ANAC 2017: Repeated Multilateral Negotiation League

Reyhan Aydoğan[1,4], Katsuhide Fujita[2], Tim Baarslag[3], Catholijn M. Jonker[4] and Takayuki Ito[5]

[1] Department of Computer Science, Özyeğin University, Istanbul, Turkey
reyhan.aydogan@ozyegin.edu.tr
[2] Tokyo University of Agriculture and Technology Tokyo, Japan
katfuji@cc.tuat.ac.jp
[3] Centrum Wiskunde & Informatica (CWI) Amsterdam, The Netherlands
T.Baarslag@cwi.nl
[4] Interactive Intelligence Group Delft University of Technology Delft, The Netherlands
C.M.Jonker@tudelft.nl
[5] Nagoya Institute of Technology Nagoya, Japan ito.takayuki@nitech.ac.jp

Abstract. The Automated Negotiating Agents Competition (ANAC) is annually organized competition to facilitate the research on automated negotiation. This paper presents the ANAC 2017 Repeated Multilateral Negotiation League. As human negotiators do, agents are supposed to learn from their previous negotiations and improve their negotiation skills over time. Especially, when they negotiate with the same opponent on the same domain, they can adopt their negotiation strategy according to their past experiences. They can adjust their acceptance threshold or bidding strategy accordingly. In ANAC 2017, participants aimed to develop such a negotiating agent. Accordingly, this paper describes the competition settings and results with a brief description of the winner negotiation strategies.

1 Introduction

In recent years, much attention has been paid to automated negotiation in which intelligent agents negotiate on behalf of their users. It is a fruitful research area consisting of a variety of research challenges such as preference elicitation and representation [4], reasoning on incomplete [1] or complex preferences [29,25], designing and developing efficient negotiation strategies [11] and protocols [28], making decision under uncertainty [21], learning opponents' interest/preferences/ strategies [12,4] and so on. In order to bring together the researchers from negotiation community and to provide benchmarks for evaluation of the proposed negotiation strategies, an annual international competition on automated negotiation namely ANAC [22] has been organized for several years in conjunction with the most prestigious AI conferences such as the International Joint Conference on Autonomous Agents (AAMAS) and Multi-Agent Systems and the International Joint Conference on Artificial Intelligence (IJCAI). As the main organizers of this competition, our aim is to introduce new research challenges every year. Up to date, the focus was mostly on bilateral multi-issue negotiation. Recently, multilateral negotiation draws the attention; consequently, ANAC 2017 encouraged the participants to pursue developing effective negotiation strategies for repeated multilateral negotiations.

In comparison to bilateral negotiation, multilateral negotiations are more complex due to their nature – requiring to establish a joint consensus among more stakeholders with varying preferences [3]. In the competition, agents followed the Stacked Alternative Offer Protocol(SAOP) in which negotiating parties exchange their bids in a turn-taking fashion till reaching an agreement or reaching the predefined deadline [2]. The challenge in the competition is not only to negotiate with multiple opponents but also to learn from previous negotiations with the same opponents. In other words, the main goal is to design effective negotiation strategies (i.e. bidding, opponent modeling and acceptance strategies) when negotiating repeatedly with the same agents in a multilateral setting.

When an agent negotiates with its opponents first time, it may not have any idea about their strategies and preferences. Therefore, it can try to model their opponents by analyzing the current bid exchanges during the undergoing negotiation. However, when the agent negotiates with the same opponents repeatedly, it can act more strategically by exploiting its previous negotiation experiences (e.g. opponent modeling about preferences/strategies). Furthermore, since the agents will negotiate multiple times, negotiating like a Hardliner (i.e. being selfish) may cause your opponent not to concede next time; therefore, on average the agent may not gain high utility. Without doubt, it depends on to what extent the opponents take previous negotiation experiences into account in their strategies. Therefore, determining to what extent collaborative behavior to be adopted is another challenge in such a negotiation setting.

ANAC 2017 has been organized in conjunction with the IJCAI 2017 in Melbourne. 18 teams from 9 different institutions participated in the competition. Agents negotiate with the same opponents on the same negotiation scenarios five times and they can access historical data from their past negotiations. That is, they know their utility distribution over the exchanged offers in any previous negotiation session according to its own utility space and previous agreements they reached. There were two winner categories: individual utility gain and social welfare in terms of product of utilities. This paper presents competition setup and the evaluation of the ANAC 2017 agents elaborately.

The rest of the paper is organized as: Section 2 explains negotiation setup while Section 3 presents the results of the competition in detailed. Related work are provided in Section 4. Lastly, we conclude the paper with the potential research challenges for upcoming competitions in Section 5.

