition, no filter base function needs to be determined beforehand – which is particularly helpful when there is no prior knowledge about which filters work properly.

The IMFs satisfy the following conditions:

1. In the whole data set (time series), the number of extrema and the number of zero crossings must either equal or differ at most by one.
2. At any point, the mean value of the envelope defined by the local maxima and the envelope defined by the local minima is zero.

Any data series can be decomposed into IMFs according to the following sifting procedure (let $k \geq 1$, k indicates the iterative decomposition level):

1. Take signal r_{k-1} as input, with r_0 representing the original signal $\chi(t)$.
 - (a) Identify all local extrema of the signal r_{k-1} .
 - (b) Construct the upper envelop $Upp(r_{k-1})$ and the lower envelop $Low(r_{k-1})$ by interpolating via a cubic spline the maximum and minimum values, respectively.
 - (c) Approximate the local average by the envelop mean $Mean(r_{k-1}) = (Upp(r_{k-1}) + Low(r_{k-1}))/2$.
 - (d) Compute the candidate implicit mode $h_{kn} = r_{k-1} - Mean(r_{k-1})$.
 - (e) If h_{kn} is an IMF, then calculate r_k as $r_k = r_{k-1} - h_{kn}$. Otherwise replace r_{k-1} with h_{kn} and repeat sifting.
2. If r_k has an implicit oscillation mode, set r_k as input signal and repeat step 1.

This sifting process serves two purposes: to eliminate riding waves and to make the wave profiles symmetric.




Fig. 3. Illustrating the prediction power of EMAR based on ARMA. The original time series χ , represented by the thick solid line, is received from negotiation with agent Agent_K2 in domain *Camera*. The prediction is depicted by the thin solid line, and the two dashed lines show the estimated upper and lower bounds of χ . The vertical thick dash-dot line indicates the time point at which EMAR calculated the prediction, and the circles right to this line are the utilities actually received in the subsequent negotiation phase.

The decomposition procedure can be repeated on all subsequent components r_j , and the result is

$$r_0 - c_1 = r_1, \quad r_1 - c_2 = r_2, \quad \dots, \quad r_{n-1} - c_n = r_n. \quad (3)$$

This procedure terminates when (1) the latest residue r_k becomes a monotonic function (from which no more IMFs can be extracted) or (2) the IMF component c_k or the residue becomes less than the predetermined value of substantial consequence. Overall, c_1 contains the signal at a fine-grained time scale and subsequent IMFs include information at increasingly longer time periods (i.e., lower frequencies). Eventually, the data series $\chi(t)$ can be expressed by

$$\chi(t) = \sum_{i=1}^n c_i + r_n \quad (4)$$

where n is the total decomposition layer (i.e., the number of IMFs), c_i is the i -th IMF component and r_n is the final residue (which represents the main trend of the data series). With that EMAR is able to achieve a decomposition of the data into n empirical modes and one residue (refer to line 7). The IMFs contained in each frequency band are independent and nearly orthogonal to each other (with all having zero means) and they change with variation of the data series $\chi(t)$, while the residue part captures the central tendency. An example can be found in Fig. 2, which shows χ and resulting IMFs along with the residue part. Observation can clearly tell that the frequency characteristic of these IMFs is becoming increasingly lower, leading to the last term, i.e., r_3 , just remaining the core tendency of the original signal.

In the next stage, ARMA is used to predict all resulting components, and then ensemble them to forecast opponent behaviors (shown in line 8). ARMA is a common regression analysis model widely used in many fields, with the formal expression as follows:

$$\left(1 + \sum_{i=1}^p \phi_i L^i\right) X_t = \left(1 + \sum_{i=1}^q \theta_i L^i\right) \epsilon_t \quad (5)$$

where L is the lag operator, ϕ_i are parameters for the p -order autoregressive term, θ are parameters for the moving average term with q order, and ϵ is a parameter capturing white noise. The parameters p and q are then chosen via the Akaike information criterion (AIC). AIC is a measure of the relative goodness of fit of a

statistical model. For a through discussion of it, interested reader is advised to refer to Ljung (1999).

Eq. (5) is applied with appropriate parameters for each component extracted via EMD (i.e., c_i and the residue) for the purpose of making accurate prediction, and then EMAR ensembles them to predict the future counter-offers of the opponent. Fig. 3 exemplifies this methodology, depicting the prediction power achieved through ARMA with a lead time of six intervals. Further details of usage are given in Section 4.2.

4.2. Adaptive decision-making

EMAR adjusts the concession rate on the basis of the generated opponent model. Thereby a dynamic conservative expectation $R(t)$ is used to avoid “irrational concession” caused by inaccurate or, more importantly, over-pessimistic predictions. This makes sense in the case of negotiation opponents that are “sophisticated and tough” and aim at avoiding (or maximally delaying) concession making in bargaining: in this case prediction can lead to a misleading, very low expectation about the utility offered by the opponent and this, in turn, could result in an adverse concession behavior. Furthermore, using global prediction could make this situation even worse. (This phenomenon is also considered in Section 5.2.)

$R(t)$ guarantees the desired minimum utility at each step, yielding values which are taken as the lower bound of the agent's expected utilities. For the purpose of adaptation to complex negotiation sessions, $R(t)$ requires two parameters e_{\min} and u_{res} . They are both periodically updated depending on the forecast of the opponent concession (see line 10). e_{\min} is defined as the minimum expectation of the compromise suggested by the opponent. Specifically, e_{\min} is set to the maximum value of $\psi^{\text{low}}(t)$, which is the estimated lower bound of the predicted χ given by the central trend. Formally:

$$\psi^{\text{low}}(t) = \omega(t)(e(r_{[0,t_1]}) - \sigma(r_{[0,t_1]})) \quad (6)$$

where ω is the predicted main tendency of χ , $r_{[0,t_1]}$ is the series including the ratio between ω over χ within $[0, t_1]$, the operator e is the expected value and σ the standard deviation.

