

Negotiating with Unknown Opponents Toward Multi-lateral Agreement in Real-Time Domains

Siqi Chen, Jianye Hao, Shuang Zhou and Gerhard Weiss

Abstract Automated negotiation has been gained a mass of attention mainly because of its broad application potential in many fields. This work studies a prominent class of automated negotiations – multi-lateral multi-issue negotiations under real-time constraints, where the negotiation agents are given no prior information about their opponents’ preferences over the negotiation outcome space. A novel negotiation approach is proposed that enables an agent to obtain efficient agreements in this challenging multi-lateral negotiations. The proposed approach achieves that goal by, (1) employing sparse pseudo-input Gaussian processes (SPGPs) to model opponents, (2) learning fuzzy opponent preferences to increase the satisfaction of other parties, and (3) adopting an adaptive decision-making mechanism to handle uncertainty in negotiation.

1 Introduction

Negotiation is ubiquitous in our daily life and serves as an important approach to facilitate conflict-resolving and reaching agreements between different parties.

This paper is a shortened version of our previous work [6].

S. Chen (✉)

School of Computer and Information Science, Southwest University,
Chongqing, China
e-mail: siqi.chen09@gmail.com

J. Hao

School of Software, Tianjin University, Tianjin, China
e-mail: jianye.hao@tju.edu.cn

S. Zhou · G. Weiss

Department of Knowledge Engineering, Maastricht University,
Maastricht, The Netherlands
e-mail: shuang.zhou@maastrichtuniversity.nl

G. Weiss

e-mail: gerhard.weiss@maastrichtuniversity.nl

© Springer International Publishing AG 2017

K. Fujita et al. (eds.), *Modern Approaches to Agent-based Complex Automated Negotiation*, Studies in Computational Intelligence 674, DOI 10.1007/978-3-319-51563-2_17

219

Development of automated negotiation techniques enables software agents to perform negotiations on behalf of human negotiators. This can not only significantly alleviate the huge efforts of human negotiators, but also aid human in reaching better negotiation outcomes by compensating for the limited computational abilities of humans when they deal with complex negotiations [13].

During negotiations, an agent usually keeps its strategy and preference as private information, in order to avoid possible exploitation. Thus one major research challenge is to effectively estimate the negotiation partner's preference profile [2, 10, 14, 21] and predicate its decision function [15, 20]. On one hand, through getting a better understanding of partners' preferences, it would increase the possibility of reaching mutually beneficial outcomes. On the other hand, with effective strategy prediction it enables negotiation agents to maximally exploit their negotiating partners and thus receive as much benefit as possible [9]. Until now, fruitful research efforts have been devoted to developing automated negotiation strategies and mechanisms in a variety of negotiation scenarios [4, 7, 13–15, 19]. However, most research efforts have been devoted to bilateral negotiation scenarios, which only models the strategic negotiation among two parties. However, in real life the more common and general way of negotiations usually involve multiple parties. It is in common agreement from the automated negotiation research community that more attention should be given to multilateral negotiations and investigate effective negotiation techniques for multilateral negotiation scenarios.

In this work, a novel negotiation approach is proposed for intelligent agents to negotiate in multilateral multi-issue real-time negotiation. During negotiation, the agents' negotiation strategies and preference profiles are their private information, and the available information about the negotiating partner is its past negotiation moves [12]. Due to the huge strategy space that a negotiating partner can consider, it is usually hard to predict which specific strategy the agent is employing based on the very limited amount of information. Toward this end, instead of predicting the exact negotiation strategies of the opponents, we adaptively adjust the non-exploitation point λ to determine the perfect timing that we should stop further exploits the opponents, and then determine the aspiration level (or the target utility) for proposing offers to opponents before and after the non-exploitation point following different rules. The value of λ is determined as the timing when the estimated expected future utility we can obtain over all opponents is maximized. The future utility that each opponent offers can be efficiently predicted using the Sparse Pseudo-inputs Gaussian Process (SPGP) technique by dividing the negotiation history into a number of atomic intervals.