2 Negotiation Setup

In the competition, three agents negotiate on multiple issues to come up with an agreement. Each issue $j \in I$ can take a value v_j from a predefined set of valid values for that issue denoted by D_j (i.e., $v_j \in D_j$). Each agent can access this domain information. A bid $b = (b_1, ..., b_{|I|})$ is an assignment of values to all issues where $b_1 \in D_1$.

Preferences of negotiating parties are represented by means of additive utility function as shown in Equation 1 where $V_a(v_j)$ denotes agent a's valuation of the value for the issue j in the given bid and the weights of that issue is represented by $w_{a,j}$. In other words, agents sum up their weighted valuation of each issue value to calculate the overall utility. It is worth noting that it is forbidden to access opponents' preferences during the negotiation. Each agent can only access their own preferences.

$$u_a(b^t) = \sum_{j \in I} V_a(b_j^t) \cdot w_{a,j} \tag{1}$$

The interaction among agents is governed by the Stacked Alternative Offers Protocol (SAOP) [2]. According to this protocol, each agent i is expected to provide an action a_i in its turn. Type of an action is denoted as $type(a)$ which can take $type(a) = \{A^{bid}, A^{accept}, A^{end}\}$; that is, the agent can make a counter offer, accept the offer in the negotiation table or walk away from the negotiation respectively.

Following SAOP, the first agent starts the negotiation with an offer that is observed by all others immediately. Whenever an offer is made the next party in line can take the following actions:

- Make a counter offer (thus rejecting and overriding the previous offer)
- Accept the offer
- Walk away (e.g. ending the negotiation without any agreement)

This process is repeated in a turn taking clock-wise fashion until reaching an agreement or reaching the deadline. Negotiation is also terminated if one of the agents walks away. To reach an agreement, all parties should accept the offer. If no agreement has been reached at the end of the negotiation, the negotiation fails and agents receive their reservation utility. If an agreement on the bid b is reached, the utility of the agent i is estimated as explained above (See Equation 1).

In the competition, the deadline is set to three minutes and each negotiation session is repeated five times. Agents are allowed to access provided historical data from their past negotiations. The historical data involves the utility distribution of the exchanged offers in previous negotiation sessions according to agent's own utility space and previous agreements if there exist. Note that agents were not allowed to keep the entire history of exchanged bids in their past negotiation. However, they can process the bid exchanges during their ongoing negotiation and model their opponent's preferences and strategies. The opponent model can be saved and updated over time.

There were 18 submissions from 9 institutions:

- Maastricht University - NL
- Nagoya Institute of Technology - Japan
- Özyegin University - Turkey
- Southwest University - United States
- Technical University Of Crete - Greece
- Tianjin University - China
- Tokyo University of Agriculture and Technology - Japan
- University of Isfahan - Iran
- University of Southampton - UK

Participation to this competition requires submitting a negotiation scenario consisting of three conflicting preference profiles with domain description. In order to evaluate the performance of the agents, We picked eight negotiations scenarios with varying domain size from the submitted negotiation scenarios. Table 1 shows number of issues, the number of possible values for each issue and size of outcome size respectively.

Table 1. Negotiation Scenarios

Name	# of Issue	# of Values (1)	# of Values (2)	# of Values (3)	# of Values(4)	# of Values (5)	Scenario Size
GeneJack	4	10	10	10	10		10000
MyDomain	5	4	10	8	14	9	40320
Music	4	9	4	7	6		1512
SuperMarket	5	8	6	5	6	4	5760
Movietime	4	6	2	3	4		144
Lunch Time	3	5	4	4			80
Taxung	4	4	5	3	6		360
SmartGrid	3	5	4	4			80

In the qualification round, finalist agents are determined. Since there were 18 submissions, running whole tournament involving 18 agents in eight domains with 5 repetitions were not feasible within the given time. Therefore, two pools were generated randomly as follows:

- **Pool-1:** Farma17, ParsCat2, taxibox, Mamenchis, MadAgent, AgentKN, PonPokoAgent, SimpleAgent, Imitator
- **Pool-2:** AgentF, GeneKing, Mosa, CaduceusDC16, ParsAgent3, TucAgent, Group3, Gin, Rubick

Genius 7.1.4 [26], automated negotiation framework was used in the competition. Agents were evaluated according to two categories: *individual utility* and *product of utilities*. The top 4 performing agents in each pool proceeded to the final for each category. Therefore, there were eight finalists for each category after the qualification round.

3 Results of the Competition

The results of the qualification and final rounds are explained in the following parts.

3.1 Results of Qualification Round

Figure 1 and Figure 2 summarize the results of the qualification round for both Pool-1 and Pool-2 respectively. The highest four score for each category is highlighted.

According to the results of the first pool, the *Mamenchis*, *AgentKN*, *PonPokoAgent* and *SimpleAgent* qualified for the final round in the individual utility category. As far as the product of the utilities are concerned, *ParsCat2*, *taxibox*, *Mamenchis*, and *AgentKN* gained the highest product (above 0.47) and qualified for the final round.

