Engineering Applications of Artificial Intelligence 34 (2014) 45–57

Contents lists available at ScienceDirect

Engineering Applications of Artificial Intelligence

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

An efficient and robust negotiating strategy in bilateral negotiations over multiple items

Jianye Hao [a,d,*], Songzheng Song [b], Ho-fung Leung [c], Zhong Ming [d,**]

[a] Massachusetts Institute of Technology, United States
[b] Nanyang Technological University, Singapore
[c] The Chinese University of Hong Kong, Hong Kong Special Administrative Region
[d] Shenzhen University, China

ARTICLE INFO

Article history:
Received 29 December 2013
Received in revised form
21 April 2014
Accepted 8 May 2014
Available online 3 June 2014

Keywords:
Bilateral negotiation
Multiagent learning
Robustness
Model checking

ABSTRACT

Multi-item negotiations surround our daily life and usually involve two parties that share common or conflicting interests. Effective automated negotiation techniques should enable the agents to adaptively adjust their behaviors depending on the characteristics of their negotiating partners and negotiation scenarios. This is complicated by the fact that the negotiation agents are usually unwilling to reveal their information (strategies and preferences) to avoid being exploited during negotiation. In this paper, we propose an adaptive negotiation strategy, called *ABiNeS*, which can make effective negotiations against different types of negotiating partners. The *ABiNeS* strategy employs the non-exploitation point to adaptively adjust the appropriate time to stop exploiting the negotiating partner and also predicts the optimal offer for the negotiating partner based on the reinforcement-learning based approach. Simulation results show that the *ABiNeS* strategy can perform more efficient exploitations against different types of negotiating partners, and thus achieve higher overall payoffs compared with the state-of-the-art strategies under negotiation tournaments. We also provide a detailed analysis of why the *ABiNeS* strategy can negotiate more efficiently compared with other existing state-of-the-art negotiation strategies focusing on two major components. Lastly, we propose adopting the *single-agent best deviation* principle to analyze the robustness of different negotiation strategies based on model checking techniques. Through our analysis, the *ABiNeS* strategy is shown to be very robust against other state-of-the-art strategies under different negotiation contexts.

1. Introduction

Negotiation is a common and important approach to resolve conflicts and reach agreements between different parties in our daily life. With the advance and popularity of Web and e-commerce, it is expected that a lot of previous negotiation activities between humans will be moved to electric platforms and greatly benefit from automated negotiation techniques (Kersten and Noronha, 1999). Automated negotiation techniques can, to a large extent, alleviate the efforts of human negotiators, and also aid human in reaching better negotiation outcomes by compensating for the limited computational abilities of humans when they are faced with complex negotiations. Until now, a lot of automated negotiation strategies and mechanisms have been proposed in different negotiation scenarios (Faratin et al., 2003; Saha et al., 2005; Hindriks and Tykhonov, 2008; Jakub and Ryszard, 2006).

The major difficulty in designing automated negotiation agent is how to achieve optimal negotiation results given incomplete information on the negotiating partner. The negotiation partner usually keeps its negotiation strategy and its preference as its private information to avoid exploitations. A lot of research effort has been devoted to better understand the negotiation partner by either estimating the negotiation partner's preference profile (Zeng and Sycara, 1998; Hindriks and Tykhonov, 2008; Coehoorn and Jennings, 2004) or predicting its decision function (Zeng and Sycara, 1996; Jakub and Ryszard, 2006). On one hand, with the aid of different preference profile modeling techniques, the negotiating agents can get a better understanding of their negotiating partners and thus increase their chances of reaching mutually beneficial negotiation outcomes. On the other hand, effective strategy prediction techniques enable the negotiating agents to maximally exploit their negotiating partners and thus receive as much benefit as possible from negotiation. However, in most of previous work, the negotiating agents are usually assumed to be situated in a

* Corresponding author at: Massachusetts Institute of Technology, United States.
** Corresponding author.
E-mail addresses: jianye@mit.edu (J. Hao), songsz@nut.edu.sg (S. Song), lhf@cuhk.edu.hk (H.-f. Leung), mingz@szu.edu.cn (Z. Ming).

http://dx.doi.org/10.1016/j.engappai.2014.05.008
0952-1976/

negotiation environment that strong limitations are put on the negotiation scenario or the negotiation opponent. For example, some work assumes (Saha et al., 2005) that the agents negotiate over a single item only, however practical negotiation scenarios usually involve multiple items to negotiate over. There also exists some work that assumes that the negotiation agents can have access to their opponents' (partial) preferences. This can be unrealistic especially in multi-issue negotiation scenarios in which the preferences of different agents may vary significantly, and the agents usually would not like to disclose their preferences to avoid being exploited. Another assumption commonly adopted (Jakub and Ryszard, 2006) is that the negotiation opponent is limited to choose from a specific set of simple strategies, e.g., time-dependent or behavior-dependent tactics. Those strategies designed under this assumption may not work well against other negotiation partners with more complex state-of-the-art strategies.

To this end, in recent years a number of advanced negotiation strategies taking advantage of existing techniques have been proposed and agents employing these strategies have participated in *automated negotiating agents competition* (*ANAC*) (Baarslag et al., 2010, 2013). The ANAC competition provides a negotiation platform which enables different negotiation agents to be evaluated against a wide range of opponents within a realistic negotiation environment. During the past three years, dozens of state-of-the-art negotiation strategies have been extensively evaluated in a variety of multi-issue negotiation scenarios and valuable insights have been obtained in terms of the advantages and disadvantages of different techniques, e.g., the efficacy of different acceptance conditions (Baarslag et al., 2011). It is still an open and interesting problem to design more efficient automated negotiation strategies against a variety of negotiating opponents in different negotiation domains.

