ANAC 2018: Repeated Multilateral Negotiation League

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

[*1] Department of Computer Science, Özyeğin University, Istanbul, Turkey
[*2] Institute of Engineering, Tokyo University of Agriculture and Technology Tokyo, Japan
[*3] Centrum Wiskunde & Informatica (CWI), Amsterdam, The Netherlands
[*4] Interactive Intelligence Group, Delft University of Technology, Delft, The Netherlands
[*5] Nagoya Institute of Technology, Nagoya, Japan

There are a number of research challenges in the field of Automated Negotiation. The Ninth International Automated Negotiating Agent Competition aimed to encourage participants to develop effective negotiating agents, which can negotiate with multiple opponents more than once. This paper discusses essential research challenges for such negotiations as well as presenting the competition set-up and results. Results showed that winner agents mostly adopt hybrid bidding strategies and take their opponents' preferences as well as their strategy into account.

1. Introduction

In multi-agent systems, agents mostly interact and collaborate with each other to achieve their goals; however, their interests and preferences may sometimes conflict. In such situations, agents can resolve their conflict and come up a consensus through negotiation. Therefore, it is important to design and develop efficient negotiation strategies for autonomous agents [Jennings 98, Baarslag 15, Fatima 14]. In order to facilitate the research on negotiation in multi-agent systems and to provide unique benchmarks for evaluation of the developed negotiation strategies, an international competition on automated negotiation namely ANAC [Jonker 17] has been organized for several years. As the main organizers of this competition, we aim to address a variety of research challenges (e.g. uncertainty about opponent, reasoning on complex preferences, negotiating with multiple opponents). In ANAC 2018, we introduced three leagues: Human-agent negotiation [Mell 18], Negotiation Strategies for Diplomacy Game [Jonge 19] and Repeated Multilateral Negotiation. This paper presents the competition setup and evaluation results of the repeated multilateral negotiation as well as pointing out the main challenges.

In multilateral negotiation, there are more than two agents searching for a consensus. Since it involves more conflicts and interactions, it is more complicated than bilateral negotiation [Aydoğan 14, Fujita 12, Fujita 14]. When those agents have a long term relation, they may need to negotiate with each other more than once. In such a situation, it is important for agents to understand their opponents' needs and strategy well and adjust their strategy accordingly so to find a better deal. It is also essential to recall that they are going to negotiate with the same opponents again. If their strategy is based on completely exploiting the other sides, their negotiation may end up with

a failure. In ANAC 2018, we encourage the participants to pursue designing effective negotiation strategies for such repeated multilateral negotiations.

Competition results showed that successful agents mostly employ a hybrid bidding strategy in order to avoid being exploited and consider their opponents' best offer in their previous negotiation. Furthermore, they usually adopt a frequency based opponent modeling. Some agents model the likelihood of an offer to be accepted by other agents by analyzing the bids in their past negotiations while others use opponents' best offers in their bidding strategy. In the rest of this paper, the essential research challenges are discussed, and competition setup and results are explained.

2. ANAC 2018 Competition Challenges

Although agents negotiate with the same opponents several times, they do not know their exact preferences and strategies. The essential research goal is to model their opponents' behaviour or preferences based on their past negotiations and to incorporate those models into their bidding strategy so as to improve their negotiation outcome. However, this is not trivial since their opponents may change their behaviour over time in spite of having the same preferences. Therefore, learnt model may mislead the agents. Furthermore, if an agent tries to exploit its opponents based on what it learnt about them, they may reciprocate in a similar way next time. It would result in decreasing utility for all agents. Moreover, when agents negotiate with the same opponents, they should establish a good relationship while still aiming at maximizing their own utility. There is a trade-off between to what extent act nicely and to what extent consider its own benefit.

3. Competition Setup

In the competition, three agents negotiate on multiple issues to reach a consensus by following the Stacked Alter-

Contact: Reyhan Aydoğan, Özyeğin University, Özyeğin University, Istanbul, Turkey, Reyhan Aydoğan

native Offers Protocol (SAOP) [Aydoğan 17]. According to this protocol, one of the agents starts the negotiation with an offer. Agents can take their action in a turn-taking fashion. When an offer is made by any agent and the next agent in line can take the following actions:

- Make a counter offer (overriding the previous offer)

- Accept the current offer

- Walk away (e.g. ending the negotiation without any agreement)