Having obtained ψ^{low} , e_{\min} can be defined as follows:

$$e_{\min} = \begin{cases} \vartheta & \text{if } \vartheta > \text{Max}(\psi^{\text{low}}(t_c + \xi)) \\ \text{Max}(\psi^{\text{low}}(t_c + \xi)) & \text{otherwise} \end{cases} \quad (7)$$

where $\text{Max}(x)$ returns the maximum value of input vector x and ϑ is the reservation value. Because counter-offers with utilities indicated by ψ^{low} have already been received or can be expected in the future opponent moves, the use of the maximum value assures an increase of the agent's potential profit at low risk of a failure.

The variable u_{res} is the dynamic reservation utility specifying the lowest expectation of the eventual benefit at the time point t_i . Formally this is captured by

$$u_{\text{res}} = \begin{cases} \vartheta & \text{if } \vartheta > \text{Max}(\psi^{\text{low}}(t_i)) \\ \frac{1}{2}(\text{Max}(\psi^{\text{low}}(t_i)) + \vartheta) & \text{otherwise} \end{cases} \quad (8)$$

Because the final negotiation outcome (failure or agreement) is more sensitive to u_{res} than e_{\min} , EMAR adopts a cautious and conservative way to specify it, where only $\psi^{\text{low}}(t_i)$ is considered.

Based on the above specifications, the dynamic conservative expectation function, in principle, should concede over time and dynamic reservation utility (u_{res}), whereas it is proportional to the minimum expectation (e_{\min}) and the discounting factor (note a small value of discounting factor means larger time pressure). Thus $R(t)$ can be characterized as a dynamic conservative

expectation function that carefully estimates utilities. $R(t)$ can be instantiated in different ways, for example:

$$R(t) = \delta^\eta \cos\left(\frac{1-\delta}{1.1}t^\lambda\right)(1-t^{1/\beta})(\text{getMaxU}(P)-u_{\text{res}}) + \frac{e_{\min}-u_{\text{res}}}{2}(1-t^{5\delta}) + e_{\min} \quad (9)$$

where β and λ are concession factors affecting the concession rate, $\text{getMaxU}(P)$ is the function specifying the maximum utility dependent on a given preference P , δ is the discounting factor, and η is the risk factor which reflects the agent's optimal expectation about the maximum utility it can achieve. The cosine function is used because it is a monotone decreasing function over $[0,1]$ with the codomain being between 0.5 and 1, and the discounting factor is multiplied by 5 to avoid quick concession already in the early negotiation stage.

The subsequent process is then to decide the target utility EMAR expects to gain, represented by line 12. The ensemble of all predicted components provides useful information about the opponent behavior in the lead time. This is essential because the observation of ω (and its estimated bound ψ) only gives the ambiguous area where opponent would make a compromise (rather than how the move might look like). Let the predicted utility series be $E(t)$, given as follows:

$$E(t) = \sum_{i=1}^n f_i(c_i(t), \xi) + f_{n+1}(r_n(t), \xi) \quad (10)$$

where $f_i(x, y)$ is the corresponding prediction model for components c_i (the IMFs) and r_n (the residue). The first parameter corresponds to the input data, the second is the lead time for prediction.

Assume that the future expectation we have obtained from $E(t)$ is optimistic (i.e., there exists an interval $\{T|T \neq \emptyset, T \subseteq [t_c, t_s]\}$), that is,

$$E(t) \geq R(t), \quad t \in T \quad (11)$$

where t_s is the end point of the predicated series and $t_s \leq t_{\text{max}}$. In the present case the time \hat{t} at which the maximal expectation \hat{u} is reached is set as follows:

$$\hat{t} = \underset{t \in T}{\text{argmax}} E(t) \quad (12)$$

Moreover, in this case \hat{u} is defined as

$$\hat{u} = E(\hat{t}) \quad (13)$$

On the other hand, now assume that the estimated opponent concession is below the agent's expectations (according to $R(t)$), that is, there exists no such time interval T as in the "optimistic case". Then it is necessary to define the probability of accepting the best possible utility that can be achieved under this pessimistic expectation. This probability should be inversely proportional to the minimum adjusted difference between $D(R(t_\nu), t_\nu)$ and $D(E(t_\nu), t_\nu)$, and the discounting factor, which is formulated here as

$$\varphi = 1 - \frac{D(R(t_\nu), t_\nu) - D(E(t_\nu), t_\nu)}{\rho \cdot \sqrt{1 - \delta} D(\text{getMaxU}(P)\delta^\eta, t_\nu)}, \quad t_\nu \in [t_c, t_s] \quad (14)$$

where ρ indicates the acceptance tolerance for the pessimistic forecast and t_ν is given by

$$t_\nu = \underset{t \in [t_c, t_s]}{\text{argmin}} (|D(E, t) - D(R, t)|) \quad (15)$$

Depending on the probability φ the best possible outcome in the "pessimistic" scenario is chosen as the target utility. The rationale behind it is that if the agent rejects the "locally optimal" counter-offer (which is not too negative in accordance with ρ), it probably gives up the opportunity to reach a fairly good agreement. In the acceptance case, \hat{u} and \hat{t} are defined as $E(t_\nu)$ and t_ν , respectively. Otherwise, \hat{u} is defined as -1 , meaning it does not take an effect, and $R(t_c)$ is used to set the target utility u' . When the agent expects

to achieve a better outcome (see Eq. (11)), it chooses the optimal estimated utility \hat{u} as its target utility (see Eqs. (12) and (13)).

Obviously, it is not rational and smart to concede immediately to \hat{u} when $u_l \geq \hat{u}$, or it is appropriate for an agent to shift to \hat{u} without delay if $u_l < \hat{u}$ (especially because the predication may be not very accurate). To deal with this, EMAR simply concedes linearly. More precisely, the concession rate is dynamically adjusted in order to be able to "grasp" every chance to maximize profit. Overall, u' is calculated as follows:

$$u' = \begin{cases} R(t_c) & \text{if } \hat{u} = -1 \\ \hat{u} + (u_l - \hat{u}) \frac{t_c - \hat{t}}{t_l - \hat{t}} & \text{otherwise} \end{cases} \quad (16)$$

where u_l is the utility of last bid before EMAR performs prediction process at time t_l .