Given the aspiration level for offering proposals, another important question is how should we select an optimal proposal to reach efficient agreements with other parties, which can also improve the possibility of accepting this offer by the negotiating partners. In this work, we measure the efficient degree of an outcome from a practical perspective – the social welfare of participants. We propose modeling the preferences of each opponent using the least square error regression technique on the basis of negotiation history. After that, the offer with the highest social welfare is selected as the offer to be proposed with certain exploration.

The remainder of the work is organized as follows. Section 2 introduces the multilateral negotiation model given in this work. In Sect. 3, our negotiation approach is then introduced in details. Finally, conclusion and future work are given in Sect. 4.

2 Multilateral Negotiation Model

To govern the complex process of a multilateral negotiation, we adopt an extension of a basic bilateral negotiation protocol [17] which is widely used in the agents field [5, 6, 8, 11, 19]. The participating agents try to establish a contract for a product (service) or reach consensus on certain matter on behalf of their parties [6]. Precisely, let $A = \{a_1, a_2, \dots, a_i, \dots, a_m\}$ be the set of negotiating agents, J be the set of issues under negotiation with j a particular issue ($j \in \{1, \dots, n\}$ where m is the number of issues) [3, 5, 6, 8]. Following the alternating bargaining model of [17], each agent, in turn, has a chance to express its opinion about the ongoing negotiation. The opinion can be communicated in a form of a contract proposal (e.g., a new offer), or an acceptance of the latest offer on the table (note that previous offers would not be accepted once there exists a new proposal), or terminating the negotiation according to its interpretation of the current negotiation situation. A simple illustration of the




Fig. 1 Multilateral negotiation protocol

multilateral negotiation process is shown in Fig. 1. Due to space constraints we refer the interested reader to the work [1] for further details [3, 5, 6, 8].

An offer is taken as 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 follows:

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

where w_j^i and O are as defined above and V_j^i is the evaluation function of agent i for issue j , mapping every possible value of issue j (i.e., $O_{j,k}$) to a real number [3, 5, 6, 8]. The weight vector \mathbf{w} denotes the weighting preference of an agent, in which w_j^i represents its preference for issue j . The issue weights of an agent are normalized (i.e., $\sum_{j=1}^n w_j^i = 1$ for each agent i). In addition an agent has a lowest expectation for the outcome of a negotiation – the reservation value ϑ [3, 5, 6, 8].

In this work we consider negotiation being conducted in a real-time way. Each negotiator has a hard deadline by when it must have completed or withdraw the negotiation [3, 5, 6, 8]. The negotiation deadline is simply denoted by t_{\max} . For domains where the value of agreements is discounted over time, the discounting factor δ ($\delta \in [0, 1]$) is defined to calculate the discounted utility as follows:

$$D(U, t) = U \cdot \delta^t \quad (2)$$

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 [3, 5, 6, 8].

3 Negotiation Approach

The proposed approach consists of three core components: deciding aspiration level, generating new offers and responding mechanism, all of which are described in detail in this section. We first give an overview of our approach shown in *Algorithm 1*. Following that, each step of *Algorithm 1* is explained in details.

3.1 Deciding Aspiration Level

Aspiration level indicates the target utility of an agent in the negotiation process. In order to respond to uncertainty in a negotiation where opponents' private information is unknown, the aspiration level is updated due to the environment (e.g., available negotiation time and discounting effect) and opponent behaviors. The agent can therefore predict opponent future moves to assist its decision by analyzing past moves

