Rethinking Frequency Opponent Modeling in Automated Negotiation

Tunali, Okan; Aydogan, Reyhan; Sanchez-Anguix, Victor

DOI
10.1007/978-3-319-69131-2_16

Publication date
2017

Document Version
Final published version

Published in
PRIMA 2017

Citation (APA)
Tunali, O., Aydogan, R., & Sanchez-Anguix, V. (2017). Rethinking Frequency Opponent Modeling in Automated Negotiation. In B. An, A. Bazzan, J. Leite, S. Villata, & L. van der Torre (Eds.), *PRIMA 2017: Principles and Practice of Multi-Agent Systems - 20th International Conference - Proceedings* (pp. 263-279). (Lecture Notes in Computer Science; Vol. 10621). Springer. https://doi.org/10.1007/978-3-319-69131-2_16

Rethinking Frequency Opponent Modeling in Automated Negotiation

Okan Tunalı[1(✉)], Reyhan Aydoğan[1,2], and Victor Sanchez-Anguix[3]

[1] Department of Computer Science, Özyeğin University, Istanbul, Turkey
okan.tunali@ozu.edu.tr
[2] Interactive Intelligence Group, Delft University of Technology, Delft, The Netherlands
R.Aydogan@tudelft.nl
[3] Coventry University, Coventry, UK
ac0872@coventry.ac.uk

Abstract. Frequency opponent modeling is one of the most widely used opponent modeling techniques in automated negotiation, due to its simplicity and its good performance. In fact, it outperforms even more complex mechanisms like Bayesian models. Nevertheless, the classical frequency model does not come without its own assumptions, some of which may not always hold in many realistic settings. This paper advances the state of the art in opponent modeling in automated negotiation by introducing a novel frequency opponent modeling mechanism, which soothes some of the assumptions introduced by classical frequency approaches. The experiments show that our proposed approach outperforms the classic frequency model in terms of evaluation of the outcome space, estimation of the Pareto frontier, and accuracy of both issue value evaluation estimation and issue weight estimation.

Keywords: Agreement technologies · Automated negotiation · Opponent modeling · Multi-agent systems

1 Introduction

In the last few years, we have seen an increasing interest on the study of agreement technologies [27]. This increasing interest goes hand in hand with an incipient acceptance of autonomy and delegation in technology, with some technologies such as self-driven cars [19] being the prime example of this trend. As delegation and autonomous systems become the norm, so will agreement technologies. The reason is simple: autonomous agents are driven by real users' preferences, and, as we all know, conflict is inherent in our world. As a consequence, we need technologies that allow autonomous agents to solve preferential conflicts and, hence, make delegation and autonomy as transparent for the user as possible. Automated negotiation [16,23,29] is considered as one of the core technologies in agreement technologies, as it provides autonomous entities with protocols and algorithms to reach agreements in a distributed way.

B. An et al. (Eds.): PRIMA 2017, LNAI 10621, pp. 263–279, 2017.
https://doi.org/10.1007/978-3-319-69131-2_16

Despite this recent and increasing interest in automated negotiation, research has been carried out for decades. Researchers have proposed a number of negotiation protocols [3,5] and negotiation strategies that guide autonomous agents on how to act in a distributed negotiation process [4,17,26,28]. There are two main families of strategies in automated negotiation process: game theoretic and heuristic approaches. On the one hand, the former focuses on achieving optimal negotiation results under the assumption of full rationality, unbounded computational resources, and, often, full disclosure of preferences. On the other hand, heuristic approaches assume that agents' resources are limited and partial or nil knowledge about the others' preferences, precluding agents from guaranteeing optimal results. This present work is categorized as a heuristic approach.

While optimal negotiation outcomes cannot be guaranteed, it is still crucial for agents to reach outcomes that are as close as possible to the optimal outcomes. There are several ways that agents can resort to optimizing the resulting negotiation outcomes, but perhaps opponent modeling is one of the most important mechanisms. Opponent modeling [9] allows us to build an approximate model of the opponents' preferences, which can be used to propose outcomes that result in win-win situations for involved parties. Hence, making outcomes more appealing and maximizing the odds of reaching an agreement.

One of the most popular opponent modeling mechanisms in automated negotiation is the frequency model [25]. The frequency model aims to build a model of the opponents' preferences assuming linear additive utility functions and steady concession towards lower utilities. For that, the frequency model uses the frequency of negotiation issue values as an indicator of both negotiation issue and value importance. Due to its simplicity and wide acceptance, the frequency model has been used in a myriad of scenarios [1,2,7,18,25]

Despite its popularity, there is no informed research on the robustness of the frequency model in a wide variety of scenarios. Most of the work in this field focuses only on the quality of the agreements and/or percentage of successful negotiations when the given negotiation strategy uses this model as opponent modeling. However, there are a number of factors having a significant impact on the negotiation outcome such as bidding strategy, acceptance strategy, how the estimated opponent's preference model is used in the underlying negotiation strategy, and so on. Therefore, gaining high utility agreements does not indicate by itself how good the opponent model is. Accordingly, this study analyzes how well the frequency model predicts opponent's preferences elaborately by comparing the estimated opponent model with the real preferences. The contributions of this paper are twofold. Firstly, we pose the problems faced by the frequency model in realistic scenarios. Secondly, we propose a new opponent modeling mechanism that deals with some of these problems and outperforms the classic frequency model mechanism.

