Expert Systems with Applications 42 (2015) 2287–2304

Contents lists available at ScienceDirect

Expert Systems with Applications

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

An approach to complex agent-based negotiations via effectively modeling unknown opponents ☆

Siqi Chen *, Gerhard Weiss

Department of Knowledge Engineering, Maastricht University, The Netherlands

ARTICLE INFO

Article history:
Available online 7 November 2014

Keywords:
Multi-agent systems
Automated multi-issue negotiation
Opponent modeling
Multi-resolution wavelet analysis
Gaussian processes
Empirical game theory

ABSTRACT

Negotiation among computational autonomous agents has gained rapidly growing interest in previous years, mainly due to its broad application potential in many areas such as e-commerce and e-business. This work deals with automated bilateral multi-issue negotiation in complex environments. Although tremendous progress has been made, available algorithms and techniques typically are limited in their applicability for more complex situations, in that most of them are based on simplifying assumptions about the negotiation complexity such as simple or partially known opponent behaviors and availability of negotiation history. We propose a negotiation approach called OMAC⋆ that aims at tackling these problems. OMAC⋆ enables an agent to efficiently model opponents in real-time through discrete wavelet transformation and non-linear regression with Gaussian processes. Based on the approximated model the decision-making component of OMAC⋆ adaptively adjusts its utility expectations and negotiation moves. Extensive experimental results are provided that demonstrate the negotiation qualities of OMAC⋆, both from the standard mean-score performance perspective and the perspective of empirical game theory. The results show that OMAC⋆ outperforms the top agents from the 2012, 2011 and 2010 International Automated Negotiating Agents Competition (ANAC) in a broad range of negotiation scenarios.

1. Introduction

Agent-based negotiation is about computational autonomous agents that attempt to arrive at joint agreements in competitive consumer-provider or buyer–seller scenarios on behalf of humans (Jennings et al., 2001). As one of the most fundamental and powerful mechanisms for solving conflicts between parties of different interests, recent years have witnessed a rapidly growing interest in automated negotiation, mainly due to its broad application range in fields as diverse as electronic commerce and electronic markets, supply chain management, task and service allocation, and combinatorial optimization. As a result, agent-based negotiation brings together research topics of artificial intelligence, machine learning, game theory, economics, and social psychology (Chen, Hao, Weiss, Tuyls, & Leung, 2014).

Dependent on the assumptions made about the negotiating agents' knowledge and the constraints under which the agents negotiate, negotiation scenarios show different levels of complexity. The following assumptions, which are reasonable in view of real-world applications and which underly our work, induce high complexity and raise particular demands on the abilities of the negotiators. First, the agents have no usable prior information about their opponents – neither about their preferences (e.g., their preferences over issues or their issue value ordering) nor about their negotiation strategies. Then, the negotiation is constrained by the amount of time being elapsed, the participants therefore do not know at any time during negotiation how many negotiation rounds there are left and they have to take into account at each time point (i) the remaining chances for offer exchange and (ii) the fact that the profit achievable through an agreement decreases over time ("negotiation with deadline and discount"). Third, each agent has a private reservation value below which an offered contract is not accepted.[1] Thereby we adopt the common view that an

☆ This article is a substantially extended version of our ECAI main track paper (Chen & Weiss, 2012). The extension primarily concerns the negotiation approach (OMAC) described in the ECAI paper. An effective negotiation approach called OMAC⋆ is proposed that advances OMAC in three significant ways. Furthermore, a comprehensive experimental analysis, as well as a useful game-theoretical robustness analysis, is presented. In the experimental analysis a large number of negotiation scenarios are considered and a comparison with the best agents from recent editions of ANAC competitions as well as with the predecessor of OMAC⋆ (i.e., OMAC) is provided.

* Corresponding author at: Department of Knowledge Engineering, Maastricht University, P.O. Box 616, 6200 MD Maastricht, The Netherlands. Tel.: +31 433883909.

E-mail addresses: siqi.chen@maastrichtuniversity.nl (S. Chen), gerhard.weiss@maastrichtuniversity.nl (G. Weiss).

[1] This reservation value is also called utility of conflict or disagreement solution.

http://dx.doi.org/10.1016/j.eswa.2014.10.048

agent obtains the reservation value even if no agreement is reached in the end. This implies that breaking-off a negotiation session would be potentially beneficial especially when the time-discounting effect is substantial and the other side is being very tough. Together these assumptions make negotiations complicated (yet realistic), where efficiently reaching agreements are particularly challenging. We thus refer to such type of negotiations as *complex negotiations* afterwards.

Although there exist many research efforts to address the problems of complex negotiations over the past years, two issues still stand out. The first one relates to learning unknown opponents' strategies. While it has been realized by early work that a successful negotiation needs to be based in one way or another on learning opponent models, the existing learning approaches either are limited in their usage in complex negotiations due to the impractical assumptions made about the environment, or have low efficacy in modeling opponents. The other issue is the absence of a decision-making mechanism that is suited for complex negotiations (i.e., the way of how to concede towards opponents in the course of negotiation). The strategies available to complex negotiation tend to consider concession in an intuitive fashion, or neglect the problem of "irrational concession" (see Section 5.2). As a result, the current decision-making methods are not adaptive and effective to respond to the high uncertainty of complex negotiations.