Algorithm 1 The overview of the proposed negotiation approach. Let t_c be the current time point, δ the time discounting factor, and t_{max} the deadline of negotiation. O_{opp} is the latest opponent offer, Ω_i the previous offers of opponent i and O_{own} a new offer to be proposed by our agent. χ is the time series including the average utilities over intervals. E denotes the expected utility of incoming counter-offers. λ is the non-exploitation time point and u' the target utility. \mathbf{W} denotes the set of learnt opponent weight vectors [3, 5, 6, 8].

```

1: Require:  $\vartheta, \delta, t_{max}$ 
2: while  $t_c \leq t_{max}$  do
3:    $O_{opp} \leftarrow receiveMessage$ ;
4:    $\Omega_i \leftarrow recordOfferSet(t_c, O_{opp}, i)$ ;
5:   if  $myTurn(t_c)$  then
6:     if  $updateModel(t_c)$  then
7:        $\chi \leftarrow preprocessData(t_c)$ 
8:        $E \leftarrow Predict(\chi, \Omega)$ ;
9:        $(\lambda, U_{min}) \leftarrow updateParas(t_c)$ ;
10:       $\mathbf{W} = updatePreferenceModels()$ ;
11:     end if
12:   end if
13:    $u' = getTargetUtility(t_c, E, \lambda)$ ;
14:    $O_{own} \leftarrow constructOffer(u', \mathbf{W})$ ;
15:   if  $isAcceptable(u'_c, O_{opp}, t_c, \delta)$  then
16:      $accept(O_{opp})$ ;
17:   else
18:      $checkTermination()$ ;
19:      $proposeNewBid(O_{own})$ ;
20:   end if
21: end while

```

of the opponent. The prediction technique we use here is a computationally efficient variant of standard Gaussian Processes (GPs) – Sparse Pseudo-inputs Gaussian Processes (SPGPs), which proves effective in negotiation context [8]. Another advantage of SPGPs over other type of regression techniques is that it not only provides accurate prediction but also the measure of confidence in the prediction.

Following the notation of GPs in [16], given a data set $\mathcal{D} = \{\mathbf{x}^{(i)}, y^{(i)}\}_{i=1}^n$ where $\mathbf{x} \in \mathbb{R}^d$ is the input vector, $y \in \mathbb{R}$ the output vector and m the number of available data points when a function is sampled according to a GP, so we write, $f(\mathbf{x}) \sim \mathcal{G} \mathcal{P}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}'))$, where $m(\mathbf{x})$ is the mean function and $k(\mathbf{x}, \mathbf{x}')$ the covariance function, fully specifying a GP. Learning in a GP setting involves maximizing the marginal likelihood of Eq. 3 [3, 5, 6, 8].

$$\log p(\mathbf{y}|\mathbf{X}) = -\frac{1}{2}\mathbf{y}^T (\mathbf{K} + \sigma_n^2\mathbf{I})^{-1} \mathbf{y} - \frac{1}{2} \log |\mathbf{K} + \sigma_n^2\mathbf{I}| - \frac{n}{2} \log 2\pi, \quad (3)$$

where $\mathbf{y} \in \mathbb{R}^{m \times 1}$ is the vector of all collected outputs, $\mathbf{X} \in \mathbb{R}^{m \times d}$ is the matrix of the data set inputs, and $\mathbf{K} \in \mathbb{R}^{m \times m}$ is the covariance matrix with $|\cdot|$ representing the determinant.

The problem with GPs is that maximizing Eq. 3 is computationally expensive due to the inversion of the covariance matrix $\mathbf{K} \in \mathbb{R}^{n \times n}$ where n is the number of data

points. For this specific reason we employ a fast and more efficient learning technique – SPGPs. The most interesting feature of SPGPs is that these approximators are capable of attaining very close accuracy in both learning and prediction to normal GPs with only a fraction of the computation cost. Using only a small amount of pseudo-inputs, SPGPs are capable of attaining very similar fitting and prediction results to normal GPs [3, 5, 6, 8].

When a counter-proposal from agent i arrives at time t_c , our agent records the time stamp t_c and the utility $U(O^i)$ that is evaluated in our agent’s utility space. To reduce misinterpretation of the opponent’s behavior as much as possible that is caused by the setting of multi-issue negotiations, the whole negotiation is divided into a fixed number (denoted as ζ) of equal intervals. The average utilities at each interval with the corresponding time stamps, are then provided as inputs to the SPGPs. Results in [18] show a complexity reduction in the training cost (i.e., the cost of finding the parameters of the covariance matrix) to $\mathcal{O}(M^2N)$ and in the prediction cost (i.e., prediction on a new set of inputs) to $\mathcal{O}(M^2)$ [3, 5, 6, 8].