This process continues in a turn taking fashion until agents reach an agreement or the given deadline, or one of the agents walks away. It is worth noting that an agreement is reached if and only if all agents accept the agreed offer. If negotiation fails, agents receive their reservation utility (i.e., BATNA). The utility of agreement for each agent is calculated with respect to their own preferences. Note that in the competition preferences of each agent are represented by means of additive utility function as shown in Equation 1 where $V_n(v_j)$ denotes agent n's valuation of the value for the issue j in the given bid and $w_{n,j}$ denotes the weights of that issue. In other words, agents sum up their weighted valuation of each issue value to calculate the overall utility. It is worth noting that each agent can only access their own preferences during the negotiation; that is, they do not know each others' truth preferences.

$$u_n(b^t) = \sum_{j \in I} V_n(b_j^t) \cdot w_{n,j} \tag{1}$$

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. In addition, agents are allowed to model their opponents' preferences by examining the bids exchanged during their negotiation and to store the learnt model where they can access in their further negotiations.

Genius 8.0.4 [Lin 14] was used to run negotiation tournaments in the competition. Agents were evaluated in four different negotiation scenarios described in Table 1. Note that all participants submitted a negotiation scenario consisting of three conflicting preference profiles and four of them were chosen for the tournaments based on varying size of their outcome space.

Table 1: Negotiation Scenarios

Name	# of Issue	# of Values	# of Outcomes
Meng wan	6	4,4,4,4,3,4	3072
BetaOne	3	4,4,4	64
IQSon	7	7,6,5,3,4,4,4	40320
Hamada	4	5,5,5,5	625

In order to complete such an extensive set of tournaments within a limited time frame, we used some high-spec computers, made available by Tokyo University of Agriculture

Figure 1: Results of Qualification Round for Pool-A

and Technology, Japan. Specifically, each of these machines contained an *Intel Core i7* CPU, at least 64GB of DDR3 memory, and a hard drive with at least 2TB of capacity.

4. Result of the Competition

We have received 21 submissions from 10 institutions in eight different countries. The performance of the agents were evaluated according to their average *individual utility* and average sum of utilities (*social welfare*). The competition consists of two stages. In the *qualification round*, finalist agents are determined while the winners for each category (*individual* and *social welfare*) are determined in the *final round*. In the following sections, results for each round are presented.

4.1 Qualification Results

Running the whole tournament involving 21 agents in four domains with 5 repetitions were not feasible within the given time. Therefore, three agent pools were generated randomly as follows:

- **Pool-A:** Meng wan, AgentHerb, IQSun2018, PonPokoRampage, FullAgent, Seto, Lancelot

- **Pool-B:** Beta One, Yeela, SMAC_Agent, AgreeableAgent2018, ConDAgent, Shiboy, Libra

- **Pool-C:** AgentNP1, GroupY, ATeamAgent, Sontag, Agent33, Agent_Hama, Exp-Rubick

The top 3 performing agents in each pool proceeded to the final. Therefore, there were nine finalists for each category after the qualification round.

According to the results of the pool-A, the *Mengwan*, *IQSun2018*, and *PonPokoRampage* qualified for the final round in the individual utility category. As far as the average sum of the utilities are concerned, *AgentHerb*, *IQSun2018*, and *FullAgent* are qualified for the final round.

According to the results of the pool-B, the *BetaOne*, *AgreeableAgent2018*, and *Shiboy* qualified for the final round in the individual utility category. As far as the average sum of the utilities are concerned, *Yeela*, *AgreeableAgent2018*, and *ConDAgent* are qualified for the final round.

According to the results of the pool-C, the *AgentNP1*, *GroupY*, and *Sontag* qualified for the final round in the individual utility category. As far as the average sum of the utilities are concerned, *AgentNP1*, *Sontag*, and *Agent33* are qualified for the final round.

Agent Name	Individual Utility	Social Welfare
Beta One	0.482887311	1.404794645
Yeela	0.399430418	1.422654384
SMAC_Agent	0.464635439	1.325276881
AgreeableAgent2018	0.505584897	1.505251211
ConDAgent	0.471002153	1.518896956
Shiboy	0.502585371	1.411124211
Libra	0.445786573	1.227738202

Figure 2: Results of Qualification Round for Pool-B

Agent Name	Individual Utility	Social Welfare
AgentNP1	0.512177707	1.52976242
GroupY	0.485323479	1.33874982
ATeamAgent	0.345297603	0.980979189
Sontag	0.484535782	1.515047114
Agent33	0.428666974	1.419233621
Agent_Hama	0.469592183	1.393009022

Figure 3: Results of Qualification Round for Pool-C

4.2 Final Results: Individual Utility Category

Figure 4 shows the average individual utility gained by each finalist over 2520 negotiations. Recall that each negotiating agent negotiates with all other agents five times for each negotiation scenario in Table 1. *AgreeableAgent2018* gained 0.59 on average and won the competition. *Meng wan* and *Beta One* agents were awarded second and third place respectively. The detailed description of the winner strategies according to the individual utility, are given.