4.3. Response to counter-offers

This stage corresponds to lines 13–18 in Algorithm 1. When the expected utility u' has been determined, the agent needs to examine whether the utility of the latest counter-offer $U(O_{\text{opp}})$ is better than u' or whether it has already proposed this offer in the earlier negotiation process. If either of these two conditions is satisfied, the agent accepts this counter-offer and finishes the current negotiation session. Otherwise, the agent constructs a new offer which has a utility within some range around u' . There are two main reasons for this kind of consideration. First, in negotiations with multiple issues it is possible to generate a number of offers whose utilities are the same or very similar for the offering agent, but have very different utilities the opposing negotiator. (Note that in real-time constraints environment there are no limits for the number of negotiation rounds, which means that an agent in principle can construct a large amount of offers having a utility close to u' and, thus, has the opportunity to explore the utility space with the purpose of improving the acceptance chance of its proposals.) Second, it is sometimes impossible to make an offer whose utility is exactly equivalent to u' . It is thus reasonable that an agent selects any offer whose utility is in the range $[(1-0.005)u', (1+0.005)u']$. If no such solution can be constructed, the agent makes its latest bid again in the next round. Moreover, with respect to negotiation efficiency, if u' drops below the utility provided by the best counter offer, the agent chooses that best counter offer as its next offer. This makes much sense because this counter offer can well satisfy the expected utility of the opponent who then will be inclined to accept it.

5. Experimental analysis

In order to evaluate the performance of EMAR, the General Environment for Negotiation with Intelligent multipurpose Usage Simulation (GENIUS) (Hindriks et al., 2009) is used as the testing platform. GENIUS is the standard platform for the annual International Automated Negotiating Agents Competition (ANAC) (Fujita et al., 2013). In this environment an agent can negotiate with other agents in a variety of domains, where the utility function is defined by the preference of each negotiating party. The performance of an agent (its negotiation strategy) can be evaluated via its utility achievements in negotiation tournaments which include a possibly large number of negotiation sessions over a variety of negotiation domains. Section 5.1 describes the overall experimental setting and Sections 5.2 and 5.3 then present the experimental results in the context of tournament performance. Furthermore, an empirical game theoretic evaluation is applied to study the robustness of the proposed method. The results of this evaluation are presented in Section 5.4.

5.1. Environmental setting

EMAR is compared against the best winners (i.e., the top five agents) of ANAC 2011 instead of all ANAC 2011 agents in order to make the setting more competitive; these are HardHeaded, Gabhoninho, IAMhaggler2011, BRAMAgent and Agent_K2 (descending order in ANAC 2011, and for technicalities of these agents refer to Baarslag et al., in press). Moreover, we use nine standard domains created for ANAC as testing scenarios. Around half of these domains were used in ANAC 2010 and others were from later ANAC competitions. It makes the setting avoid advantageous bias for EMAR (note that the developers of the 2011 winners knew the ANAC 2010 domains and could optimize their agents accordingly). Additionally, to evaluate the performance in domains where agent performance are affected by time-discounting factors, we equip the domains with a discounting factor. Each domain thus has two versions, that is, a non-discounting and a discounting version. The domains used for the experimental analysis are carefully chosen so that they cover a very broad range of important domain features. Specifically, the domains vary w.r.t. level of competition/opposition, size of outcome space, and time pressure. Moreover, the set of chosen domains includes many test domains used in recent related works (e.g., see Williams et al., 2011; Hao and Leung, 2012; Chen and Weiss, 2012a), which makes EMAR directly comparable to other approaches. The application domains are overviewed in Table 1 with respect to four key aspects. For each

version of the domains, we run a tournament consisting of six agents (i.e., the five 2011 winners and the EMAR agent) 10 times to get results with high statistical confidence, where each agent negotiates against all other agents in different roles. These roles are predefined in ANAC domains and correspond to conflictive preferences like “buyer” and “seller”. The agents do not have any information about their opponents’ strategies and they are prohibited to take advantage of knowledge they might have acquired in previous negotiation sessions about their opponents. The duration of a negotiation session is 180 s.

Furthermore, the experiments are also done using the top agents from ANAC 2012 as benchmarks. The 2012 competition, as the latest international negotiation competition so far, was held later than the finalization of EMAR. It is therefore of interest to compare how EMAR performs playing against those most recent and advanced agents. The results are presented in Section 5.3 in addition to the primary results (with the best winners of ANAC 2011) given in Section 5.2.

The EMAR agent divides the overall duration of a session into 100 consecutive intervals of 1.8 s each. The lead time ξ is 6, the pair concession coefficients of (β, λ) is (0.04, 3) and the risk factor η is 0.2, the tolerance coefficient ρ is 0.05. These values work well in practice, but we have not intended to tweak them to stay away the issues of over-fitting and unfair competition.

5.2. Primary competition results

We show the experimental results achieved by each agent in terms of raw score (i.e. the score or utility from the experimental results without being normalized) averaged over the non-discounting and discounting version of each domain in Fig. 4. As depicted in the figure, the results clearly highlight the excellent bargaining skills of EMAR. Precisely, EMAR wins in all 10 domains, with achieving a mean score of 18% higher than that of the other five competing agents. Most notably and impressively, it outperforms others in the domain with the largest outcome space—Energy, by a margin of 31% over the mean score of the ANAC 2011 agents.

