

Collective Voice of Experts in Multilateral Negotiation

Taha D. Güneş^{1,2}, Emir Arditi², and Reyhan Aydoğan^{2,3,4}

¹ Agents, Interaction and Complexity Group, University of Southampton, UK,
ORCID: 0000-0003-2774-7027 t.d.gunes@soton.ac.uk

² Department of Computer Science, Özyeğin University, Istanbul, Turkey
emir.arditi@ozu.edu.tr and reyhan.aydogan@ozyegin.edu.tr

³ Interactive Intelligence Group, Delft University of Technology, The Netherlands

⁴ Frontier Research Institute for Information Science, Nagoya Institute of Technology

Abstract. Inspired from the ideas such as “*algorithm portfolio*”, “*mixture of experts*”, and “genetic algorithm”, this paper presents two novel negotiation strategies, which combine multiple negotiation experts to decide what to bid and what to accept during the negotiation. In the first approach namely *incremental portfolio*, a bid is constructed by asking each negotiation agent’s opinion in the portfolio and picking one of the suggestions stochastically considering the expertise levels of the agents. In the second approach namely *crossover strategy*, each expert agent makes a bid suggestion and a majority voting is used on each issue value to decide the bid content. The proposed approaches have been evaluated empirically and our experimental results showed that the crossover strategy outperformed the top five finalists of the ANAC 2016 Negotiation Competition in terms of the obtained average individual utility.

Keywords: agreement technologies, automated negotiation, multilateral negotiation, negotiation competition, multi-agent systems

Automated negotiation is a fundamental solution approach in multi-agent systems where there is a conflict of interest among parties [12]. It takes place to resolve the underlying conflicts for coming up with a joint agreement. Although there have been a number of negotiating agents designed in this area with varying strategies, heuristics and assumptions [7, 21, 8, 5, 4, 20], the empirical evaluations in the negotiation literature have shown that each strategy performs significantly different in varying scenarios [6, 10, 22, 17], due to their opponent’s strategies and diverse characteristics of negotiation scenarios (e.g. size of outcome space, degree of conflicts). Some negotiation strategies may outperform other strategies in some negotiation scenarios while they may not perform well in other scenarios. Based on those observations, we investigate the question of how to devise an approach that combines the strengths of expert agents’ strategies in order to achieve higher performance overall in a variety of negotiation scenarios.

The question of which strategy our agent should employ in the underlying negotiation is intriguing. Picking the best-performing strategy for a specific scenario resembles to the widely known “algorithm selection problem” [19], where

an algorithm amongst a set of algorithms for a particular problem, which is expected to perform better than others, is selected. One common approach to this problem known as “winner-take-all” [14], is to select the algorithm whose overall performance is the best on a given problem distribution. Following this approach in the context of negotiation might not be appropriate since the performance of the negotiation strategies may significantly vary regarding to their opponents as well as the negotiation scenarios.

Inspired from “*algorithm portfolio*” approach [15], Ilany and Gal [11] develop a meta-agent, which aims to guess the performance of a set of bilateral negotiation strategies based on distinctive features of domains such as the size of the outcome space, the competitiveness of the given scenario, and degree of conflicts among preferences and so on, and accordingly adopt the strategy expected to perform best for the given scenario. However, more complex automated negotiations such as multilateral negotiations are not covered by this work.

Instead of following a one-shot strategy determination, an agent may combine multiple strategies of expert agents on-the-fly to achieve higher performance similarly to “*mixture of experts*”, combining multiple learners approach in machine learning [1]. Accordingly, this paper studies how to combine multiple experts in multilateral negotiation settings in which more than two agents aim to reach a joint consensus by following Stacked Alternating Offer Protocol (SAOP) [2]. In the light of aforementioned motivations, this paper introduces two techniques to merge the strength of the experts: *incremental portfolio strategy* and *crossover voting strategy*. In the former approach (incremental portfolio), our strategy assigns a weight to each expert denoting how reliable they are, and choose one of the experts to ask what action to take at each negotiation round by taking their associated weights into account. Inspired from genetic algorithm and mixture of expert, the latter strategy, namely crossover voting strategy, asks each expert’s opinion on what to bid and merges the given bids by voting the values per each issue. That is, a bid is constructed from other bids like generating a genome by using genetic algorithm except there is no mutation operation in our case and crossing over is done by means of voting. It is worth noting that our agent employing incremental portfolio strategy, *Caduceus* participated in ANAC 2016 competition⁵, and got the first place amongst ten finalist agents.