In the second pool, *Agent F*, *CaduceusDC16*, *ParsAgent3* and *Rubick* agents are qualified for the final round according to the gained individual utility. In terms of the product of utilities, the agents with the highest products were *GeneKing*, *Mosa*, *ParsAgent3* and *Rubick*.

Agent Name	Individual Utility	Product of Utilities
Farma17	0.680307731	0.38284705
ParsCat2	0.681106633	0.474018267
taxibox	0.692776131	0.487286208
Mamenchis	0.735889079	0.494775099
MadAgent	0.671866389	0.437101437
AgentKN	0.724925194	0.473772255
PonPokoAgent	0.710761165	0.448005582
SimpleAgent	0.698954726	0.437494805
Imitator	0.623136568	0.442123768

Fig. 1. Results of Qualification Round for Pool-1

Agent Name	Individual Utility	Product of Utilities
AgentF	0.722600455	0.39668025
GeneKing	0.683409649	0.448725557
Mosa	0.70047027	0.436250542
CaduceusDC16	0.709912278	0.421041454
ParsAgent3	0.72097358	0.435728313
TucAgent	0.675038202	0.344809219
Group3	0.623009688	0.433441349
Gin	0.696220696	0.423431405
Rubick	0.742502203	0.426399043

Fig. 2. Results of Qualification Round for Pool-2

3.2 Results of Final Round

In the final round, we have separate pools for each evaluation category.

Individual Utility Category In this category, we ran a tournament with the following qualified agents:

- AgentF, Nagoya Institute of Technology
- AgentKN, Nagoya Institute of Technology
- CaduceusDC16, Özyeğin University & University of Southampton
- Mamenchis, Southwest University & Maastricht University & Tianjin University
- ParsAgent3, University of Isfahan
- PonPokoAgent, Tokyo University of Agriculture and Technology
- Rubick, Özyeğin University
- SimpleAgent, Özyeğin University

Figure 3 shows the average individual utility gained by each finalist over 1680 negotiations. Recall that each negotiating agent negotiates with all other agents five times for each negotiation scenario in Table 1. *PonPokoAgent* gained 0.75 on average and won the competition. *CaduceusDC16* and *Rubick* agents were awarded second and third place respectively.

Three main components of the negotiation agents contribute to their overall ranking: the bidding strategy, the opponent model and the acceptance strategy [7]. In the following part, the descriptions of the agents are given.

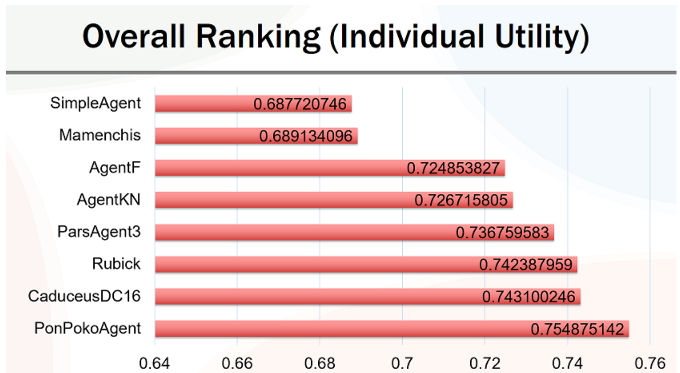

Fig. 3. Overall Ranking w.r.t. Individual Utility

Winners w.r.t. Individual Utility In this section, the detailed description of the winner strategies according to the individual utility, are given.

– ***PonPokoAgent by Takaki Matsune:*** PonPokoAgent tries to propose new offers to compromise the opponents offers as the time passes. However, it has several different bidding strategies and selects the one of them randomly because it is hard for the opponents to predict our agent strategy through the previous sessions. In addition, it tries to obtain the high individual utilities just before the deadline.
In the bidding strategy, it makes the list including all bids available in the domain at the beginning of the negotiation. It has 2 parameters to decide the highest/lowest utility of our agents offers, and 5 patterns decide these parameters (time means the negotiation time elapsed $[0, 1]$).
Pattern(1): $HIGHEST\_UTIL = 1.0 - 0.1 * time$
$LOWEST\_UTIL = 1.0 - 0.1 * time - abs(sin(40 * time))$
Pattern(2): $HIGHEST\_UTIL = 1.0 - 0.1 * time$
$LOWEST\_UTIL = 1.0 - 0.22 * time$

Pattern(3): $HIGHEST\_UTIL = 1.0 - 0.1 * time$
$LOWEST\_UTIL = 1.0 - 0.15 * time - abs(sin(20 * time))$