The rest of this paper is structured as follows: Sect. 2 provides an overview of related work and Sect. 3 addresses the potential problems with the frequency model. The proposed opponent model is explained in Sect. 4. Section 5 provides a detailed analysis of the frequency model as well as the proposed opponent model empirically. Finally, Sect. 6 concludes the paper.

2 Related Work

A variety of opponent modeling mechanisms have been proposed in the automated negotiation literature. Some opponent modeling mechanisms aim to provide an educated guess over the opponents' reservation value or the opponent's concession strategy, while other opponent modeling approaches take an educated guess on the opponents' preferences with respect to outcomes. This paper is enclosed in the latter family. As far as learning techniques for opponent's preferences are concerned, two main approaches namely, probabilistic models (e.g. Bayesian) and frequency approaches, come to the forefront.

Bayesian approaches [10,11,21,28] usually employ Bayes' update rule and a set of hypotheses to model the opponents' preferences. For instance, Bui *et al.* [11] propose a multi-party cooperative negotiation mechanism for the distributed meeting scheduling domain. Agents follow an iterative process that gradually partitions the negotiation space into acceptable areas by expressing their preferences on suggested partitions. In order to speed up the negotiation process, the agents employ Bayesian classifiers to learn other agents' preferences according to the information gathered from the current and past negotiations. Another example of the use of Bayesian learning in negotiation is presented by Buffett *et al.* [10]. In the proposed model, agents negotiate over a set of limited objects that can either be included or excluded from the final deal. Bayesian classifiers are employed to classify opponent's preferences into classes of preference relations over the objects in the negotiation domain. Bayesian learning was also used by Hindriks and Tykhonov [21] in order to predict the shape of the opponent's utility function (i.e., downhill, uphill and triangular), as well as the corresponding rank of issue values and issue weights. Sanchez-Anguix *et al.* [28] used Bayesian classifiers to learn the acceptability of partial offers for each team member in a negotiation team, and their opponent.

On the other hand, frequency approaches [1,2,7,18,25] usually model opponents' preferences by counting the frequency of issue values and the frequency of changes in negotiation issues of the given bids, without considering an explicit set of initial hypotheses in mind. The most popular frequency model was introduced by the HardHeaded agent [25], whereby issue weights are updated when issue values do not change in consecutive pairs of opponent offers, and issue value weights are estimated by counting the occurrences of values in opponent's offers. A more detailed description of this model can be found in Sect. 3. In [18], the authors propose a frequency model similar to HardHeaded's frequency model. The main difference between those approaches is how they estimate the issue weights. The approach in [18] estimates the issue weights based on the relative frequency of the most offered values. Afiouni [1] adapted the classic frequency model to a real-time strategy for a video game (i.e., Civilization IV) where bilateral negotiation is used to exchange resources between parties. The frequency model showed to be applicable in real time, while also shortening the negotiation time/interactions between parties. HardHeaded agent, and thus the frequency opponent modeling, was also employed in [2] to study the efficiency of different agents in cloud computing negotiations. Furthermore, Ikarashi and Fujita proposed a weighted counting method, which aims to learn opponent's preferences by taking what time the bids made by the opponent into account [22]. That study focuses on learning from past negotiations so that the proposed approach was applied on the history of their opponent's bids in the previous negotiation.

There are also other remarkable approaches that are not classified under those two families. Aydoğan *et al.* proposed a concept based learning algorithm to figure out what offers are more likely acceptable for the opponent during the negotiation [6]. Kernel density function was used to predict issue weights of the opponent's preferences by Coehoorn and Jennings [12].

Recently, Baarslag [7–9] showed that, despite their simplicity, frequency models tend to outperform in practice more complex approaches like Bayesian opponent models. Part of this success can be attributed to the fact that frequency approaches tend to make less assumptions about the opponent behavior, and the fact that frequency models allow for more exploration of the negotiation space due to its quicker computation with respect to Bayesian approaches. In this paper, we further study how to improve the efficiency of the classic frequency model by alleviating the effect of some of its assumptions.

3 The Classical Frequency Model

As mentioned in the previous sections, frequency models have been widely used as opponent modeling mechanisms in automated negotiation [1,2,7,18,25]. Apart from being reported as one of the most effective families of opponent modeling techniques, frequency approaches have the advantage of being simple and offering a good balance of time/exploration [7].

There have been multiple implementations and variations of frequency models in the automated negotiation literature, but perhaps the most popular implementation is the frequency model in HardHeaded's agent [25]. The model was proposed with the following assumptions in mind: (1) the opponent steadily restricts the offers proposed to a possibly moving and decreasing utility range; (2) the opponent prefers to explore the negotiation space rather than repeating the same offer(s) over and over; (3) Opponents tend to concede less on the most preferred issues, keeping them unchanged.