We evaluate the performance of our strategies against a variety of negotiation scenarios (i.e., different domains, different opponents) in **General Environment for Negotiation with Intelligent multi-purpose Usage**, [16] **GENIUS**⁶, a testbed also used in International Competition Automated Negotiating Agents Competition (ANAC) [13, 9]. The benchmark provided by the competition has been used in our tests. The conducted results show that our agent employing *crossover voting strategy* yields higher performance overall in a variety of different negotiation settings regarding the gained average individual utility.

The rest of the paper is organized as: First, we elaborate on related work in Section 1 and our problem setting in Section 2. Our negotiation strategies for

⁵ <http://web.tuat.ac.jp/%7Ekatfujii/ANAC2016/>

⁶ <http://ii.tudelft.nl/genius/>

multilateral negotiation setting are explained in Section 3. Section 4 provides empirical evaluation of our findings. Lastly, we conclude our work with future work directions in Section 5.

1 Related Work

Alpaydin proposes to combine multiple machine learning algorithms in order to increase the performance of the overall learning process [1]. One of the techniques mentioned in this paper is voting where each learner has a weight and their votes have an impact on the overall decision accordingly. In our work, we adopt a similar approach. In our incremental portfolio approach, each negotiation strategy has a weight denoting their expertise level and the bid is made stochastically by asking each strategy in our portfolio by taking their weights into account. In crossover approach, we again use voting but slightly different way. Each agent suggests a bid and we use the majority voting on each issue to decide their values in the final bid to be made.

Another study [11] also discusses the positive effects of using different algorithms in different domains in order to achieve better outcomes in the context of automated negotiation. The study mainly focuses on selecting a negotiation strategy based on the characteristics of the negotiation strategy. They apply machine learning algorithms to predict which strategy would work better in the given negotiation scenario. The main difference between that study and our work is that they pick the best negotiation strategy guessed by a machine learning algorithm, and employ this study in the entire negotiation whereas our approach is using multiple strategies and these strategies collectively determine what to bid and what to accept during the negotiation.

Furthermore, Aydoğan *et al.* pursue the problem of selecting the most effective negotiation mechanism for a given negotiation scenario [3]. While Ilany and Gal address the problem of selecting the best negotiation strategy, Aydoğan *et al.* studies which negotiation mechanism (i.e., negotiation protocol and compatible strategies) should be adopted under the given negotiation problem. They define a set of scenario metrics to capture the characteristics of the negotiation problems and apply machine learning techniques to guess which mechanisms work better with which scenarios. That work also focuses on bilateral negotiations while our approach is for multilateral negotiations.

Faratin *et al.* propose a number of negotiation tactics such as time-based concession tactics and behavior-based tactics, which mimics opponent's behavior [21]. They also provide a meta-strategy, which makes a counter proposal by linear combination of tactics. This approach is similar to our approach since both approaches take more than one tactics/strategies into account to make a counter offer during the negotiation. However, our approach also asks each expert's opinion in the decision of acceptance of a given bid. Another difference is that they consider only numeric issues in a specific range [min, max] so that they can construct a counter offer by taking the linear combination of the bid values

suggested by each tactic. On the other hand, our approach works well with the discrete issues as well as the integer issues.

Matsune and Fujita combines multiple existing opponent modeling algorithms by using boosting based on the least square method and nonlinear programming in multilateral negotiation settings [9]. In that sense, our work is complementary. While we are combining bidding and acceptance strategies of multiple negotiation strategies, they are combining existing opponent models.

2 Problem Setting

In this work, we focus on non-mediated multilateral negotiations in which more than two negotiating agents, $\text{Agents}=\{1, 2, \dots, n\}$, aim to reach a joint agreement on a set of issues denoted by L . Each issue $j \in L$ can take a value v_j from a predefined set of valid values for that issue denoted by D_j ; that is $v_j \in D_j$. This domain information is shared between each negotiating party. However, preferences of a negotiating party is private; that is, none of the agents does not know their opponents' preferences. Particularly, the utility function of an opponent is unknown to the negotiating agent. A bid $b = (b_1, \dots, b_{|L|})$ is an assignment of values to all issues where $b_1 \in D_1$.