In this paper, we propose an adaptive negotiation strategy *ABiNeS* for automated agents to negotiate in bilateral multi-issue negotiation environments following the settings adopted in ANAC 2012 (The Third International Automated Negotiating Agent Competition, 2012). Bilateral multi-issue negotiations surround people's daily life and have received lots of attention in the negotiation literature. During negotiation, both the agents' negotiation strategies and preference profiles are their private information, and for each agent the only available information about the negotiating partner are its past negotiation moves. Considering the diversity of the available negotiation strategies that the negotiating agents can adopt, it is usually very difficult (or impossible) to predict which specific strategy the negotiating partner is using based on this limited information. To effectively cope with different types of opponents, we introduce the concept of non-exploitation point λ to adaptively adjust the degree that an *ABiNeS* agent exploits its negotiating opponent. The value of λ is determined by the characteristics of the negotiation scenario and the concessive degree of the negotiating partner, which is estimated based on the negotiation history. Besides, to maximize the possibility that the offer the *ABiNeS* agent proposes will be accepted by its negotiating partner, it can be useful to make predictions on the preference profile of the negotiating partner. Instead of explicitly modeling the negotiation partner's preference profile, we propose a reinforcement-learning based approach to determine the optimal proposal for the negotiating partner based on the current negotiation history.

We evaluate the performance of the *ABiNeS* strategy compared with a number of state-of-the-art negotiation strategies from two different perspectives: *efficiency* in terms of the average payoff obtained under a particular negotiation tournament and *robustness* in terms of how likely the agents have the incentive to adopt our strategy rather than other strategies. First, the *efficiency* evaluation is conducted under the negotiation tournament setting following ANAC 2012 using GENIUS[1] (Lin et al., 2012) platform. Simulation results show that the *ABiNes* strategy can make more effective exploitations against a variety of negotiation partners and thus obtain higher average payoffs during negotiation tournaments and it is worth mentioning that the *ABiNes* strategy wins the champion of ANAC 2012 known as *CUHKAgent*. Second, we give a detailed analysis of the *ABiNes* strategy by investigating the influence of its two major novel components. Through the detailed analysis, we aim at providing a clear understanding of why the *ABiNes* strategy can win the champion of ANAC 2012, and more importantly, offering valuable insights for the automated negotiation community for the future negotiation strategy design. Third, we propose adopting the *single-agent best deviation* principle to analyze the robustness of different negotiation strategies based on model checking techniques. According to the *robustness* analysis, the *ABiNes* strategy is shown to be very robust against other state-of-the-art strategies under different negotiation contexts.

The remainder of the paper is organized as follows. In Section 2, we give a description of negotiation model we consider in this paper. In Section 3, the negotiation strategy *ABiNes* we propose is introduced. In Section 4, we give detailed evaluation of the negotiation *efficiency* and *robustness* of *ABiNes* compared with the state-of-the-art negotiation strategies under different negotiation contexts. An overview of related work on automated negotiation strategies is given in Section 5. Lastly conclusion and future work are given in Section 6.

2. Negotiation model

In this section, we describe the negotiation model we consider in this work, which follows the settings adopted in ANAC 2012 (The Third International Automated Negotiating Agent Competition, 2012). We focus on bilateral negotiations, i.e., negotiations between two agents. Specifically, the alternating-offers protocol is adopted to regulate the interactions between the negotiating agents, in which the agents take turns to exchange proposals. For each negotiation scenario, both agents can negotiate over multiple issues (items), and each item can have a number of different values. Let us denote the set of items as $\mathcal{M}$, and the set of values for each item $m_i \in \mathcal{M}$ as $\mathcal{V}_i$.[2] We define a negotiation outcome ω as a mapping from every item $m_i \in \mathcal{M}$ to a value $v \in \mathcal{V}_i$, and the negotiation domain is defined as the set Ω of all possible negotiation outcomes. For each negotiation outcome ω, we use $\omega(m_i)$ to denote the corresponding value of the item m_i in the negotiation outcome ω. We assume that the knowledge of the negotiation domain is known to both agents beforehand, and is not changed during the whole negotiation session.

For each negotiation outcome ω, different agents may have different preferences. Here we assume that each agent i's preference can be modeled by a utility function u_i such that $\forall \omega \in \Omega$, it is mapped into a real-valued number in the range of [0,1], i.e., $u_i(\omega) \in [0, 1]$. In practical negotiation environments, it is usually associated with certain cost in each negotiation. To take this factor into consideration, a real-time deadline is imposed on the negotiation process and each agent's actual utilities over the negotiation outcomes are decreased by a discounting factor δ over time. Following the setting adopted in ANAC'12, each negotiation session is allocated 3 min, which is normalized into the range of [0,1], i.e., $0 \le t \le 1$. Formally, if an agreement is reached at time t before the deadline, each agent i's actual utility function $U_i^t(\omega)$ over this

[1] GENIUS is short for General Environment for Negotiation with Intelligent multi-purpose Usage Simulation.

[2] Here $\mathcal{V}_i$ can be either discrete values or continuous real values.