After learning a suitable model, SPGPs makes forecast about the future concession of the opponent as shown in line 7 of Algorithm 1. Our agent keeps track of the expected discounted utility based on the predictive distribution at a new input t_* , which is given by:

$$p(u_*|t_*, \mathcal{D}, \bar{\mathbf{X}}) = \int p(u_*|t_*, \bar{\mathbf{X}}, \bar{\mathbf{f}}) p(\bar{\mathbf{f}}|\mathcal{D}, \bar{\mathbf{X}}) d\bar{\mathbf{f}} = \mathcal{N}(u_*|\mu_*, \sigma_*^2), \quad (4)$$

where

$$\begin{aligned} \mu_* &= \mathbf{k}_*^T \mathbf{Q}_M^{-1} (\mathbf{A} + \sigma^2 \mathbf{I})^{-1} \mathbf{u} \\ \sigma_*^2 &= \mathbf{K}_{**} - \mathbf{k}_*^T (\mathbf{K}_M^{-1} - \mathbf{Q}_M^{-1}) \mathbf{k}_* + \sigma^2 \\ \mathbf{Q}_M &= \mathbf{K}_M + \mathbf{K}_{MN} (\mathbf{A} + \sigma^2 \mathbf{I})^{-1} \mathbf{K}_{NM} \end{aligned}$$

With given probability distribution over future received utilities and the effect of the discounting factor, the expected utility E_t is then formulated by

$$E_t = \frac{1}{C} \int_0^1 D(u \cdot p(u; \mu_t, \sigma_t), t) du \quad (5)$$

where μ_* and σ_* are the mean and standard deviation at time t_* , and the normalizing constant C is introduced to preserve a valid probability distribution [3, 5, 6, 8].

Our agent employs the target utility function as given in Eq. 6 to determine the aspiration level over time. The function adopts a tough manner (i.e., slowly conceding) before the non-exploitation time point (λ) for seeking higher expected profits, then it quickly goes to the expected minimal utility such that negotiation failure/disagreement could be avoided. λ is tweaked according to the behavior of the negotiation participants. More precisely, the higher the average opponent concession (measured in the our own utility space), the later our agent begins to compromise.

$$u' = \begin{cases} U_{\max} - \Delta \left(\frac{t_c}{\lambda}\right)^{1+\delta} & \text{when } t_c \leq \lambda, \\ (U_{\max} - \Delta) \left(1 - \frac{t_c - \lambda}{t_{\max} - \lambda}\right)^{1+\delta} & \text{otherwise} \end{cases} \quad (6)$$

where U_{\max} is the maximal utility, U_{\min} is the minimal utility ($U_{\min} = \max(\vartheta, \gamma)$) and γ the received lowest opponent concession), constant Δ is the maximal concession amount (i.e., $U_{\max} - U_{\min}$), with

$$\lambda = \operatorname{argmax}_{t \in T} \frac{1}{|A| - 1} \sum_{i \in A \setminus o} \frac{1}{C_i} \int_0^1 D_\delta(u \cdot p(u; \mu_i, \sigma_i), t) du \quad (7)$$

with o representing our agent and $T \in [t_c, t_{\max}]$.