Briefly, this model works as follows. To estimate the weight of an issue value (e.g. Dell, HP, MAC for "laptop brand"), the frequency model computes how often each issue value appears in the opponent's bids. The weight of the issue value is then normalized by the most repeated issue value. For instance, consider that Dell, HP and MAC appear 20 times, 10 times and 15 times respectively. In that case, the model estimates the issue value weights as $V(\text{Dell}) = \frac{20}{20}$, $V(\text{HP}) = \frac{10}{20}$ and $V(\text{MAC}) = \frac{15}{20}$. The frequency model analyzes how often the value of an issue changes. At the beginning, the assumption is that each issue has the same importance. For example, if we have four issues, the weight of each issue is set as 0.25. For each successive pair of offers made by the opponent, if the value of an issue did not change, then the model increases the weight of that issue.

While these assumptions may sound appropriate for some scenarios, the truth is that many state of the art agents do not fully comply with those assumptions. Firstly, agents may prefer to follow a more flexible concession and bidding strategy that allows them to stochastically explore a wide portion of the negotiation space. Despite the fact that the general trend for the opponent is conceding, consecutive offers may not reflect this general trend due to the stochastic and flexible nature of agents. As a consequence, opponents may make a range of negotiation steps (e.g., concession, trade-off, unfortunate move, etc.), misleading the learning mechanism in the classic frequency model.

Secondly, another common behavior of the state-of-the-art agents is repeating the same set of offers for a long period of time. This is true for those agents that try to avoid exploitation by not leaking significant and full information about their utility functions. This goes against the classic frequency's assumptions and the original model is not ready for dealing with these cases. In fact, we experimentally found that, when the same bid is repeated for a significant number of rounds, the update and normalization rules in the classic frequency implementation presents a convergence problem: all the issue weights converge towards $\frac{1}{n}$, where n is the number of issues in the negotiation. Being hard headed and repeating the same offer does not mean that all of the issues are equally valued. Hence, other mechanisms are necessary to tackle these situations.

Last, but not least important, it is true that in the first negotiation rounds most agents do not tend to vary the value for those issues that are the most important. The reason for this is because most agents start by demanding the best offers for themselves, and these offers entail very little changes in the most important issues. However, as the negotiation proceeds, opponents may concede. At some point, it is possible to reach one's own aspirations by varying the values for the most important issues and maximizing less important issues. In fact, this behavior can be observed in many state of the art agents who trade-off issues to achieve one's own aspirations. Therefore, the assumption that opponents tend to concede less on the most preferred issues may hold for hardheaded agents or agents that do not steadily concede, but it may result fruitless in other scenarios.

4 Distribution-Based Frequency Model

Our proposed frequency model relies on the comparison of frequency distributions across negotiation windows. Hence, we have taken the liberty of naming it *distribution-based frequency model*. Next, we describe the details of our opponent modeling mechanism.

4.1 Negotiation Setting

For the sake of simplicity we assume that two agents negotiate following the alternating offers protocol. Nevertheless, the model can also be extrapolated to other protocols, including multilateral scenarios. The agents negotiate in a time-bounded scenario where T delimits the end of the negotiation. If the deadline is reached without any agreement, the agents get their reservation utility.

The negotiation scenario consists of $\mathcal{AT} = \{1, 2, \ldots, n\}$ negotiation issues whose domain values are represented by $\mathcal{D} = \{D_1, \ldots, D_n\}$. An offer is represented by o, while $\mathcal{O}^*$ represents the set of all the offers in the negotiation domain. The agents' preferences are represented by means of linear additive utility functions in the form of $\mathcal{U}(o) = \sum_{i \in \mathcal{AT}} w_i \times V_i(o_i)$, where w_i represents the importance of the negotiation issue i, o_i represents the value for issue i in offer o, and $V_i(.)$ is the valuation function for issue i, which returns the desirability of the issue value. Without losing generality, it is assumed that $\sum_{i \in \mathcal{AT}} w_i = 1$ and the domain of $V_i(.)$ is (0,1) for any i.

There are two main components that an opponent model should estimate in a linear additive function scenario: a vector of weights $\hat{\boldsymbol{\mathcal{W}}} = (\hat{w_1}, \ldots, \hat{w_n})$ representing the estimation of the importance given by the opponent to the different negotiation issues, and an estimation for every possible valuation function $\hat{V_i}(.)$. Thus, our model defines update mechanisms for both.

4.2 Value Function Estimation

Firstly, we describe how the the valuation functions $\hat{V_i}(.)$ are estimated. It should be highlighted that we employ a similar strategy to the one outlined in [25]. The rationale behind our estimation is that, in opponents' offers, the most preferred issue values should appear more frequently than less preferred issue values. Hence, a frequency count of the issue values should provide an educated guess on the real valuation functions $V_i(.)$. We define the estimation of the valuation functions as:

$$\hat{V_i}(j) = \frac{(1 + \sum_{o \in \mathcal{O}_{1 \rightarrow t}} \delta_i(j, o))^{\gamma}}{max_{k \in \mathcal{AT}}(1 + \sum_{o \in \mathcal{O}_{1 \rightarrow t}} \delta_i(k, o))^{\gamma}} \tag{1}$$

where $\delta_i(j, o)$ is 1 if the value j is used for issue i in offer o and 0 otherwise. Please note that the frequency count is smoothed by using a Laplace approach. The rationale behind the smoothing is avoiding crisp distributions and giving importance to issue values that do not appear in $\mathcal{O}$, as they may not appear due to the limited nature of $\mathcal{O}$. On the other hand, both denominator and numerator are passed by an exponential filter with $0 < \gamma \leq 1$ exponential filter. The idea is that of slowing the growth of unbalanced value distributions when opponents send the same offer over and over for a significant part of the negotiation. When $\gamma = 1$ the value estimation is equivalent to the value estimation proposed in the classic frequency model plus a Laplace smooth.