- *AgreeableAgent2018 by Sahar Mirzayi (University of Tehran, Iran):* This agent tries to learn its opponents' preferences during the ongoing negotiations by using a frequency based modeling. It uses a time-based bidding strategy, which takes its opponents' preferences into account. Basically, it generates all candidate bids above estimated target utility and sorts those candidates according to opponent models. It makes the bid selected by using Roulette Wheel Selection. It employs ACNext acceptance strategy. It also accepts it opponents' bid, which is higher than reservation utility if it is almost deadline.
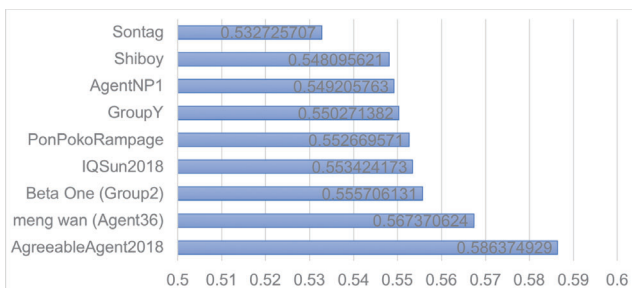
- *Meng wan by Meng Wan, Hui Cui (University*

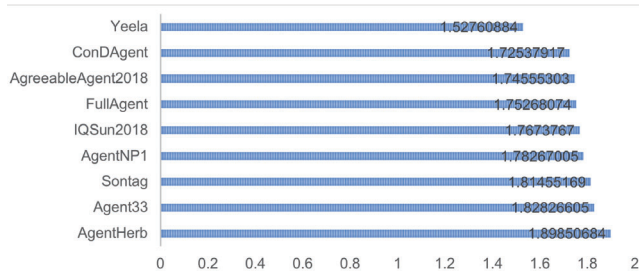Figure 4: Overall Ranking w.r.t. Individual Utility

Figure 5: Overall Ranking wrt. Social Welfare

of Southampton, UK): It employs a hybrid bidding strategy composed of three bidding strategy. The first strategy generates a random bid above an estimated target utility (time-based) while the second strategy chooses a random bid among opponents' best offers so far. The third bidding strategy makes a bid whose utility is higher than target utility, which also maximizes its opponents' utility based on estimated frequency-based opponent models. It employs a time-based acceptance strategy, which is not inclined to accept any offer at the beginning. By the end of negotiation, this agent tends to ask opponents' best offers so far.

- *Beta One by Alper Sekerci, Abdulkadir Nurkalem (Özyeğin University):* This agent has a tendency to be more stubborn against stubborn opponents while more generous against generous opponents. For the opponent modelling, the agent uses statistical analysis to decide whether it should concede or not. The personality of the agent is parameterized. By analyzing the history, the agent is able to tune its parameters (e.g. selfish ratio) to try to perform better in the upcoming negotiations.

4.3 Social Welfare Category

Figure 5 shows the average sum of utilities of the agreements reached by each agent. *AgentHerb* gained 1.89 on average and won the competition where *Agent33* and *Sontag* agents took the second and third place respectively. The description of the winner strategies are given below.

- *AgentHerb by Alon Stern, Amit Moryossef, Yehudit Reyzer, Karin Dahan (Bar Ilan University, Israel):* This agent records the history of bids for each opponent and whether they were accepted or rejected by the opponents. Accordingly, Agent Herb uses a logistic regression model to predict the likelihood of acceptance of a bid by the opponents. It makes its bids based on its chances of acceptance by all the opponents while taking its own utility into account. The agent employs ACNext acceptance strategy with a discount factor.

- *Agent33 by Liu Shan (Nagoya Institute of Technology, Japan):* The bidding Strategy aims to propose bids around Nash bargaining solution (NBS). Agent33 uses a novel heuristic method in order to find

Table 2: Detailed Results of Winners Tournaments

	Agreeable	MengWan	BetaOne	Social Welfare	Dist. to Pareto	Dist. to Nash
Individual Category	0.769856823	**0.805217872**	0.787584175	2.362659083	0.019035083	0.223667667
	Agent Herb	**Agent33**	**Sontag**	**Social Welfare**	**Dist. to Pareto**	**Dist. to Nash**
Social Category	0.634797437	0.688591936	**0.899989331**	2.223378333	0.027012083	0.32230525

the promising bids around the NBS. The proposed heuristic method aims to construct a list of the opponents prior issues, which is incrementally updated throughout the negotiation process by calculating the standard deviation of each issue value's frequency and the standard deviation of each issue's value.