Table 2 shows the mean (raw) scores of all agents averaged over the non-discounting and discounting versions of the domains. Our agent, on average, is the best, both in the context of negotiations where the time-discounting pressure takes effect or not. In more detail, in the non-discounting domains, the average

Table 1
Overview of application domains.

| Domain | Issues | Size | Opposition | Discounting factor | |
|---------------------|--------|---------|------------|--------------------|--------------|
| | | | | Non-dis. version | Dis. version |
| Travel | 7 | 188,160 | Medium | 1.0 | 0.4 |
| Itex vs Cypress | 4 | 180 | Strong | 1.0 | 0.5 |
| SuperMarket | 6 | 98,784 | Strong | 1.0 | 0.5 |
| England vs Zimbabwe | 5 | 576 | Medium | 1.0 | 0.6 |
| Energy | 8 | 390,625 | Strong | 1.0 | 0.6 |
| NiceOrDie | 1 | 3 | Strong | 1.0 | 0.6 |
| Amsterdam party | 6 | 3024 | Medium | 1.0 | 0.7 |
| Grocery | 6 | 1600 | Weak | 1.0 | 0.7 |
| Camera | 5 | 3600 | Weak | 1.0 | 0.89 |

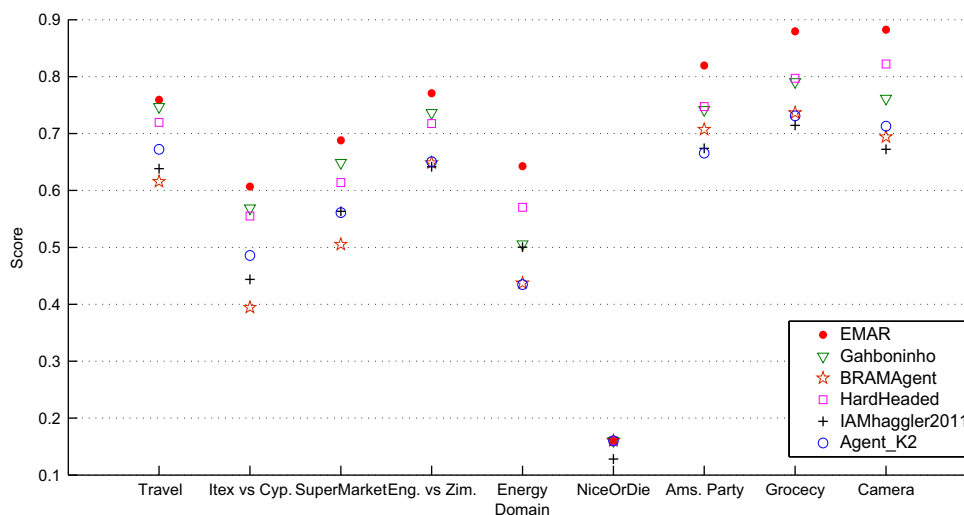


Fig. 4. Raw scores of all agents in the 10 domains, averaged over the respective non-discounting and discounting domains. The vertical axis gives the scores and the horizontal axis shows the domains.

performance of the five opponents reaches 82.6% of ours, and EMAR outperforms the other agents by 14% across the discounting version of the domains. Moreover, EMAR shows the smallest standard deviation among all agents. The overall performance is summarized in Table 3. Normalization is done in the standard way, using the maximum and minimum utilities obtained by all agents. According to the overall performance shown in this table, EMAR is ranked number one, with an average normalized score of 0.718. This is 15.2% above the second best agent – HardHeaded – and 30% above the mean score of all five opponents. Moreover, the performance of EMAR is the most stable, with merely 54.2% of the mean standard deviation of the other agents. EMAR is followed by HardHeaded and Gahboninho; these two agents made the first two places in ANAC 2011. Agent_K2, which is an updated version of the champion (named Agent_K) in ANAC 2010, made the fourth place. We notice that the ranking is somewhat different from the

Table 2
Average raw score of each agent for the non-discounting and discounting domains.

| Agent | Non-discounting domains | | Discounting domains | |
|----------------|-------------------------|--------------------|---------------------|--------------------|
| | Mean | Standard deviation | Mean | Standard deviation |
| EMAR | 0.801 | 0.0038 | 0.579 | 0.0026 |
| Gahboninho | 0.776 | 0.0085 | 0.481 | 0.0081 |
| HardHeaded | 0.744 | 0.0138 | 0.523 | 0.0063 |
| Agent_K2 | 0.622 | 0.0083 | 0.505 | 0.0061 |
| IAMhaggler2011 | 0.582 | 0.0044 | 0.524 | 0.0027 |
| BRAMAgent | 0.582 | 0.0078 | 0.506 | 0.0069 |

Table 3
Overall performance. The bounds are based on the 95% confidence interval. The initial letter (**bold**) of each strategy is taken as the identifier for the later EGT analysis.

| Agent | Normalized score | | | |
|----------------|------------------|--------------------|--------------|--------------|
| | Mean | Standard deviation | Lower bound | Upper bound |
| EMAR | 0.718 | 0.0058 | 0.706 | 0.729 |
| HardHeaded | 0.623 | 0.0128 | 0.598 | 0.649 |
| Gahboninho | 0.617 | 0.0119 | 0.593 | 0.641 |
| Agent_K2 | 0.520 | 0.0107 | 0.498 | 0.541 |
| IAMhaggler2011 | 0.507 | 0.0063 | 0.494 | 0.519 |
| BRAMAgent | 0.494 | 0.0118 | 0.471 | 0.518 |

final results of ANAC 2011 for other agents. Based on the experimental results we think that this is mainly caused by the participation of EMAR, leading to changes in the relative strength among the negotiating agents. To sum up, these results show that EMAR is pretty efficient and significantly outperforms in a variety of negotiation scenarios the state-of-the-art automated negotiators (resp. negotiation strategies) currently available.