4.3 Issue Weight Estimation

The main differences between our opponent modeling technique and that described in [25] resides in the estimation of the issue weights $\hat{\boldsymbol{\mathcal{W}}}$. In order to provide a more robust estimation of the issue weights, our strategy analyzes consecutive and disjoint windows of the negotiation history instead of individual offers. As mentioned, many of the most popular negotiating agents do not steadily concede but often fluctuate in the demanded utility, even though as a general trend they may concede. By analyzing pairs of offers, the agent may be misled by such stochastic fluctuations and it may end up updating the model incorrectly. However, when analyzing disjoints windows of the negotiation history the effect of such stochastic fluctuations should be alleviated, and general trends better observed.

We divide the current negotiation history into consecutive and disjoint windows of k offers received from the opponent, as it can be observed in Fig. 1. The rationale behind our strategy is comparing the offers in the last window with the offers in the previous window[1]. If the distribution of offers is different between both windows, then it suggests that the opponent has moved its negotiation strategy (e.g., concession, trade-off, etc.). By comparing distributions of offers, one alleviates the problem of stochastic variations between pairs of offers in the strategy of opponents, and it also helps to observe general trends and changes in the opponent strategy.

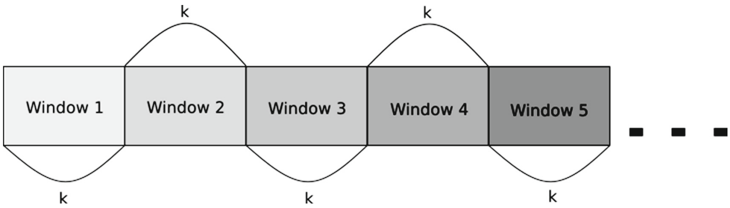

Fig. 1. The negotiation history divided into disjoint windows of offers, each containing k offers sent by the opponent

The classic frequency approach considers that negotiation issues that remain the same between pairs of offers are normally those that are the most relevant. Despite the fact that this may be true in the initial rounds of the negotiation, as the negotiation proceeds opponents may decide to concede on the most preferred issues and achieve its aspirations with less important issues. A classic frequency approach can be misled by this type of behavior, which is not so uncommon in many state of the art agents. As a countermeasure to this behavior, we introduce an issue weight update rule whose effect decays over time. The update rule can be observed in Eq. 2. This update rule will be used to update weights whose value distribution did not change over consecutive windows.

$$\Delta(t) = \alpha \times (1 - t^{\beta}) \tag{2}$$

The issue weight estimation is triggered whenever a new window of k disjoint opponent offers is completed. The outline of the mechanism can be observed in Algorithm 1. As mentioned, the algorithm takes the two latest consecutive and disjoint windows of k offers ($\mathcal{O}$, $\mathcal{O}'$), the current negotiation time t, and the current estimation of the issue weights $\boldsymbol{\mathcal{W}'}$. Before explaining the algorithm in

[1] Please note that windows are not overlapping and not sliding.

Data: t: The current time in the negotiation, $\mathcal{O}'$: The previous partition of k offers, $\mathcal{O}$: The current partition of k offers, $\mathcal{O}_{1\to t}$: All the offers received so far, $\mathcal{W}' = \{w'_1, \ldots, w'_n\}$: The current weights for the opponent model
Result: $\mathcal{W} = \{w_1, \ldots, w_n\}$: The new weights for the opponent model

1 $e \leftarrow \emptyset$;
2 concession$\leftarrow$ $False$;
3 **foreach** $i \in \mathcal{N}$ **do**
4     $w_i \leftarrow w'_i$
end
5 **foreach** $i \in \mathcal{AT}$ **do**
6     $\boldsymbol{F'_i} \leftarrow (Fr_i(1, \mathcal{O}'), \ldots, Fr_i(n, \mathcal{O}'))$;
7     $\boldsymbol{F_i} \leftarrow (Fr_i(1, \mathcal{O}), \ldots, Fr_i(n, \mathcal{O}))$;
8     $p_{val} \leftarrow \mathcal{X}^2$-test$(\boldsymbol{F_i} = \boldsymbol{F'_i})$;
9   **if** $p_{val} > 0.05$ **then**
10         $e \leftarrow e \cup \{i\}$;
    else
11         $\boldsymbol{\mathcal{V}_i} \leftarrow (\hat{V}_i(1), \ldots, \hat{V}_i(n))$;
12         $E[\mathcal{U}_i(\mathcal{O}')] \leftarrow \boldsymbol{\mathcal{V}_i} \times \boldsymbol{F'_i}$;
13         $E[\mathcal{U}_i(\mathcal{O})] \leftarrow \boldsymbol{\mathcal{V}_i} \times \boldsymbol{F_i}$;
14   **if** $E[\mathcal{U}_i(\mathcal{O})] < E[\mathcal{U}_i(\mathcal{O}')]$ **then**
15   concession $\leftarrow$ $True$;
        end
    end
end
16 **if** $|e| \neq n$ ***and*** *concession*$=$ $True$ **then**
17   **foreach** $i \in e$ **do**
18         $w_i \leftarrow w'_i + \Delta(t)$
    end
end