Pattern(4): $HIGHEST\_UTIL = 1.0 - 0.05 * time$
$LOWEST\_UTIL = 1.0 - 0.1 * time (if time <= 0.99)$
$LOWEST\_UTIL = 1.0 - 0.3 * time (if time > 0.99)$

Pattern(5): $HIGHEST\_UTIL = 1.0 - 0.15 * time * abs(sin(20 * time))$
$LOWEST\_UTIL = 1.0 - 0.21 * time * abs(sin(20 * time))$

It selects a pattern randomly at the beginning of the negotiation. It randomly selects the next offer between HIGHEST_UTIL and LOWEST_UTIL in utility from the list. In the acceptance strategy, it accepts the bid when the utility of the opponents bid is higher than LOWEST_UTILITY. It does not use the history information from previous negotiation sessions at all for machine learning.

– ***CaduceusDC16 by Taha Gunes and Emir Arditi:*** *CaduceusDC16* extends the Caduceus agent [20], which inspired from the ideas such as "*algorithm portfolio*", "*mixture of experts*", and "*genetic algorithm*". Simply, this agent asks the expert agents' opinion on whether to make a counter offer or to accept the given offer as well as what to bid. The expert agents are the ANAC agents succeeded in the previous years.

  When CaduceusDC16 needs to take an action, it asks the expert agents what to do, and if half or more of them states that the incoming offer should be accepted, it accepts the offer. If the majority suggests making a counter offer, then CaduceusDC16 combines bids from expert agents by picking the most favorable values of an issue by applying majority voting on each issue value. That is, the agent relies on the expert agents' bidding, opponent modeling and acceptance strategies.

  Furthermore, CaduceusDC16 aims to learn from its previous negotiations. It has an acceptance threshold and if the opponent's bids below this threshold, the agent immediately decides to make a counter offer irrespective of expert's suggestion (i.e. reject the opponent's bid). This threshold is updated over time based on the previous negotiation experience, particularly the utility gained in the previous negotiations. For instance, if the current threshold is 0.6 and the agent had an agreement with a utility of 0.8 in its previous negotiation, then it updates the threshold to 0.8.

– ***Rubick by Okan Tunalı:*** Rubick Agent is a complex time based conceder enriched by derivations of well studied heuristics in automated negotiation field. The main component of the agent is the target utility, which is actually lower boundary in the bid generation and in acceptance condition. If the history is not available yet, target utility is initialized as the utility of the first received bid and updated to the highest utility received from any of the opponent parties. On the other hand, if it detects a negotiation history with the same opponents, it sets the lower bound to be the highest utility ever received throughout the negotiation; thinking that the opponents will be designed in a myopic way. Both bid generation and acceptance strategies include randomness; they follow a Boulware strategy but the pace of concession is randomized. Technically, they sample from one sided Gaussian distributions whose standard deviation increase over time, increasing the likelihood of sending bids close to lower boundary. The opponent model resolves the bids into issue evaluation values and considering their occurrence frequencies, searches for bids that holds the target utility requirements while having the most common values. That is, it employs a frequency-based opponent modeling. Finally, the model keeps a list of bids accepted by only one of the opponents in the previous negotiations, which is sorted according to the gained utility by the Rubick agent. Elements of this list is used in bid generation only if there is almost no time left. The target utility is estimated according to the following formula:

target-utility=$max - receivedUtility + (1 - max - receivedUtility) * 1 - ((current - time^k) * |\frac{Gaussian()}{constant}|$

Product of Utilities Category In this category, we ran a tournament with the following qualified agents:

- AgentKN, Nagoya Institute of Technology
- GeneKing,Nagoya Institute of Technology
- Mamenchis, Southwest University & Maastricht University & Tianjin University
- Mosa, Southwest University
- ParsAgent3, University of Isfahan
- ParsCat2, University of Isfahan
- Rubick, Özyeğin University
- Tangxun, Nagoya Institute of Technology
- AgentKN, Nagoya Institute of Technology

Figure 4 shows the average product of utilities of the agreements reached by each agent. *ParsCat2* gained 0.509 on average and won the competition where *AgentKN* and *ParsAgent3* agents took the second and third place respectively (0.506 and 0.505). As seen from the results, the performance of these agents are very closed to each other. In the following part, the descriptions of those agents are given.

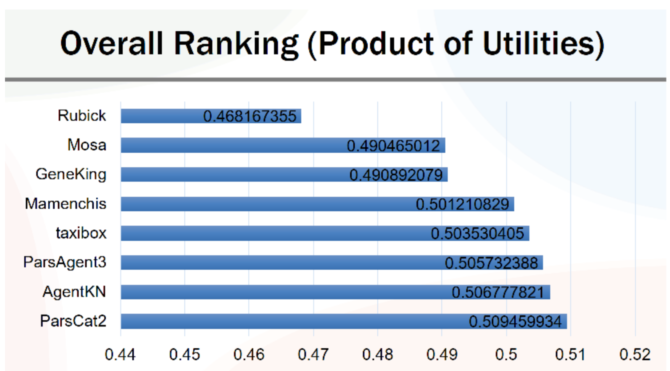

Fig. 4. Overall Ranking wrt Product of Utility

Winners w.r.t. Product of Utilities In this section, the detailed description of the winner strategies according to the product utilities, are given.