3.2 Generating Offers

Given an aspiration utility level to achieve, our agent next needs to consider what offer to send such that the likelihood of an offer being accepted could be maximized. Performing this task would require certain knowledge about opponents' preferences. However, negotiation opponents unfortunately have no motivation to reveal their true likings over proposals (or their utility functions) to avoid exploitation. In order to address this problem, we model the opponent concession tactics as time-dependent tactics (originated in [11]) shown in Eq. 8, which are classic tactic in the current literature.

$$\tilde{u} = U_{\max} - (U_{\max} - \vartheta)(t_c/t_{\max})^\alpha \quad (8)$$

where α is the concession factor controlling the style of concessive behavior (for example, bouldware behavior ($\alpha < 1$) or concenter behavior ($\alpha > 1$)). Time-dependent tactics are widely used in automated negotiation community to decide concession toward opponents since a negotiator needs to make more or less compromise over time so as to resolve conflicts of the parties. In more detail, bouldware tactic maintains the target utility level until the late stage of a negotiation process, whereupon it concedes to the reservation utility. By contrast, concenter tactic makes quick compromise to other parties once a negotiation session starts. For linear tactic, it simply reduces the target utility from the maximal utility to the reservation utility in a linear way.

Learning opponent preferences, while useful, is indeed challenging because information about opponent preferences over different issues (e.g., the weight vector \mathbf{w}) is severely lacking. To tackle this issue, researchers typically assume that opponent concession tactic is fully known or preferences follow a certain distribution. In many real-world applications, it is however difficult or costly to acquire the exact information about opponent concession.¹ Therefore we make a mild assumption that we

¹Note that the opponent concession is the amount of concession measured in the utility space of the opponent instead of ours.




Fig. 2 A toy example of opponent concession ranges given by the pairs of concession factors (0.5, 2) at time 0.3, (0.5, 4) at 0.5 and (2, 0.25) 0.7, respectively

could enquire of domain experts about the approximate concession range of an opponent. This fuzzy knowledge is provided in form of a pair of concession factors that indicate the upper and lower concession an opponent makes at each time point. This idea is illustrated in Fig. 2. Thus, the agent can estimate opponent preferences with the aid of the fuzzy information about opponent concession. Specifically, the preferences are learnt through minimizing the loss function L , which gives the expected loss associated with estimating opponent concession based on a weight vector \mathbf{w} ; moreover an additional penalty is imposed by φ when an expected utility for \mathbf{w} exceeds the upper and lower bounds of opponent concession. When calculating an offer’s utility for opponent i , yet the valuation of each issue choice is needed. We here simply assume that the importance order of issue choices is known, and approximate the valuation like [14] as follow,

$$V_{j,k}^i(O_{j,k}) = \frac{2r_{j,k}^i}{K(K+1)} \tag{9}$$

where K is the number of possible choices for issue j , while $r_{j,k}^i$ denotes the ranking of the issue choice $O_{j,k}$.

Let the opponent utility of an offer for a weight vector \mathbf{w} be $\hat{u}_{\mathbf{w}}$. With the opponent concession tactic given in Eq. 8 and the two concession factors (which denote the

approximate concession range suggested by experts), our agent can estimate the weight vector of opponent i by means of linear least squares. This can be achieved by minimizing the following loss function,

$$L^i(\mathbf{w}) = \begin{cases} \left| \frac{(u_{upper}^i + u_{lower}^i)}{2} - \hat{u}_{\mathbf{w}} \right| + \varphi(u_{lower}^i, \hat{u}_{\mathbf{w}}), & \hat{u}_{\mathbf{w}} \leq u_{lower}^i \\ \left| \frac{(u_{upper}^i + u_{lower}^i)}{2} - \hat{u}_{\mathbf{w}} \right| + \varphi(\hat{u}_{\mathbf{w}}, u_{upper}^i), & u_{upper}^i \leq \hat{u}_{\mathbf{w}} \\ \left| \frac{(u_{upper}^i + u_{lower}^i)}{2} - \hat{u}_{\mathbf{w}} \right|, & \text{otherwise} \end{cases} \quad (10)$$

with u_{upper}^i and u_{lower}^i being the upper and lower bound of concession made by opponent i at time t , and φ the penalty function as below,

$$\varphi(x, y) = \beta |x - y|^{\frac{1}{2}} \quad (11)$$

where β denotes the confidence of the expert, and the lower the value, the more confidence the expert has about the prediction (to limit further complexity, we let β be 1).