Algorithm 1. The issue weight update mechanism

detail, we need to define the following equation that defines the frequency of a negotiation value j of issue i in a window of offers $\mathcal{O}$:

$$Fr_i(j, \mathcal{O}) = \frac{1 + \sum_{o \in \mathcal{O}} \delta_i(j, o)}{n + |\mathcal{O}|} \tag{3}$$

The equation above counts the number of times that a value j appears in a window of offers, and divides by the total number of offers in the window. Again, the count is smoothed using Laplacian smoothing. This formula will be used in Algorithm 1 to provide a frequency distribution for issue values given a window of offers. Next, we explain the algorithm for updating issue weights in detail.

Initially, the new estimation for the weights $\boldsymbol{\mathcal{W}}$ takes the value of the current estimation $\boldsymbol{\mathcal{W}'}$ in lines 3 and 4. Then, the algorithm iterates over every single negotiation issue i from lines 5 to 15. In this loop, we calculate the frequency distribution of the issue values in the previous window $\boldsymbol{\mathcal{F}'_i}$ and the frequency distribution of the issue values in the current window $\boldsymbol{\mathcal{F}_i}$. Both frequency distributions are calculated by applying the expression in Eq. 3 to every single possible value in the domain of issue i. Then, a Chi-squared test is carried out with the null hypothesis being that both frequency distributions, $\boldsymbol{\mathcal{F}_i}$ and $\boldsymbol{\mathcal{F}'_i}$, are statistically equivalent. The main goal behind this test is checking whether or not the distribution of issue values for i has changed from the previous window of offers to the current one. This information will help us to determine if, overall, the opponent has changed the type of offers sent. In the case that the null hypothesis cannot be rejected (lines 9 and 10), we add the issue i to the set of issues e whose distribution did not change from the previous to the current window.

When the null hypothesis is rejected (lines 11 to 15), it means that the frequency distribution for issue i has been different from the past to the current window. The question is in what direction the change points for that issue (e.g., concession, increase of utility). More specifically, inspired by classic frequency approaches, we are interested in checking if the opponent has conceded in the issue, because then we can update the weights for those issues that remained the same. Again, the assumption is that opponents tend not to change the most important issues more often than less preferred issues. In order to estimate if the opponent has conceded in the issue, we employ the frequency distribution for issue i during the whole negotiation $\mathcal{V}_i$ as an approximation of the real valuations, as specified in Eq. 1. Then, the expected utility obtained in issue i for the previous window of opponent offers $E[\mathcal{U}_i(\mathcal{O}')]$ is calculated in line 12. The same procedure is applied to obtain the expected utility obtained in issue i for the current window $E[\mathcal{U}_i(\mathcal{O})]$. Then, both expected utilities are compared to assess if a concession has been carried out in the issue i.

We take an agressive strategy to detecting overall concessions over two consecutive windows of opponent offers. We consider that there is a concession as long as the opponent has conceded in one of the issues (line 16). In that case, we update the importance for those issues that stayed in the same frequency distribution (lines 17 and 18). We understand that there are other strategies to detect an overall concession, and we are currently exploring the performance of more conservative approaches and probabilistic approaches.

5 Experiments

In this section we evaluate the performance of the proposed frequency mechanism, the distribution-based frequency model, and compare it with the classic frequency opponent modeling mechanism. The goal of this section is assessing whether or not the proposed strategy is capable of overcoming the shortcomings highlighted in the previous sections. First, we describe how the experiments were designed, and then we analyze the results gathered.

5.1 Experimental Design

Given the fact that the goal of this paper is comparing the performance of two learning mechanisms (classic frequency model, and distribution-based frequency model), we decided to use the same bidding and concession strategies for both strategies. This setting allows us to study both learning mechanisms in fair and equal conditions. More specifically, we chose HardHeaded's concession and bidding strategy. The rationale for selecting this strategy is twofold. First, as the agent employs a Boulware strategy [15], it guarantees that the agent will not rapidly end the negotiation. A quick and abrupt end of the negotiation would preclude learning mechanisms from being studied effectively, as they have not been exposed to sufficient bids. Second, the HardHeaded's bidding strategy actively employs the opponent model to propose bids to the opponent. This is important, as many times opponent modeling will have an impact on the opponent's actions and bidding steps. By taking this realistic setting, we are also able to observe whether or not one's opponent model influences the negotiation towards actions that further improve one's opponent model. For the sake of simplicity, we decided to employ the alternating offers protocol although the opponent modeling mechanism should be applicable to other settings as long as multiple offers are exchanged between parties over time. Accordingly, we employed bilateral negotiation domains to experimentally test the performance of our opponent modeling mechanism. More specifically, we decided to test our modeling mechanism under a wide range of domain characteristics. These characteristics include different domain sizes (i.e., number of possible outcomes) and different degrees of competition between agents, measured by the distance from the Kalai point to the complete satisfaction point (1,1) [9]. The list of domains can be found in Table 1. Discount factors were ignored (i.e., removed) in this experimental setting and they are regarded as a matter of future study.