In our setting, the utility of a given bid is modelled in terms of additive utility function as shown in Equation 1 where $o_i(v_j)$ denotes agent i 's valuation of the value of issue j in the given bid and the weights of that issue is represented by $w_{i,j}$. In other words, agents sum up their weighted valuation of each issue value to calculate the overall utility.

$$u_i(b^t) = \sum_{j \in L} o_i(b_j^t) \cdot w_{i,j} \quad (1)$$

Each party complies to *Stacked Alternating Offers Protocol* (SAOP) [2], which is a turn-taking fashion protocol where each agent i is expected to provide an action a_i in its turn during a negotiation session. Type of an action is denoted as $\text{type}(a)$ which can take $\text{type}(a) = \{I^{bid}, I^{accept}, I^{end}\}$; that is, the agent can make a counter offer, accept the offer in the negotiation table or walk away from the negotiation respectively. The session ends whenever the offer on the table is accepted by all parties, the deadline is reached or a party walks away. If an agreement on the bid b is reached, the utility of the agent i at the time t is estimated as formulated in Equation 2. As seen obviously, the utility of a given bid goes down over time.

$$u_i(b^t) = [\sum_{j \in L} o_i(b_j^t) \cdot w_{i,j}] \cdot \delta_i^t \quad (2)$$

In the case of not reaching a consensus, the utility that every agent is simply equal to $r_i \cdot \delta_i^t$ where r_i is the reservation value and δ_i^t is the discount factor of at time t . The reservation value and the discount factor information is provided within the preference profile. Each agent is self-interested: the agent wants to maximize its individual utility $u_i(b^t)$ at the end of a negotiation session.

We follow the negotiation settings of ANAC 2016 competition for our empirical evaluations. The winner agent is picked after a *tournament* where the participants enter a series of negotiation sessions with all possible settings (e.g. a variety of negotiation domains with varying preference profiles, different combination of agents, etc.). It is important to note that the agents are not allowed to save any domain information throughout the tournament and identities of the participants are not shared. According to the competition, the best agent is the one that achieves the highest total utility after all negotiation rounds.

3 Proposed Negotiation Strategies

Here, we describe our meta strategies based on the opinions of multiple negotiation experts. *Incremental portfolio* strategy in which actions of experts agents are gathered and only a single action from all collected actions is picked stochastically by taking the expertise level of the experts into account while *crossover* strategy combines the suggested actions by using majority voting.

3.1 Incremental Portfolio Strategy

As shown in Algorithm 1.1, *Incremental Portfolio* strategy separates negotiation process into two major phases. In the first phase, our agent chooses the best possible offer for itself (i.e. the offer with the highest utility) and sends this bid to its opponents up to $T \cdot \beta$, where T is the maximum negotiation duration (deadline) and β is the parameter controlling how long the first phase will endure (Lines 6–9). To decide the best value for β , we did some experiments and chose the value giving the highest utility. Note that we observed that involving such an eager phase enables our agent to yield higher utilities in practice.

In the second phase, each expert agent is informed by the received action a , which is one of our opponents' most recent action, and asked to make a decision on which action to be taken: *accept* or *make a counter bid* (Line 12). The response of each agent a'_i can be either a counter-offer I^{bid} or an accept action I^{accept} . In the case of I^{bid} , action is added into candidate bid set, B . Votes of expert agents for actions ℓ_{bid} and ℓ_{accept} are counted (Lines 13–15). Finally, the output action is decided based on those votes. If majority of the votes are in favor of accept action, the agent accepts the opponent's offer a (Line 16); otherwise, a bid is chosen from the set B by Monte Carlo sampling where the weights of each expert agent are taken in account (Line 18).

3.2 Crossover Strategy

Similar to *Incremental Portfolio* Strategy, *Crossover* strategy also consists of two phases. Since the first phase is identical in both strategies, we will focus on how agent will act in the second phase. The valuable information during the generation of a counter-offer action is not preserved and utilised well enough in our *incremental portfolio* strategy since only a single bid is chosen from a set

Algorithm 1.1 *Incremental Portfolio Strategy*

```
1: Input: Action  $a$     {one of an opponent action}
2: Output: Action  $b'$    {for the negotiation table}
3: Initialization:  $A \leftarrow \{\text{set of expert agents}\};$ 
4:  $\Omega \leftarrow (\omega_i, \dots, \omega_{|A|});$  {impact weights of each expert agent}
5:  $B \leftarrow \{\};$  {set of bids suggested by the expert agents}
6: Waiting phase:
7: if  $t < T \cdot \beta$  then
8:   return  $B^{max};$  {return best bid for agent itself}
9: Voting phase:
10: for each  $i \in A$  do
11:    $a'_i \leftarrow$  query agent  $i$  with action  $a$  in time  $t$ ;
12:   if  $\text{type}(a'_i) = I^{bid}$  then
13:      $B \leftarrow B \cup \{a'_i\}$     $\ell_{bid} \leftarrow \ell_{bid} + 1;$ 
14:   else
15:      $\ell_{accept} \leftarrow \ell_{accept} + 1;$ 
16: if  $\ell_{accept} > \ell_{bid}$  then return  $I^{accept};$ 
17: else
18:   return a bid  $b'$  from weighted random sample set  $B$  by  $\Omega;$ 
```