- ***ParsCat2 by Delaram Javdani and Maedeh Najar:*** ParsCat2 mirrors changes in agent's behavior by the passage of time. Thus, it is a time-dependent agent. The main idea is that it swings the offers because it may appear flexible taking the offer as the time passes. Therefore, it keeps the offer that has good utility and would use them again.

Acceptance Strategy The acceptance strategy of ParsCat2 is comprised of multiple parts base on time: Mainly, it makes the same offer as does this agent or more than that, accepts the offer. In other cases, it calculates their utility based on time. If the utility of the agent is larger than the estimated utility, it accepts the offer from the opponent; otherwise, makes a counter offer. At the beginning, it accepts an offer which has the utility of highest for this agent. By the passage of time, it cuts out on its expectation. On the run of time, reaching numbered seconds of the tournament, raising up the expectation again.

Bidding Strategy ParsCat2 is just going to make an offer which is good for this agent and it is good for others agents. It generates a bid, sets a threshold on the call of time, and checks whether the utility of the bid calculated are between the threshold. If not, it generates another bid and repeats it on. It sometimes on the chain of search of details not finding the bid with that utility, increases the interval. In addition, it doesn't offer a bid with the utility of less than 0.3. If the first mover, it offers the bid with maximum utility; otherwise, it makes sure the utility of the bid to offer is bigger than the offers the other agents made. If not, it makes a bid having the same utility with that offer. It offers with swinging by the time and reduces the average of its utility of offering.

– ***AgentKN by Keita Nakamura:*** AgentKN searches 10 bids that maximize self utility value while randomly shifting the initial bid. In bidding strategy, it decides the bid according to the self utility value and frequency that opponents have offered. First, it searches 10 bids that maximize self utility. Next, it sorts 10 bids by self utility and issue values, frequency, and offers. In searching the local-optimal bids, it uses Simulated Annealing, and searches 10 bids that maximize self utility value while randomly shifting the initial bid. It sorts by the following scores: $(utility) + 0.1^{(\log_{10} frequency+1)} * frequency$. $utility$ is the individual utility, $frequency$ is the sum of the number of values of issues opponents offered in the previous rounds.

  In acceptance strategy, it accepts when the utility value of the opponents bid exceeded the threshold at that time. The threshold is decided as the following equation: $threshold(t) = 1-(1-emax(t))*t^{\alpha}$ ($emax(t)$ is the estimated oppomnent's maximum utility, and α(¿ 1) is the parameter of consecting to the emax(t). It estimates max utility value that opponents may offer: $emax(t) = \mu(t)+(1-\mu(t))d(t)$ $d(t) = \frac{\sqrt{3}\sigma(t)}{\sqrt{\mu(t)(1-\mu(t))}}$

  ($\mu(t)$: Average of utility values that the opponent have offered, $d(t)$: Estimated width of the bid range of the opponent, which uses the concession strategy, $\sigma(t)$: Standard deviation of utility value that the opponent has offered)

– ***ParsAgent3 by Zahra Khosravimehr and Faria Nassiri-Mofakham:*** ParsAgent3 proposes a negotiation strategy in the multilateral negotiation domain. It is a concession-based negotiation strategy. It begins with a concession. If its opponents concede, it continues with more concessions, otherwise, it changes its strategy and behaves hard headed. To this end, ParsAgent3 first calculates its temporary target utility. It uses this value for using in bidding and acceptance strategies according to the following equations:

$$u(t) = P_{min} + (P_{max} - P_{min})(1 - F(t))$$

$$F(t) = k + (1 - k)t^{1/\beta}$$

In these equations, t is time, β indicates the agent's concedence speed during the negotiation process. The variables, P_{min} and P_{max} are respectively the minimum and maximum values of the function $F(t)$, so that P_{min}, P_{max}, $k \in [0, 1]$.
ParsAgent3 clusters every 100 bids received from each individual opponent, and estimates the temporary reservation value of each opponent [24]. This value is the minimum acceptable utility for the opponent at the given time. The agent then sets its reservation value according to the current estimated reservation values of all the opponents. If the opponents concede, the average of their reservation values is employed. Otherwise, the maximum value among the estimated reservation values of all the opponents is chosen as ParsAgent3 reservation value.
It calculates the difference between maximum and minimum utilities of the bids received from each opponent in the previous session using the history information. Then, if this value is less than 0.3, ParsAgent3 categorize the opponents as hard headed and so the agent will not concede in the current session.