Table 1. The domains chosen for testing our opponent modeling mechanism

	Laptop	CypressI.	EngvsZim.	Grocery	Amsterdam	Camera	S.market	Travel
Size	Small	Small	Medium	Medium	Medium	Large	Large	Large
Conflict	Low	High	Medium	Low	Low	Low	High	Medium

Another important decision to take for the experimental setting was deciding on the opponent agents to negotiate with. There were some factors that influenced our decision in this matter. First of all, any agent with offline opponent modeling was discarded as these may introduce interdependences between the outcomes of different negotiations, leading to effects in our opponent model that may be the by-product of past negotiation interactions. Second, we wanted to expose our opponent modeling mechanism to a variety of concession behaviors and bidding strategies representing the state of the art in negotiation. For that, we employed the following opponent agents:

– AgentK [24]: This agent was the winning agent of the 2010 ANAC negotiation competition. It is a conceder agent whose concession speed is regulated by the average utility of all received bids and its standard deviation. In terms of bid proposal, it just selects any offer from above the current aspiration. Hence, pairs of consecutive offers may not present an obvious decreasing trend. This behavior is in conflict with assumption 1 in Sect. 3.

- IAmHaggler2011 [30]: A negotiating agent that uses Gaussian processes to predict the future concession of its opponent, and then adjust its concession rate accordingly to get the most from the negotiation. The goal of this agent is that of optimizing one's own utility while also trying to reduce the utility received by the opponent. Bids are only selected from a small range around the target utility. This aggressive stance usually results in the agent repeating the same offers. This is contradiction with assumption 2 in Sect. 3.
- TheNegotiatorReloaded [13]: This agent was the best performing agent in undiscounted domains for the ANAC 2012 agent competition. The agent divides the negotiation into non-sliding windows, similarly to our approach. For each window, the agent estimates the type of agent behavior that it is facing and adjusts its concession rate accordingly. The most similar bids to the current target utility are sent back to the opponent. This agent was selected for the experiment for similar reasons to IAmHaggler2011.
- Boulware agent [15]: A classic negotiation agent that adjusts its aspiration levels according to time, only conceding in the later stages of the negotiation. Bids close to the target utility are selected and sent to the opponent. This agent was included as an example of scenarios where none of the aforementioned assumptions are strictly violated.
- Conceder agent [15]: A classic negotiation agent that adjusts its aspiration levels according to time. However, concessions are carried out early on in the negotiation process. Bid are randomly selected from above the threshold defined by the concession strategy. Due to the rapid concessions at the start of the negotiation, assumption 3 in Sect. 3 may be invalid very quickly.

The platform that supported our experiments was Genius [20]. We compared our opponent modeling with the performance of the classic frequency model. For that, both opponent models faced all of the opponents in every single domain, which included two preference profiles per domain. In order to capture stochastic variations in negotiations, each possible case was repeated a total of 20 times. This gives a total of 3200 negotiations[2].

In order to assess the quality of our opponent modeling mechanism, we employed the following quality metrics:

- Pearson correlation of bids: It aims to compare the estimated outcome space with the real outcome space. For that, the Pearson correlation of bids is calculated and averaged. This metric is employed due to the fact that it has a strong correlation with overall opponent modeling performance [7,8].
- Difference in surface of Pareto frontiers: Another metric that is employed to assess the overall performance of an opponent modeling mechanism is the absolute difference between the area under the real Pareto optimal curve and the area under the estimated Pareto optimal curve. The rationale behind this metric is that some claim that it is enough to accurately estimate the

[2] 2 models $\times$ 5 opponents $\times$ 8 domains $\times$ 2 profiles $\times$ 20 repetitions.

Pareto optimal frontier for succesful negotiations with the opponent. Again, this metric was shown to have a strong correlation with overall opponent modeling performance [7,8].