of expert agents' bids. To address this, we extended our strategy by combining offers from the expert agents. We call this new approach, *crossover* strategy by being inspired by “*genetic algorithms*” [18]. The *crossover* strategy, as shown in Algorithm 1.2, combines bids from expert agents by picking the most favourable values of an issue by applying majority voting on each issue value.

As previously described, all expert agents have a weight value to impact the final bid action. The value of an each issue is voted (Lines 19–22) and the value that gets the majority of the votes is then picked in this strategy (Lines 23–24). Therefore not only which action to take is decided collectively by expert agents, but also what to offer is collectively made with this strategy.

4 Performance Evaluation

To evaluate our strategies empirically, we conducted a number of experiments in GENIUS negotiation testbed environment. We compare the performance of the proposed negotiation strategies against the best five agents from ANAC 2016 Negotiation Competition⁷ namely: ParsCat, YXAgent, Farma, MyAgent and Atlas3. We ran 14400 different negotiation sessions (6 agents, 3 participants per sessions, 5 repetitions per session, 6 domains, 3 preference profiles per domain and four differently configured *Caduceus*.) in GENIUS platform with version 7.1.2. Each run per domain took around 24 hours, and conducted within similar machine configurations.

⁷ <http://web.tuat.ac.jp/~katfujii/ANAC2016/>

Algorithm 1.2 *Crossover Strategy*

```
1: Input: Action  $a$  {one of an opponent action}
2: Output: Action  $b'$  {for the negotiation table}
3: Initialization:  $A \leftarrow \{\text{set of expert agents}\};$ 
4:  $\Omega \leftarrow (\omega_i, \dots, \omega_{|A|});$  {impact weights of each expert agent}
5:  $B \leftarrow \{\};$  {set of bids suggested by the expert agents}
6: Same as Algorithm 1.1;
7: if  $\ell_{\text{accept}} > \ell_{\text{bid}}$  then return  $I^{\text{accept}};$ 
8: else  $b' \leftarrow (\emptyset, \dots, \emptyset)$  {create an empty bid}.
9:   for each  $j \in L$  do {each issue in negotiation scenario}
10:    for each  $v_j \in D_j$  do {each valid value for issue j}
11:      for each  $i \in A$  do {each expert agent}
12:         $\ell_{v_j} = \ell_{v_j} + \mathbb{1}_{b_i}(v_j) \cdot \omega_i;$ 
13:   for each  $j \in L$  do {each issue in the domain}
14:      $b'_j = \arg \max_{v_j} \ell_{v_j};$  {pick the value that got the highest vote}
15:   return  $b';$ 
```

Figure 1 shows the average individual utilities gained by each agent over 14400 different negotiations. We created two versions of each negotiation strategy. First, we use equal weights for each agent in our portfolio while in the second version we set different weights for each agent regarding to their past performance. It is obviously seen that Caduceus agent with *crossover* strategy with equal weights (*CE*) and with ordered weights (*CO*) outperforms all of the agents in overall. The performance of our *incremental portfolio* strategy suffered slightly with the given equal weights, which were giving more chance to all consisting expert agents (0.727 versus 0.721).




Fig. 1: The performance of the agents

It is worth noting that the ranking order of the other top five agents from the competition in our experiments changed, but the utility range was similar

and stayed above 0.7 as in competition. This is because the performance of the negotiating agents highly depend on their opponent sets. In the competition, there were 10 finalist agents. However, we have 5 top agents and the variants of *Caduceus*. Since this competition is harder than ANAC 2016, some of the agents' rank may be different from their rank in the competition.