An interesting observation is that there is the noticeable gap between EMAR and IAMhaggler2011. More specifically, this agent only achieves 70.6% of the performance of EMAR in terms of normalized utility. As described in Williams et al. (2011), similar to EMAR IAMhaggler2011 aims at predicting an opponent's future in order to be able to adjust its own behavior appropriately. Unlike EMAR, IAMhaggler (i) applies Gaussian process as prediction tool and (ii) adapts its concession rate on the basis of a global prediction view (i.e., on the basis of the whole preceding negotiation process). Our experimental studies suggest that a main reason for this performance gap lies in the global prediction view: this view seems to be vulnerable to "irrational concession making" induced by pessimistic predictions (see also Section 4.2). The phenomenon of irrational concession becomes increasingly apparent when IAMhaggler2011 bargains with "sophisticated and tough" opponents like HardHeaded, Gahboninho, and EMAR. For instance, when competing against these opponents in a highly competitive domain (i.e., *Itex vs Cypress*), IAMhaggler2011 only achieves a utility of 0.313 while the three opponents achieve 0.903 on average.

5.3. Additional competition results

The above setting focuses on non-discounting domains. As an additional evaluation of the negotiation performance of EMAR in discounting settings, we ran a tournament with the five best winners of ANAC 2012; these agents are CUHKAgent, AgentLG, OMACagent, TheNegotiator reloaded, and BRAMAgent_2. For this tournament, we sampled the discounting factor of each domain from a uniform distribution in the interval [0.5,1] to obtain conclusive performance results for different time-pressure scenarios. Note that we do not mix ANAC 2011 and 2012 agents, because the concept of reservation values was not yet used in the 2011 competition and the agents thus were not sensitive to this concept. Instead, we isolate the effect of reservation value on the agents' performance by averaging over the results for the same reservation value. The reservation values ϑ used here are 0, 0.25

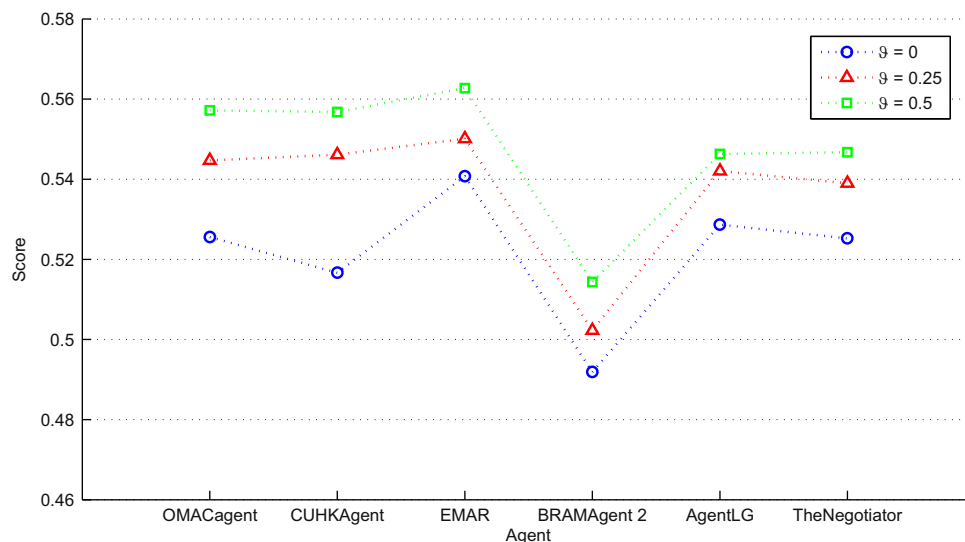


Fig. 5. Agent performance under different levels of reservation value when the ANAC 2012 agents are considered.

and 0.5. Fig. 5 shows the mean score of each agent for each value of ϑ . All agents, unsurprisingly, manage to achieve more profit as the reservation value increases. EMAR is the dominating agent regardless of the value of ϑ . The top three agents of ANAC 2012 – CUHKAgent, AgentLG and OMACagent – finish second for different reservation values. Table 4 summarizes the outcome of this tournament. As can be seen from this table, EMAR is superior to the other agents in that it achieves the highest score among all strategies; overall, it exceeds the performance averaged over the best winning strategies of ANAC 2012 by a margin of 3.6%.

5.4. Empirical game theoretical analysis

The experimental studies analyze the strategy performance from the common competition perspective (i.e., in terms of the mean scores achieved by the agents). This analysis is limited in that it does not provide a clear indication of the robustness of these strategies. In particular, it does not answer the question whether a strategy is so robust that it keeps its performance level in negotiation settings where the opponents are allowed to deviate, i.e., to switch to another strategy in search for a better outcome. To address this robustness criterion, empirical game theory (EGT) analysis (Jordan et al., 2007), which was initially developed to analyze the Trading Agent Competition (TAC), is applied to the competition results. We consider the statistically significant deviations as in Williams et al. (2011), where there is an incentive for an agent to unilaterally change its strategy in order to achieve a statistically significant increase of its own profit. The aim of using EGT is to search for pure Nash equilibria, in which no agent has an incentive to deviate from its current strategy.

We investigated strategy robustness by means of EGT in the two following scenarios:

1. a single negotiation between two players, and
2. a tournament consisting of six agents, each using one of the top three strategies reported in Table 3.

The former scenario captures bilateral agent-agent negotiation situations (i.e., with only two players participating in the game), and the latter captures such a kind of negotiation in the context of

Table 4
Overall performance averaged across all domains with the best agents of ANAC 2012 included.

| Agent | Mean | 95% Confidence interval | |
|------------------------|--------------|-------------------------|--------------|
| | | Lower bound | Upper bound |
| EMAR | 0.551 | 0.541 | 0.561 |
| OMACagent | 0.542 | 0.529 | 0.556 |
| CUHKAgent | 0.540 | 0.527 | 0.576 |
| AgentLG | 0.539 | 0.526 | 0.552 |
| TheNegotiator reloaded | 0.537 | 0.523 | 0.551 |
| BRAMAgent_2 | 0.503 | 0.483 | 0.522 |

Table 5
Agent-pair payoff matrix, where the score pair in each entry is averaged over all domains, with the first score representing the row player and the second for the column player. (The first letter of each agent, in bold, is used as the identifier.)