– Spearman rank correlation of the issue weights: The previous two metrics offer an insight into how the opponent modeling performs overall. However, both our opponent modeling mechanism and the classic frequency model have two components: the weight and the issue value update mechanisms. Therefore we decided to include extra metrics to assess the performance of each of these individual components. This metric compares the rank correlation between the issue weights learned by an opponent modeling mechanism and the target issue weights. The value is between −1 and 1, with 1 being used for a perfect ranking of the issue weights and −1 for a completely opposite ranking. The rationale for selecting a ranking metric is that a ranking of the issues is normally enough to trade-off [14,21].
– Weighted Root Mean Squared Error of the issue values: Given an estimated model and the target model, this metric computes a weighted version of the root mean squared error (RMSE) per issue. Predicted issue values are compared against the target issue values and weighted according to the importance of the issue value. This can be observed in Eq. 4

$$\text{WRMSE}(i) = \sqrt{\sum_{j \in D_i} w_j \times (\hat{V}_i(j) - V_i(j))^2} \tag{4}$$

In our case, the weights for issue values w_j were set to $\frac{V_i(j)}{max_{k \in D_i} V_i(k)}$ so that more weight is given to those issue values that provide more utility to the opponent. Then, after each negotiation, the metric is averaged with all the issues. In this case we employed a metric that both captures ranking and value accuracy. Although a ranking of issues is enough for carrying out trade-offs, one needs to have an accurate estimation of values for successfully providing appealing offers to the opponent.

As for the parameters of our model, we set $\alpha = 10$ and $\beta = 5$ in Eq. 2. This means that weight updates will have a greater magnitude at the start of the negotiation, and they will gradually be reduced as the negotiation finishes. This type of update is meant to avoid incorrect updates when the opponent starts changing the most important issues relatively soon in the negotiation. With regards to Eq. 1, γ was set to 0.25 to slow the growth of value importance when the opponent tends to repeat the same offer repeatedly. These values were found as good in a previous experimental setup. However, no exhaustive search was carried out over them. Therefore, the performance depicted in these experiments should be considered as a lower bound for the best achievable performance with this opponent modeling mechanism.

5.2 Results

In this section we analyze the performance of our opponent modeling mechanism with respect to the classic frequency model. As mentioned we employ four

Table 2. Results obtained for the Pearson correlation of bids (Prs. B.), the difference in surface of the Pareto optimal frontier (Par. Fr. D.), the Spearman rank correlation of the issue weights (Spr. W.), and the weighted root mean squared error of the issue values (WRMSE), aggregated by domain.

	Distribution-based frequency model				Frequency model			
	Prs. B.	Par. Fr. D.	Spr. W.	WRMSE	Prs. B.	Par. Fr. D.	Spr. W.	WRMSE
EngZimb	**0.91**	**0.003**	**0.35**	**0.053**	0.80	0.007	0.32	0.080
Cypress	**0.83**	**0.043**	**0.46**	**0.042**	0.60	0.092	0.30	0.063
Travel	**0.85**	**0.015**	**0.73**	**0.043**	0.70	0.019	0.21	0.057
Amst.	**0.91**	**0.004**	**0.61**	**0.044**	0.86	0.015	0.31	0.072
Grocery	**0.90**	**0.005**	**0.96**	**0.041**	0.86	0.010	0.54	0.046
Laptop	0.87	0.006	**0.84**	0.094	**0.89**	0.006	0.59	0.10
Camera	**0.89**	0.006	**0.77**	**0.035**	0.86	0.005	0.57	0.041
S.Market	**0.87**	**0.044**	**0.93**	**0.030**	0.69	0.067	0.59	0.050

Table 3. Results obtained for aforementioned metrics, aggregated by opponent.

	Distribution-based frequency model				Frequency model			
	Prs. B.	Par. Fr. Dis.	Spr. W.	WRMSE	Prs. B.	Par. Fr. Dis.	Spr. W.	WRMSE
AgentK	**0.91**	**0.007**	**0.76**	**0.040**	0.83	0.019	0.58	0.065
Haggler	**0.93**	**0.005**	**0.78**	0.035	0.88	0.014	0.32	0.036
TNR	**0.79**	0.024	**0.63**	**0.079**	0.74	0.018	0.50	0.092
Boulw.	**0.91**	**0.007**	**0.70**	**0.033**	0.78	0.017	0.35	0.072
Conc.	**0.84**	**0.035**	**0.67**	0.050	0.68	0.070	0.38	0.053

metrics: two that measure the overall quality of the model (i.e., Pearson correlation of bids, and the difference between the surfaces defined by the Pareto optimal frontiers) and two other metrics that assess the quality of the two individual components of the opponent modeling (i.e., Spearman rank correlation of the issue weights, and the weighted root mean squared error). Table 2 aggregates the results obtained by domain, while Table 3 aggregates the results obtained by opponent. Those results that are statistically better than its counterpart are highlighted with a bold font. For statistical significance, a one-tailed Mann-Whitney test was carried out with $\alpha = 0.05$.

First, we will analyze the results per domain. Both the Pearson correlation of bids, and the difference between Pareto optimal frontiers tend to indicate that our opponent modeling provides a more accurate and overall estimation of the opponent's preferences. In the case of the Pearson correlation of bids, our approach is statistically better for all domains except for the Laptop domain, where the classic frequency model obtains a better estimation. However, the difference between both metrics is small for that domain. Similarly, our model outperforms when it comes to estimating the Pareto optimal frontier in all domains except for the Laptop and Camera domain, where there is no difference between our and the classic frequency model. Both the Laptop and the Camera domain are two of the less competitive domains. This may suggest that our opponent modeling may not necessarily outperform the classic frequency model for domains with low conflict. However, further experiments will be needed to make that conclusion. In this very same table, we can also observe that both our weight and issue value estimation are statistically more accurate than the classic frequency model. This is supported by statistically better results in both the Spearman rank correlation of issue weights and the weighted root mean squared error for issue values. Only in the Laptop domain the issue value estimation is no better, but also no worse, than the classic frequency model.