5 Conclusion and Discussion

In this paper, we have introduced two novel negotiation strategies, which ask negotiation expert's opinion to collectively decide which action to be taken. To our knowledge, this is the first work that uses existing expert agents progressively in multilateral negotiations. Our results have shown that the *crossover* strategy that employs previous years top five agents from ANAC 2015 with varying impact weights outperformed top five agents (excluding *Caduceus* from ANAC 2016). Thus, we have shown the feasibility of utilising pre-existing expert agents to devise better negotiation behaviors to achieve higher utilities, even by using dated expert agents. When calculating the votes for the acceptance, the weights were not taken into account. We can try to consider weights in this process too. Furthermore, the same weights of experts are used for the entire negotiation. As a future work, the agent may aim to observe its opponent's behavior and adapt these weights accordingly during the negotiation.

Acknowledgments

We thank Burak Atalay and Bahadır Kırdan for their help in implementation of the initial agent. This work was supported by the ITEA M2MGrids Project, grant number ITEA141011.

References

1. Alpaydin, E.: Techniques for combining multiple learners. In: Proceedings of Engineering of Intelligent Systems. pp. 6–12 (1998)
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, Springer International Publishing. publisher (2017)
3. Aydoğan, R., Marsa-Maestre, I., Klein, M., Jonker, C.M.: A machine learning approach for mechanism selection in complex negotiations. In: Proceedings of The 8th International Workshop on Agent-based Complex Automated Negotiations (2015)
4. Aydoğan, R., Yolum, P.: Effective negotiation with partial preference information. In: Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2010), Toronto, Canada, May 10-14, 2010, Volume 1-3. pp. 1605–1606 (2010)
5. 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)

6. Baarslag, T., Hindriks, K., Jonker, C., Kraus, S., Lin, R.: The first automated negotiating agents competition. In: *New Trends in agent-based complex automated negotiations*, Springer. publisher (2012)
7. Beam, C., Segev, A.: Automated negotiations: A survey of the state of the art. *Wirtschaftsinformatik* 39(3), 263–268 (1997)
8. Fatima, S., Kraus, S., Wooldridge, M.: Cambridge University Press. publisher (2014)
9. Fujita, K., Aydoğan, R., Baarslag, T., Hindriks, K., Ito, T., Jonker, C.: The sixth automated negotiating agents competition. In: Fujita, K., Bai, Q., Ito, T., Zhang, M., Ren, F., Aydoğan, R., Hadfi, R. (eds.) *Modern Approaches to Agent-based Complex Automated Negotiation*, Springer International Publishing. publisher, Cham (2017)
10. Fujita, K., Ito, T., Baarslag, T., Hindriks, K., Jonker, C., Kraus, S., Lin, R.: The second automated negotiating agents competition. In: *Complex Automated Negotiations: Theories, Models, and Software Competitions*, Springer. publisher (2013)
11. Ilany, L., Gal, Y.: Algorithm selection in bilateral negotiation. *Autonomous Agents and Multi-Agent Systems* 30(4), 697–723 (Jul 2016)
12. Jennings, N., Faratin, P., Lomuscio, A., Parsons, S., Wooldridge, M., Sierra, C.: Automated negotiation: Prospects, methods and challenges. *Group Decision and Negotiation* 10(2), 199–215 (2001)
13. Jonker, C.M., Aydoğan, R., Baarslag, T., Fujita, K., Ito, T., Hindriks, K.: Automated negotiating agents competition. In: *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI-17)*. AAAI Press. publisher (2017)
14. Lazzaro, J., Ryckebusch, S., Mahowald, M.A., Mead, C.A.: Winner-take-all networks of $O(n)$ complexity. In: *Advances in neural information processing systems*. pp. 703–711 (1989)
15. Leyton-Brown, K., Nudelman, E., Andrew, G., McFadden, J., Shoham, Y.: A portfolio approach to algorithm selection. In: *IJCAI*. vol. 1543, p. 2003 (2003)
16. 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)
17. 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)
18. Melanie, M.: An introduction to genetic algorithms. Cambridge, Massachusetts London, England, Fifth printing 3, 62–75 (1999)
19. Rice, J.R.: The Algorithm Selection Problem. In: *Advances in Computers*, vol. 15, Elsevier. publisher (1976)
20. Sanchez-Anguix, V., Aydogan, R., Julian, V., Jonker, C.: Unanimously acceptable agreements for negotiation teams in unpredictable domains. *Electronic Commerce Research and Applications* 13(4), 243–265 (2014)
21. Sierra, C., Faratin, P., Jennings, N.R.: A service-oriented negotiation model between autonomous agents. In: *Collaboration between human and artificial societies*, Springer. publisher (1999)
22. 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. In: *Novel Insights in Agent-based Complex Automated Negotiation*, Springer. publisher (2014)