After the estimation of weight vectors of other parties has been done, our agent chooses an offer being capable of maximizing the social welfare (e.g., the sum of the utility of all participants in the negotiation) given a aspiration level, shown as below:

$$\begin{aligned} & \underset{o}{\operatorname{argmax}} \frac{1}{|A| - 1} \sum_{i \in A \setminus o} (\hat{u}_{\mathbf{w}}^i(o) - \vartheta)^2 \\ & \text{subject to} \\ & U^o(o) \geq u' \end{aligned} \quad (12)$$

Although opponent preferences could be learnt on the basis of the provided concession tactics, it sometimes may be ineffective due to the fuzzy nature of the information; therefore our agent needs an alternative approach to choosing new offers. Fortunately, a real-time negotiation typically allows agents to exchange a large number of offers, thereby giving them many opportunities to explore the outcome space. Therefore, the proposed approach generates a new offer for next round following an ε -greedy strategy. The strategy selects either a greedy action (i.e., exploit) with $1 - \varepsilon$ probability ($\varepsilon \in [0, 1]$) or a random action with a probability of ε [3, 5, 6, 8]. It is worth noting that random action means choosing one offer from the set whose utility is above the given aspiration level by chance. The greedy action aims at choosing an offer that are expected to satisfy other sides' preferences most in order to improve their utilities over the negotiation outcome and the chance of the offer being accepted through fuzzy preference learning. With a probability $1 - \varepsilon$, the approach randomly picks one of those offer whose utility is equal or larger than the given aspiration level. In the latter case, the agent just chooses a new offer that has an utility within some range around u' .

3.3 Responding Mechanism

This responding mechanism of the proposed approach corresponds to lines 15–20 of Algorithm 1. After receiving a counter-proposal, the agent should decide whether to accept the proposal by checking two conditions. Firstly the agent validates whether the utility of the latest offer from opponents is higher than u' , while in the second, the agent needs to determine whether it had already proposed this very offer (i.e., the opponent's counter-offer) earlier. If either one of these two conditions is satisfied, the agent then accepts the offer as shown in line 16 and the negotiation will be completed if the proposal is also supported by the remaining agents [3, 5, 6, 8].

Moreover, when the negotiation situation becomes hard and might offer our agent a utility even lower than the reservation utility, the agent should consider whether to leave the negotiation course to receive the predefined reservation utility or not. Thus the reservation value is taken as an alternative offer from opponents with a constant utility. Thus the agent needs to check if the aspiration utility is smaller than the reservation utility. If positive, our agent is going to leave the negotiation table in the next round. If our agent decides neither to accept the latest counter-proposal nor to leave the negotiation, it proposes a new offer following the steps of lines 19 of Algorithm 1.

4 Conclusion

This work introduced a novel approach for multilateral agent-based negotiation in complex environments (i.e., multiple issues, real time-constrained, and unknown opponents). Our proposed approach, based on the adaptive decision-making scheme and the effective preference learning method, outclassed recent top ANAC agents. Empirical evaluation (see [6]) shows that our agent not only generates a higher mean individual utility but also leads to better social welfare compared to the state-of-the-art negotiation agents, and further game-theoretic analysis clearly manifests the high robustness of the proposed approach.

Acknowledgements This work is supported by National Natural Science Foundation of China (Program number: 61602391), and also by Southwest University and Fundamental Research Funds for the Central Universities (Grant number: SWU115032, XDJK2016C042).

References

1. R. Aydođan, D. Festen, K.V. Hindriks, C.M. Jonker, Alternating offers protocols for multilateral negotiation, in *Modern Approaches to Agent-based Complex Automated Negotiation*, ed. by K. Fujita, Q. Bai, T. Ito, M. Zhang, R. Hadfi, F. Ren, R. Aydođan (Springer, To be published)
2. T. Baarslag, K. Hindriks, C. Jonker, Acceptance conditions in automated negotiation, in *ACAN'11* (2011)