4 Related Work

The automated negotiating agent competition follows in the footsteps of a series of successful competitions that aim to advance the state-of-the-art in artificial intelligence (other examples include the Annual Computer Poker Competition and the various Trading Agent Competitions, the Alimama International Advertising Algorithm Competition, and Angry Birds Competition). Over the years the competition has addressed various topics: varying the number of negotiators, the complexity of the negotiation domains, repeated negotiations with the same set of opponents, negotiations in special domains, and negotiating against humans.

Every year, new features have been incorporated into the competition environment to increase realism and to encourage the development of flexible and practical negotiation agents [5]. The first installment of ANAC was in 2010 and started as a jointly organized competition between Delft University of Technology and Bar-Ilan University [13]. Since the first edition, the local organization has rotated between various international institutions. In 2011, ANAC was organized by the 2010 winner, the Nagoya Institute of Technology [9]. Next editions were organized by University of Southampton (2012) [30], Ben Gurion University of the Negev (2013) [19], Nagoya Institute of Technology and Tokyo University of Agriculture and Technology (2014) [18], and Delft University of Technology (2015) [6].This paper describes the competitions from the next subsequent years. Every installment has seen a growing number of participants: ANAC 2010 started with 7 teams from 5 different institutes, while in 2017, it featured 18 teams from 9 institutes.

By analyzing the results of ANAC, the stream of the strategies of ANAC and important factors for developing the competition have been shown. Baarslag et al. [8]

presented an in-depth analysis of ANAC 2011. This paper mainly analyzes the different strategies using classifications of agents with respect to their concession behavior against a set of standard benchmark strategies and empirical game theory (EGT) to investigate the robustness of the strategies. It also shows that the most adaptive negotiation strategies, while robust across different opponents, are not necessarily the ones that win the competition. Furthermore, our EGT analysis highlights the importance of considering metrics.

Kawaguchi et al. [23] proposed a strategy for compromising the estimated maximum value based on estimated maximum utility. These papers have been important contributions for bilateral multi-issue closed negotiation; however, they don't deal with multi-time negotiation with learning and reusing the past negotiation sessions. This agent is the winner of ANAC2010. After that, Fujita [17,16] proposed the compromising strategy with adjusting the speed of making agreements using the Conflict Mode, and focused on multi-time negotiations. However, these strategies only focused on the linear utility functions.

In addition, some studies have focused on the divided parts of negotiating strategies in the alternating offering protocol: proposals, responses, and opponent modeling. Effective strategies can be achieved by combinations of top agents' strategies in the competitions depending on the opponent's strategies and negotiation environments. Many of the sophisticated agent strategies that currently exist are comprised of a fixed set of modules. Therefore, the studies for proposing the negotiation strategies focusing on the modules are important and influential [14,31].

5 Conclusion & Future Work

This paper describes the ANAC 2017 Repeated Multilateral Negotiation Competition League. Learning from previous negotiations and enhancing the current negotiation strategy based on those experiences are the main challenges introduced in this competition. According to the results of the qualification and final round, it can be observed that employing different negotiation strategies against the same opponents in a repeated negotiation setup may let the agent gain higher utility on average (e.g., PonPoko Agent). Although the participants introduced new ideas regarding the underlying challenges, there is still some room for improving and designing more sophisticated agents. For this reason, organizers have introduced the same challenge in ANAC 2018 to enrich the benchmark in repeated multilateral negotiation research.

There are a number of new challenges planned for future iterations of the competition. For 2019, we aim to introduce the challenge of *negotiating under preference uncertainty* [10,15,27]. Instead of a fully specified utility function, agents in 2019 will be tasked with performing a successful negotiation when only partial preferences is available in the form of ordinal preference comparisons (i.e. "outcome x is preferred over outcome y"). The overarching goal is to design negotiating agents that will eventually be able to represent users in a negotiation, even when limited preference elicitation can be performed (e.g. due to user bother).

Acknowledgments

This work is part of the Veni research programme with project number 639.021.751, which is financed by the Netherlands Organisation for Scientific Research (NWO) and by the ITEA M2MGrids Project, grant number ITEA141011.