| Strategy | EMAR | Gahboninho | BRAMAgent | HardHeaded | IAMhaggler2011 | Agent_K2 |
|----------------|---------------|---------------|---------------|---------------|----------------|----------------|
| EMAR | (0.542,0.542) | (0.673,0.508) | (0.655,0.382) | (0.640,0.532) | (0.796,0.444) | (0.684, 0.487) |
| Gahboninho | (0.508,0.673) | (0.639,0.639) | (0.645,0.468) | (0.539,0.592) | (0.790,0.526) | (0.662,0.581) |
| BRAMAgent | (0.382,0.655) | (0.468,0.644) | (0.544,0.544) | (0.556,0.613) | (0.661,0.604) | (0.654,0.548) |
| HardHeaded | (0.533,0.640) | (0.592,0.539) | (0.613,0.556) | (0.578,0.578) | (0.748,0.526) | (0.683,0.495) |
| IAMhaggler2011 | (0.444,0.796) | (0.526,0.790) | (0.604,0.661) | (0.526,0.748) | (0.658,0.658) | (0.663,0.708) |
| Agent_K2 | (0.487,0.684) | (0.581,0.662) | (0.548,0.654) | (0.495,0.683) | (0.708,0.664) | (0.627,0.627) |

tournaments (i.e., with more than two players being involved). In the following, we do the game-theoretic analysis on the basis of the primary competition results shown in Section 5.2. The initial letter of each strategy is used as the strategy identifier (e.g., **H** means Hardheaded). The strategy set, S , is thus given by $S=\{\mathbf{E}, \mathbf{G}, \mathbf{B}, \mathbf{H}, \mathbf{I}, \mathbf{A}\}$.

For the first scenario, each agent is allowed to choose one strategy from S . We define a *profile* as the two strategies used by the players in the game (note that they may use the same strategy). Furthermore, the *score* of a specific strategy in a particular profile is calculated as its averaged payoff achieved when playing against the other strategy. The payoff matrix is given in Table 5, where each entry is a pair of scores composed of the score for the row player and the score of the column player. Using this payoff matrix, the results are generated and depicted in Fig. 6. The first row of each node represents a strategy profile being a mix of two strategies from S , and the second row shows the average score of the two strategies. This score is used as a measure of the social welfare achieved by the strategies, which can represent the overall benefit brought by a profile to the whole set of participants. An arrow indicates the statistically significant deviation between strategy profiles. For instance, the most left arrow in the figure means that there is a switch from the strategy profile (**H,I**) to (**E,I**); as a consequence, the agent deviating from strategy **H** (with utility 0.748 according to Table 5) to **E** (with utility 0.796) achieves a utility increase of 0.048. According to the EGT analysis, there exists one pure Nash equilibrium, namely, the strategy profile (\mathbf{E}^* versus \mathbf{E}^*). The equilibrium thus contains only EMAR. From this it follows that for any non-Nash equilibrium strategy profile there exists a path of statistically significant deviations (strategy changes) that leads to this equilibrium. This shows that EMAR is the most robust strategy in this scenario.

As a second scenario, we consider tournaments consisting of six agents, where each agent is permitted to choose one of the top three strategies shown in Table 3. The results of the EGT analysis of these tournaments are shown in Fig. 7. Here a *profile* is defined as the mixture of strategies used by the six players in a tournament. The nodes consist of two rows: the top row shows the set of strategies selected by agents in the tournament, and the second row shows the number of agents choosing each strategy. As can be seen from Fig. 7, just like in the previous scenario there only exists one pure Nash equilibrium and in this equilibrium all agents use EMAR. This means, in particular, that for any non-Nash equilibrium strategy profile there exists a path of statistically significant deviations leading to this equilibrium. With that the results for the second scenario also confirm the high robustness of EMAR.

We also conducted the EGT analysis for the unrestricted scenario of 6-agent tournaments with *all* strategies. The achieved results confirm the outcome of the two simpler scenarios described above. Due to the large number of profiles involved in the unrestricted scenario, a clear visualization of the results in an EGT graph is not possible (this graph contains $\binom{|p|+|s|-1}{|p|} = \binom{11}{6} = 462$ distinct nodes, where $|p|$ means the number of players and $|s|$ the number of strategies).

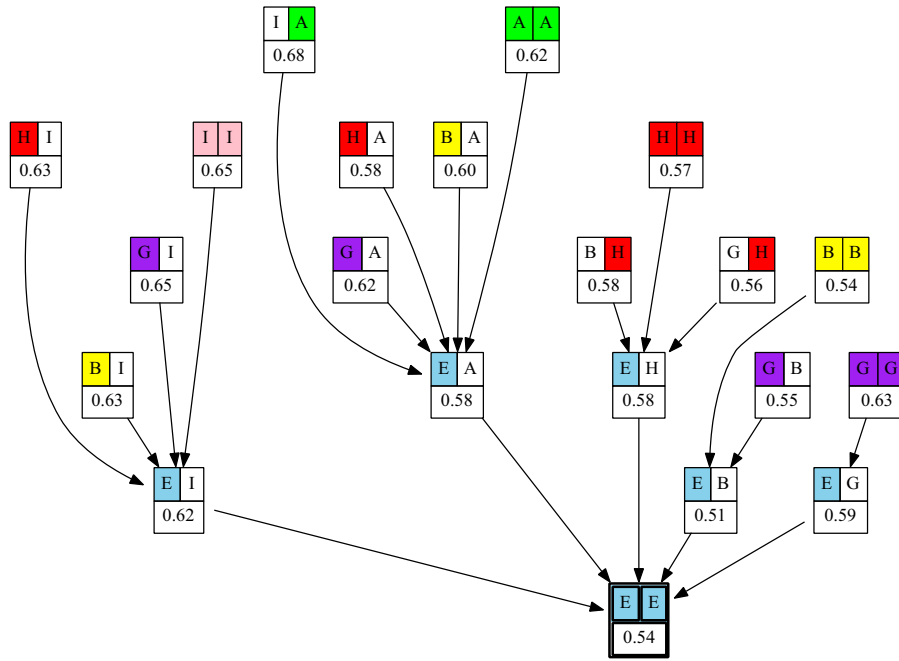


Fig. 6. Deviation analysis for the two-player negotiation. Each node shows a strategy profile and the average score of the two involved strategies, where the higher scoring strategy is marked by a colored background. An arrow indicates a statistically significant deviation between strategy profiles. The node with no outgoing arrow is the only equilibrium. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this paper.)