Based on the above motivation, this work proposes a novel strategy called OMAC★ for complex negotiations to address the aforementioned two issues that could further improve performance of a negotiating agent. In particular, it extends the OMAC negotiation strategy, which we introduced in Chen and Weiss (2012), in several important aspects (as detailed in Section 2). The proposed approach manages to integrate two key aspects of a successful negotiation: efficient opponent modeling and adaptive decision-making. Opponent modeling realized by OMAC★ aims at predicting the utilities of opponent future counter-offers (for itself) and is achieved through two standard mathematical techniques known as discrete wavelet transformation (DWT) and Gaussian processes (GPs). Adaptive decision-making realized by OMAC★ consists of two components, namely, concession making and counter offer responding, and it employs the learnt opponent model to automatically adjust the concession behavior and the response to counter-offers from opponents.

The remainder of this article is structured as follows. Section 2 overviews important related work. Section 3 provides the negotiation environment that we have considered. Section 4 describes the main mathematical techniques exploited by OMAC★. Section 5 shows the technicalities of the proposed strategy. Sections 6 and 7 offer a careful empirical evaluation and game-theoretic analysis of OMAC★. Section 8 discusses some interesting experimental results and other related aspects of agent-based negotiation. Finally, Section 9 identifies some important research lines induced by the work.

2. Related work

Negotiation has traditionally been investigated in game theory (Osborne & Rubinstein, 1994; Raiffa, 1982) and in previous years it has also developed into a core topic of multiagent systems (e.g., Lopes, Wooldridge, & Novais, 2008; Mor, Goldman, & Rosenschein, 1996; Weiss, 2013). Numerous approaches to automated negotiation have been proposed that, like the one described in this work, explore the idea to equip an agent with the ability to build a model of its opponent and to use this model for optimizing its negotiation behavior. Modeling the opponent's behavior, however, is practically challenging because negotiators usually do not reveal their true preferences and/or negotiation strategies in order to avoid that others exploit this information to their advantage (e.g., Coehoorn & Jennings, 2004; Raiffa, 1982). Current methods however tend to make simplifying assumptions about the negotiation settings. For example, there are approaches that deal with single-issue negotiation and others that assume that the opponents have a rather simple (e.g., non-adaptive) behavior, or the negotiations take place in scenarios with a low dimension (e.g., a small number of issues and possible choices for each of them). In the following, representative model-based negotiation approaches are overviewed.

Many of the available approaches aim at learning opponents' preferences or the reservation value. Faratin, Sierra, and Jennings (2002) propose a trade-off strategy to increase the chance of getting own proposals accepted without decreasing the own profit. The strategy applies the concept of fuzzy similarity to approximate the preference structure of the opponent and uses a hill-climbing technique to explore the space of possible trade-offs for its own offers that are most likely to be accepted. The effectiveness of this method highly depends on the availability of prior domain knowledge that allows to determine the similarity of issue values. Coehoorn and Jennings (2004) propose a method using Kernel Density Estimation for estimating the issue preferences of an opponent in multi-issue negotiations. It is assumed that the negotiation history is available and that the opponent employs a time-dependent tactic (i.e., the opponent's concession rate depends on the remaining negotiation time, see, e.g., Faratin, Sierra, & Jennings (1998) for details on this kind of tactic). The distance between successive counter-offers is used to calculate the opponent's issue weights and to assist an agent in making trade-offs in negotiation. Some approaches use Bayesian learning in automated negotiation. For instance, Zeng and Sycara (1998) use a Bayesian learning representation and updating mechanism to model beliefs about the negotiation environment and the participating agents under a probabilistic framework; more precisely, they aim at enabling an agent to learn the reservation value of its opponent in single-issue negotiation. Another approach based on Bayesian learning is presented in Lin, Kraus, Wilkenfeld, and Barry (2008). Here the usage of a reasoning model based on a decision-making and belief-update mechanism is proposed to learn the likelihood of an opponent's profile; thereby it is assumed that the set of possible opponent profiles is known as a priori. Hindriks and Tykhonov (2008) present a framework for learning an opponent's preferences by making assumptions about the preference structure and rationality of its bidding process. It is assumed that (i) the opponent starts with optimal bids and then moves towards the bids close to the reservation value, (ii) its target utility can be expressed by a simple linear decreasing function, and (iii) the issue preferences (i.e., issue weights) are obtainable on the basis of the learned weight ranking. Moreover, the basic shape of the issue evaluation functions is restricted to downhill, uphill or triangular. In order to further reduce uncertainty in high-dimensional domains, issue independence is assumed to scale down the otherwise exponentially growing computational complexity. Oshrat, Lin, and Kraus (2009) developed an effective negotiating agent for effective multi-issue multi-attribute negotiations with both human counterparts and automated agents. The successful negotiation behavior of this agent is, to a large extent, grounded in its general opponent modeling component. This component applies a technique known as Kernel Density Estimation to a collected database of past negotiation sessions for the purpose of estimating the probability of an offer to be accepted, the probability of the other party to propose a bid, and the expected averaged utility for the other party. The estimation of these values plays a central role in the agent's decision making. While the agent performs well, the approach taken is not suited for the type of negotiation we are considering (real-time, no prior knowledge, etc.) because opponent modeling is done offline and requires knowledge about previous negotiation traces.