References

1. Aydoğan, R., Baarslag, T., Hindriks, K.V., Jonker, C.M., Yolum, P.: Heuristics for using cp-nets in utility-based negotiation without knowing utilities. Knowledge and Information Systems **45**(2), 357–388 (2014)
2. Aydoğan, R., Festen, D., Hindriks, K.V., Jonker, C.M.: Alternating offers protocols for multilateral negotiation. In: Fujita, K., Bai, Q., Ito, T., Zhang, M., Ren, F., Aydoğan, R., Hadfi, R. (eds.) Modern Approaches to Agent-based Complex Automated Negotiation, pp. 153–167. Springer International Publishing (2017)
3. Aydoğan, R., Hindriks, K.V., Jonker, C.M.: Multilateral mediated negotiation protocols with feedback. In: Novel Insights in Agent-based Complex Automated Negotiation, pp. 43–59. Springer (2014)
4. Aydoğan, R., Yolum, P.: The effect of preference representation on learning preferences in negotiation. In: Ito, T., Zhang, M., Robu, V., Fatima, S., Matsuo, T. (eds.) New Trends in Agent-Based Complex Automated Negotiations, Studies in Computational Intelligence, vol. 383, pp. 3–20. Springer Berlin Heidelberg (2012). https://doi.org/10.1007/978-3-642-24696-8_1, `http://dx.doi.org/10.1007/978-3-642-24696-8_1`
5. Baarslag, T.: What to Bid and When to Stop. Dissertation, Delft University of Technology (Sep 2014). https://doi.org/10.4233/uuid:3df6e234-a7c1-4dbe-9eb9-baadabc04bca, `http://dx.doi.org/10.4233/uuid:3df6e234-a7c1-4dbe-9eb9-baadabc04bca`
6. Baarslag, T., Aydoğan, R., Hindriks, K.V., Fuijita, K., Ito, T., Jonker, C.M.: The automated negotiating agents competition, 2010-2015. AI Magazine **36**(4), 115–118 (12/2015 2015), `http://www.aaai.org/ojs/index.php/aimagazine/article/view/2609`
7. Baarslag, T., Dirkzwager, A.S., Hindriks, K.V., Jonker, C.M.: The significance of bidding, accepting and opponent modeling in automated negotiation. In: 21st European Conference on Artificial Intelligence. Frontiers in Artificial Intelligence and Applications, vol. 263, pp. 27–32. IOS Press (2014). https://doi.org/10.3233/978-1-61499-419-0-27, `http://ebooks.iospress.nl/volumearticle/36911`
8. Baarslag, T., Fujita, K., Gerding, E., Hindriks, K., Ito, T., Jennings, N.R., Jonker, C., Kraus, S., Lin, R., Robu, V., Williams, C.: Evaluating practical negotiating agents: Results and analysis of the 2011 international competition. Artificial Intelligence Journal (AIJ) **198**, 73–103 (2013)
9. Baarslag, T., Fujita, K., Gerding, E.H., Hindriks, K., Ito, T., Jennings, N.R., Jonker, C., Kraus, S., Lin, R., Robu, V., Williams, C.R.: Evaluating practical negotiating agents: Results and analysis of the 2011 international competition. Artificial Intelligence **198**(0), 73 – 103 (May 2013). https://doi.org/10.1016/j.artint.2012.09.004, `http://dx.doi.org/10.1016/j.artint.2012.09.004`
10. Baarslag, T., Gerding, E.H.: Optimal incremental preference elicitation during negotiation. In: Proceedings of the Twenty-fourth International Joint Conference on Artificial Intelligence. pp. 3–9. IJCAI'15, AAAI Press (2015), `http://dl.acm.org/citation.cfm?id=2832249.2832250`