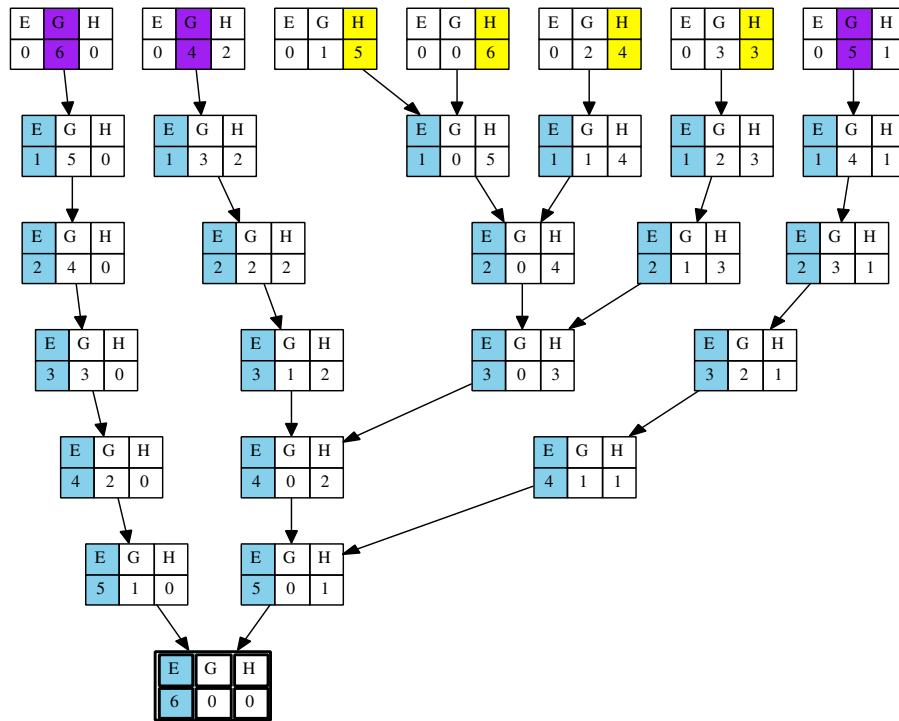


Fig. 7. Deviation analysis for six-player tournaments composed of the three top strategies. Each node shows a strategy profile, where the colored strategy in each node represents the one that achieved the highest score. An arrow indicates a statistically significant deviation between strategy profiles. The node with no outgoing arrow is the only equilibrium. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this paper.)

6. Conclusion

This paper introduced an effective strategy called EMAR for automated bilateral negotiation in complex environments (multi-issue, time-constrained, unknown opponents, and no prior domain knowledge). EMAR outperforms state-of-the-art agents

chosen from the International Automated Negotiation Agents Competition 2011 and 2012, both from the perspective of mean-score analysis and EGT analysis.

Research described in this paper opens several interesting research avenues and questions, and we consider the following three questions as most promising. First, are there opponent

modeling techniques which are more efficient than the one used by EMAR under the considered negotiation framework? Techniques that can be considered here are, for instance, Gaussian processes or GMDH networks (Madala and Ivakhnenko, 1994). As regards Gaussian processes, recently Chen et al. (2013) proposed a method using sparse pseudo-input Gaussian processes to reduce the computational complexity of an opponent modeling. To compare this method and EMAR w.r.t. their effectiveness and applicability under different negotiation circumstances can yield valuable insights on automated bilateral negotiation, and this comparison is on our current research agenda. Next, is it possible to extend opponent modeling of EMAR, which focuses on modeling the opponents' negotiation strategies, toward modeling of the opponents' preferences? We believe that such an extension would lead to a significant increase in negotiation power. And, last but not least, how can an agent transfer knowledge it gained in a negotiation task or domain to other, possibly new negotiation tasks and domains? In our current work we approach this question from the perspective of transfer reinforcement learning.

Acknowledgments

We are grateful to the China Scholarship Council (CSC) for providing a PhD scholarship to Siqi Chen. Moreover, we greatly appreciate the fruitful discussions with various members of the DKE SwarmLab (<http://swarmlab.unimaas.nl/>). Special thanks also goes to the anonymous reviewers of this paper for their valuable comments.