- **Sontag by Ryohei Kawata (Tokyo University of Agriculture and Technology, Japan):** Sontag makes concessions in early to increase social welfare. Sontag does not model opponents and no learning from negotiation history. In other words, it expects that the opponents will make a concession finally. The proposal and acceptance of the bids are based on the following equation: $f(t) = t/2.5 - \log(t/2 + 0.1)$.

4.4 Further Analysis of ANAC2018 Winners

After competition, we ran another tournament among the winners per each category. The settings of tournaments are same as ANAC2018; however, the agents in the tournament consists of only the winners of each category (Individual Utility Category: AgreeableAgent2018, MengWan, BetaOne; Social Welfare Category: AgentHerb, Agent33, Sontag). According to the results listed in Table 2, Meng-Wan received the highest average utility although this agent took the second place in the competition. Similarly, when only winners of the social welfare category negotiate with each other, Songtag outperforms other winners. The result should not surprise us because negotiation outcome highly depends on whom we are negotiating with. Similarly, the winners in individual category received higher average social welfare than the winners of social category. In addition, it can be seen that winners made agreements that are closer to the Pareto Frontier, received higher social welfare.

5. Conclusion

This paper discusses main research challenges in repeated multilateral negotiations as well as explaining the ANAC 2018 competition setup and results briefly. As a future work, it would be interesting to analyze the performance of agents elaborately in a more extensive setup. Furthermore, we are planning to run additional tournaments including ANAC 2017 winner agents and compare their performances.

References

[Aydoğan 17] Aydoğan, R., Festen, D., Hindriks, K. V., and Jonker, C. M.: Alternating offers protocols for multilateral negotiation, in Fujita, K., Bai, Q., Ito, T., Zhang, M., Ren, F., Aydoğan, R., and Hadfi, R. eds., *Modern Approaches to Agent-based Complex Automated Negotiation*, pp. 153–167, Springer Publishing (2017)

[Aydoğan 14] Aydoğan, R., Hindriks, K. V., and Jonker, C. M.: Multilateral mediated negotiation protocols with feedback, in *Novel Insights in Agent-based Complex Automated Negotiation*, pp. 43–59, Springer (2014)

[Baarslag 15] Baarslag, T., Gerding, E. H., Aydoğan, R., and Schraefel, M.: Optimal negotiation decision functions in time-sensitive domains, in *IEEE International Conference on Web Intelligence and Intelligent Agent Technology*, Vol. 2, pp. 190–197IEEE (2015)

[Fatima 14] Fatima, S., Kraus, S., and Wooldridge, M.: *Principles of Automated Negotiation*, Cambridge University Press, New York, NY, USA (2014)

[Fujita 12] Fujita, K., Ito, T., and Klein, M.: A Secure and Fair Protocol that Addresses Weaknesses of the Nash Bargaining Solution in Nonlinear Negotiation, *Group Decision and Negotiation*, Vol. 21, No. 1, pp. 29–47 (2012)

[Fujita 14] Fujita, K., Ito, T., and Klein, M.: Efficient issue-grouping approach for multiple interdependent issues negotiation between exaggerator agents, *Decision Support Systems*, Vol. 60, pp. 10–17 (2014)

[Jennings 98] Jennings, N. R., Faratin, P., Lomuscio, A., Parsons, S., Sierra, C., and Wooldridge, M.: Automated Negotiation : Prospects , Methods and Challenges (1998)

[Jonge 19] Jonge, de D., Baarslag, T., Aydoğan, R., Jonker, C. M., Fujita, K., and Ito, T.: The Challenge of Negotiation in the Game of Diplomacy, in *Proceedings of the 6th Conference on Agreement Technologies* (2019)

[Jonker 17] Jonker, C. M., Aydoğan, R., Baarslag, T., Fujita, K., Ito, T., and Hindriks, K. V.: Automated Negotiating Agents Competition (ANAC)., in *AAAI*, pp. 5070–5072 (2017)

[Lin 14] Lin, R., Kraus, S., Baarslag, T., Tykhonov, D., Hindriks, K., and Jonker, C. M.: Genius: An integrated environment for supporting the design of generic automated negotiators, *Computational Intelligence*, Vol. 30, No. 1, pp. 48–70 (2014)

[Mell 18] Mell, J., Gratch, J., Baarslag, T., Aydoğan, R., and Jonker, C. M.: Results of the First Annual Human-Agent League of the Automated Negotiating Agents Competition, in *Proceedings of the 18th International Conference on Intelligent Virtual Agents*, IVA '18, pp. 23–28, New York, NY, USA (2018), ACM