11. Baarslag, T., Gerding, E.H., Aydoğan, R., Schraefel, M.: Optimal negotiation decision functions in time-sensitive domains. In: Web Intelligence and Intelligent Agent Technology (WI-IAT), 2015 IEEE/WIC/ACM International Conference on. vol. 2, pp. 190–197. IEEE (2015)
12. Baarslag, T., Hendrikx, M.J., Hindriks, K.V., Jonker, C.M.: Learning about the opponent in automated bilateral negotiation: a comprehensive survey of opponent modeling techniques. Autonomous Agents and Multi-Agent Systems pp. 1–50 (2015). https://doi.org/10.1007/s10458-015-9309-1, `http://dx.doi.org/10.1007/s10458-015-9309-1`
13. Baarslag, T., Hindriks, K., Jonker, C.M., Kraus, S., Lin, R.: The first automated negotiating agents competition (ANAC 2010). In: Ito, T., Zhang, M., Robu, V., Fatima, S., Matsuo, T. (eds.) New Trends in Agent-based Complex Automated Negotiations, Series of Studies in Computational Intelligence. pp. 113–135. Springer-Verlag, Berlin, Heidelberg (2012). https://doi.org/10.1007/978-3-642-24696-8_7, `dx.doi.org/10.1007/978-3-642-24696-8_7`
14. Baarslag, T., Hindriks, K.V.: Accepting optimally in automated negotiation with incomplete information. In: Proceedings of the 12th International conference on Autonomous agents and multi-agent systems (AAMAS 2013). pp. 715–722 (2013)
15. Baarslag, T., Kaisers, M.: The value of information in automated negotiation: A decision model for eliciting user preferences. In: Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems. pp. 391–400. AAMAS '17, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC (2017), `http://dl.acm.org/citation.cfm?id=3091125.3091185`
16. Fujita, K.: Automated strategy adaptation for multi-times bilateral closed negotiations. In: Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014). pp. 1509–1510 (2014)
17. Fujita, K.: Compromising adjustment strategy based on tki conflict mode for multi-times bilateral closed negotiations. Computational Intelligence (2017). https://doi.org/10.1111/coin.12107
18. Fujita, K., Aydoğan, R., Baarslag, T., Ito, T., Jonker, C.: The fifth automated negotiating agents competition (ANAC 2014). In: Fukuta, N., Ito, T., Zhang, M., Fujita, K., Robu, V. (eds.) Recent Advances in Agent-based Complex Automated Negotiation, pp. 211–224. Springer International Publishing, Cham (2016). https://doi.org/10.1007/978-3-319-30307-9_13, `http://dx.doi.org/10.1007/978-3-319-30307-9_13`
19. Gal, K., Ito, T., Jonker, C., Kraus, S., Hindriks, K., Lin, R., Baarslag, T.: The forth international automated negotiating agents competition (ANAC2013) (2013)
20. Güneş, T.D., Arditi, E., Aydoğan, R.: Collective voice of experts in multilateral negotiation. In: In PRIMA 2017: Principles and Practice of Multi-Agent Systems. PRIMA 2017. Lecture Notes in Computer Science, vol 10621. Springer, Cham. pp. 450–458 (2017)
21. Hara, K., Ito, T.: A mediation mechanism for automated negotiating agents whose utility changes over time. In: Twenty-Seventh AAAI Conference on Artificial Intelligence (2013)
22. Jonker, C.M., Aydoğan, R., Baarslag, T., Fujita, K., Ito, T., Hindriks, K.: Automated negotiating agents competition (anac). In: Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI-17). pp. 5070–5072. AAAI Press (2017)
23. Kawaguchi, S., Fujita, K., Ito, T.: Compromising strategy based on estimated maximum utility for automated negotiation agents competition (anac-10). In: 24th International Conference on Industrial Engineering and Other Applications of Applied Intelligent Systems (IEA/AIE d2011). pp. 501–510 (2011)
24. Khosravimehr, Z., Nassiri-Mofakham, F.: Pars agent: Hybrid time-dependent, random and frequency-based bidding and acceptance strategies in multilateral negotiations. In: Fujita, K., Bai, Q., Ito, T., Zhang, M., Ren, F., Aydogan, R., Hadfi, R. (eds.) Modern Approaches

to Agent-based Complex Automated Negotiation, p. 175183. Springer Publishing Company, Incorporated, 1st edn. (2017)
25. Klein, M., Faratin, P., Sayama, H., Bar-Yam, Y.: Negotiating complex contracts. Group Decision and Negotiation **12**(2), 58–73 (2003)
26. Lin, R., Kraus, S., Baarslag, T., Tykhonov, D., Hindriks, K., Jonker, C.M.: Genius: An integrated environment for supporting the design of generic automated negotiators. Computational Intelligence **30**(1), 48–70 (2014)
27. Luo, X., Jennings, N.R., Shadbolt, N.: Acquiring user strategies and preferences for negotiating agents: A default then adjust method. International Journal of Human Computer Studies **64**(4), 304–321 (2006), http://eprints.soton.ac.uk/264474/
28. Marsa-Maestre, I., Klein, M., Jonker, C.M., Aydoğan, R.: From problems to protocols: Towards a negotiation handbook. Decision Support Systems **60**, 39–54 (2014)
29. Sánchez-Anguix, V., Valero, S., Julián, V., Botti, V., García-Fornes, A.: Evolutionary-aided negotiation model for bilateral bargaining in ambient intelligence domains with complex utility functions. Information Sciences **222**, 25–46 (2013)
30. Williams, C.R., Robu, V., Gerding, E.H., Jennings, N.R.: An overview of the results and insights from the third automated negotiating agents competition (ANAC 2012). In: Marsa-Maestre, I., Lopez-Carmona, M.A., Ito, T., Zhang, M., Bai, Q., Fujita, K. (eds.) Novel Insights in Agent-based Complex Automated Negotiation, Studies in Computational Intelligence, vol. 535, pp. 151–162. Springer, Japan (2014). https://doi.org/10.1007/978-4-431-54758-7_9, http://dx.doi.org/10.1007/978-4-431-54758-7_9
31. Zafari, F., Mofakham, F.N.: POPPONENT: highly accurate, individually and socially efficient opponent preference model in bilateral multi issue negotiations. In: Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI 2017). pp. 5100–5104 (2017)