References

- Adhau, S., Mittal, M., Mittal, A., 2012. A multi-agent system for distributed multi-project scheduling: an auction-based negotiation approach. *Engineering Applications of Artificial Intelligence* 25, 1738–1751.
- ANAC, 2012. (<http://anac2012.ecs.soton.ac.uk/>).
- Baarslag, T., Fujita, K., Gerding, E.H., Hindriks, K.V., Ito, T., Jennings, N.R., Jonker, C.M., Kraus, S., Lin, R., Robu, V., Williams, C.R., 2013. Evaluating practical negotiating agents: results and analysis of the 2011 international competition. *Artificial Intelligence* 198, 73–103.
- Box, G., Jenkins, G.M., Reinsel, G.C., 1994. *Time Series Analysis: Forecasting and Control*, third ed. Prentice-Hall.
- Bzostowski, J., Kowalczyk, R., 2006. Predicting partner's behaviour in agent negotiation. In: *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems*. ACM, New York, NY, USA, pp. 355–361.
- Carbonneau, R., Kersten, G.E., Vahidov, R., 2008. Predicting opponent's moves in electronic negotiations using neural networks. *Expert Systems with Applications* 34, 1266–1273.
- Chen, S., Ammar, H.B., Tuyls, K., Weiss, G., 2013. Optimizing complex automated negotiation using sparse pseudo-input Gaussian processes. In: *Proceedings of the 12th International Joint Conference on Autonomous Agents and Multi-Agent Systems*. ACM, Saint Paul, Minnesota, USA, pp. 707–714.
- Chen, S., Weiss, G., 2012a. An efficient and adaptive approach to negotiation in complex environments. In: *Proceedings of the 20th European Conference on Artificial Intelligence*. IOS Press, Montpellier, France, pp. 228–233.
- Chen, S., Weiss, G., 2012b. A novel strategy for efficient negotiation in complex environments. In: *Proceedings of the 10th German Conference on Multiagent System Technologies*. Springer-Verlag, Germany, pp. 68–82 (Best Paper Award).
- Coehoorn, R.M., Jennings, N.R., 2004. Learning on opponent's preferences to make effective multi-issue negotiation trade-offs. In: *Proceedings of the Sixth International Conference on Electronic Commerce*. ACM, New York, NY, USA, pp. 59–68.
- Dang, J., Huhns, M.N., 2006. Concurrent multiple-issue negotiation for internet-based services. *IEEE Internet Computing* 10, 42–49.
- Faratin, P., Sierra, C., Jennings, N.R., 1998. Negotiation decision functions for autonomous agents. *Robotics and Autonomous Systems* 24, 159–182.
- Flandrin, P., Rilling, G., Goncalves, P., Basics, I.E., 2004. Empirical mode decomposition as a filter bank. *IEEE Signal Processing Letters* 11, 112–114.
- Fujita, K., Ito, T., Baarslag, T., Hindriks, K.V., Jonker, C.M., Kraus, S., Lin, R., 2013. The Second Automated Negotiating Agents Competition (ANAC2011). *Studies in Computational Intelligence*, vol. 435. Springer, Berlin/Heidelberg, pp. 183–197.
- Hao, J., Leung, H., 2012. ABiNeS: an adaptive bilateral negotiating strategy over multiple items. In: *Proceedings of the 2012 IEEE/WIC/ACM International Conference on Intelligent Agent Technology (IAT 2012)*. Macau, China, pp. 95–102.
- Hendriks, M., 2011. *A Survey of Opponent Models in Automated Negotiation*. Technical Report, Delft University of Technology, The Netherlands.
- Hindriks, K., Jonker, C., Kraus, S., Lin, R., Tykhonov, D., 2009. Genius: negotiation environment for heterogeneous agents. In: *Proceedings of the Eighth International Joint Conference on Autonomous Agents and Multiagent Systems*, pp. 1397–1398.
- Hou, C., 2004. Predicting agents tactics in automated negotiation. In: *IEEE/WIC/ACM International Conference on Intelligent Agent Technology*. IEEE Computer Society, Los Alamitos, CA, USA, pp. 127–133.
- Huang, N.E., Shen, S.S., 2005. Hilbert–Huang Transform and its Applications. World Scientific.
- Huang, N.E., Wu, M.L., Qu, W., Long, S.R., Shen, S.S., 2003. Applications of Hilbert–Huang transform to non-stationary financial time series analysis. *Applied Stochastic Models in Business and Industry* 19, 245–268.
- Hunag, N.E., Shen, Z., Long, S.R., 1998. The empirical mode decomposition and the Hilbert spectrum for nonlinear and non-stationary time series analysis. *Proceedings of the Royal Society A*, 903–995.
- Jordan, P.R., Kiekintveld, C., Wellman, M.P., 2007. Empirical game-theoretic analysis of the TAC supply chain game. In: *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multi-Agent Systems*, pp. 1188–1195.
- Lai, G., Li, C., Sycara, K., Giampapa, J., 2004. Literature Review on Multi-attribute Negotiations. Technical Report, Robotics Institute, Carnegie Mellon University.
- Lau, R.Y., Li, Y., Song, D., Kwok, R.C.W., 2008. Knowledge discovery for adaptive negotiation agents in e-marketplaces. *Decision Support Systems* 45, 310–323.
- Lin, R., Kraus, S., Wilkenfeld, J., Barry, J., 2008. Negotiating with bounded rational agents in environments with incomplete information using an automated agent. *Artificial Intelligence* 172, 823–851.
- Ljung, L., 1999. *System Identification: Theory for the User*. Prentice-Hall PTR, Upper Saddle River, NJ.
- Madala, H.R., Ivakhnenko, A.G., 1994. *Inductive Learning Algorithms for Complex Systems Modeling*. CRC Press, Boca Raton.
- Park, S., Yang, S., 2008. An efficient multilateral negotiation system for pervasive computing environments. *Engineering Applications of Artificial Intelligence* 21, 633–643.
- Ragone, A., Noia, T., Sciascio, E., Donini, F., 2008. Logic-based automated multi-issue bilateral negotiation in peer-to-peer e-marketplaces. *Autonomous Agents and Multi-Agent Systems* 16, 249–270.
- Rubinstein, A., 1982. Perfect equilibrium in a bargaining model. *Econometrica* 50, 97–109.
- Saha, S., Biswas, A., Sen, S., 2005. Modeling opponent decision in repeated one-shot negotiations. In: *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems*. ACM, New York, NY, USA, pp. 397–403.
- Wang, M., Wang, H., Vogel, D., Kumar, K., Chiu, D.K., 2009. Agent-based negotiation and decision making for dynamic supply chain formation. *Engineering Applications of Artificial Intelligence* 22, 1046–1055.
- Williams, C., Robu, V., Gerding, E.H., Jennings, N.R., 2011. Using Gaussian processes to optimise concession in complex negotiations against unknown opponents. In: *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*. AAAI Press, pp. 432–438.
- Yu, L., Wang, S., Lai, K.K., 2008. Forecasting crude oil price with an EMD-based neural network ensemble learning paradigm. *Energy Economics* 30, 2623–2635.