3. S. Chen, H.B. Ammar, K. Tuyls, G. Weiss, 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* (Saint Paul, Minnesota, ACM, USA, 2013), pp. 707–714
4. S. Chen, H.B. Ammar, K. Tuyls, G. Weiss, Using conditional restricted boltzmann machine for highly competitive negotiation tasks, in *Proceedings of the 23th International Joint Conference on Artificial Intelligence* (AAAI Press, 2013), pp. 69–75
5. S. Chen, J. Hao, G. Weiss, K. Tuyls, H.-F. Leung, Evaluating practical automated negotiation based on spatial evolutionary game theory, in *KI 2014: Advances in Artificial Intelligence*, Lecture Notes in Computer Science, vol. 8736, ed. by C. Lutz, M. Thielscher (Springer, New York, 2014), pp. 147–158
6. S. Chen, J. Hao, G. Weiss, S. Zhou, Z. Zhang, Toward efficient agreements in real-time multilateral agent-based negotiations, in *Proceedings of the 2015 IEEE 27th International Conference on Tools with Artificial Intelligence (ICTAI), ICTAI'15* (IEEE Computer Society, Washington, DC, USA, 2015), pp. 896–903
7. S. Chen, G. Weiss, An efficient and adaptive approach to negotiation in complex environments, in *Proceedings of the 20th European Conference on Artificial Intelligence*, vol. 242 (IOS Press, Montpellier, France, 2012), pp. 228–233
8. S. Chen, G. Weiss, An intelligent agent for bilateral negotiation with unknown opponents in continuous-time domains. *ACM Trans. Auton. Adapt. Syst.* **9**(3), 16:1–16:24 (2014)
9. S. Chen, G. Weiss, An approach to complex agent-based negotiations via effectively modeling unknown opponents. *Expert Syst. Appl.* **42**(5), 2287–2304 (2015)
10. R.M. Coehoorn, N.R. Jennings, Learning an opponent's preferences to make effective multi-issue negotiation trade-offs, in *ICEC'04* (2004), pp. 59–68
11. P. Faratin, C. Sierra, N.R. Jennings, Negotiation decision functions for autonomous agents. *Rob. Autom. Syst.* **24**(4), 159–182 (1998)
12. J. Hao, H.-F. Leung, ABiNeS: an adaptive bilateral negotiating strategy over multiple items, in *Proceedings of the 2012 IEEE/WIC/ACM International Conference on Intelligent Agent Technology* (IEEE Computer Society, Macau, China, 2012), pp. 95–102
13. J. Hao, S. Song, H.-F. Leung, Z. Ming, An efficient and robust negotiating strategy in bilateral negotiations over multiple items. *Eng. Appl. Artif. Intell.* **34**, 45–57 (2014)
14. K. Hindriks, D. Tykhonov, Opponent modeling in automated multi-issue negotiation using bayesian learning, in *AAMAS'08* (2008), pp. 331–338
15. B. Jakub, K. Ryszard, Predicting partner's behaviour in agent negotiation, in *AAMAS '06* (2006), pp. 355–361
16. C.E. Rasmussen, C.K.I. Williams, *Gaussian Processes for Machine Learning* (The MIT Press, Cambridge, 2006)
17. A. Rubinstein, Perfect equilibrium in a bargaining model. *Econometrica* **50**(1), 97–109 (1982)
18. E. Snelson, Z. Ghahramani, Sparse gaussian processes using pseudo-inputs, *Advances in Neural Information Processing Systems* (MIT press, Cambridge, 2006), pp. 1257–1264
19. C.R. Williams, V. Robu, E.H. Gerding, N.R. Jennings, Negotiating concurrently with unknown opponents in complex, real-time domains, in *ECAI'12*, pp. 834–839 (2012)
20. D. Zeng, K. Sycara, Bayesian learning in negotiation, in *AAAI Symposium on Adaptation, Co-evolution and Learning in Multiagent Systems*, pp. 99–104 (1996)
21. D. Zeng, K. Sycara, Bayesian learning in negotiation. *Int. J. Hum. Comput. Syst.* **48**, 125–141 (1998)