We can observe very similar results if we focus on the results aggregated by opponent. Overall, our opponent model produces a statistically better and more accurate model of the opponent's preferences (i.e., Pearson correlation of bids, and Pareto frontier estimation). Only in the case of The Negotiation Reloaded (TNR) the estimation of the Pareto optimal frontier is no better, but also no worse than the classic frequency model. Component by component, we can appreciate that our weight update mechanism is consistently more accurate at detecting the relative importance of issues (i.e., Spearman rank correlation of weights). We also tend to produce a better value estimation for issue values for most opponent agents (i.e., WRMSE). Only we produce a statistically equivalent estimation against IAmHaggler2011 and the Conceder agent.

Overall, it can be appreciated that our opponent modeling mechanism tends to produce more accurate models of the opponent's preferences, regardless of opponent and domain. The differences tend to be more acute in the weight update mechanisms (i.e., Spearman rank correlation of weights) than in the issue value update mechanism (i.e., WRMSE). This suggests that, most likely, an important part of our improvement is due to the weight update mechanism. Other issue value update mechanisms may be necessary to further improve the classic frequency model. Further exploring other issue value update mechanism is highlighted as future areas of improvement for our current opponent modeling.

6 Conclusions

In the last few years, frequency modeling has been shown to outperform more sophisticated opponent modeling techniques like Bayesian approaches. The reason for this result is, among others, weaker assumptions on the opponent's behavior. Nevertheless, frequency models still rely on some underlying assumptions that may not be fully realistic in many scenarios. In this paper we have presented a new frequency approach to opponent modeling in automated negotiation. This new approach, which we named as distribution-based frequency model, soothes the effect of some of the assumptions in the classic frequency model. More specifically, the main characteristics of our opponent modeling are: (i) comparison of windows of offers instead of consecutive pairs of offers, offering a more robust estimation on the opponent's behavior; (ii) decayed weight update to avoid incorrect updates when the opponent starts conceding on the most important issues; and (iii) slow growth of issue values importance, avoiding unbalanced issue value distributions when the opponent offers the same offer repeatedly.

This paper advances the state of the art in frequency approaches to opponent modeling in automated negotiation by showing that the model proposed in this work outperforms the accuracy of the opponent obtained by the classic frequency model. The increased accuracy is observable in the learned outcome space, the estimated Pareto optimal frontier, and both in the learned issue weights and issue values. The difference in accuracy is specially acute in estimated issue weights, where our approach produces more accurate rankings of issues.

References

1. Afiouni, E.N., Øvrelid, L.J.: Negotiation for strategic video games. Master's thesis, NTNU (2013)
2. Alsrheed, F., El Rhalibi, A., Randles, M., Merabti, M.: Intelligent agents for automated cloud computing negotiation. In: IEEE International Conference on Multimedia Computing and Systems, pp. 1169–1174. IEEE (2014)
3. An, B., Gatti, N., Lesser, V.: Extending alternating-offers bargaining in one-to-many and many-to-many settings. In: Proceedings of the 2009 IEEE/WIC/ACM International Conference on Intelligent Agent Technology, vol. 2, pp. 423–426 (2009)
4. Aydoğan, R., Baarslag, T., Hindriks, K.V., Jonker, C.M., Yolum, P.: Heuristics for using cp-nets in utility-based negotiation without knowing utilities. Knowl. Inf. Syst. **45**(2), 357–388 (2015)
5. 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. SCI, vol. 674, pp. 153–167. Springer, Cham (2017). doi:10.1007/978-3-319-51563-2_10
6. Aydoğan, R., Yolum, P.: Ontology-based learning for negotiation. In: IEEE/WIC/ACM International Conference on Intelligent Agent Technology, pp. 177–184 (2009)
7. Baarslag, T.: Measuring the performance of online opponent models. In: Exploring the Strategy Space of Negotiating Agents. ST, pp. 111–127. Springer, Cham (2016). doi:10.1007/978-3-319-28243-5_6
8. Baarslag, T., Hendrikx, M., Hindriks, K., Jonker, C.: Predicting the performance of opponent models in automated negotiation. In: International Joint Conference on Web Intelligence and Intelligent Agent Technologies, vol. 2, pp. 59–66. IEEE (2013)
9. 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. Auton. Agent Multi Agent Syst. **30**, 849–898 (2016)
10. Buffett, S., Spencer, B.: Learning opponents' preferences in multi-object automated negotiation. In: Proceedings of the 7th International Conference on Electronic Commerce, pp. 300–305 (2005)
11. Bui, H.H., Kieronska, D., Venkatesh, S.: Learning other agents' preferences in multiagent negotiation. In: Proceedings of the National Conference on Artificial Intelligence, pp. 114–119 (1996)
12. Coehoorn, R.M., Jennings, N.R.: Learning an opponent's preferences to make effective multi-issue negotiation tradeoffs. In: The 6th International Conference on E-Commerce, pp. 59–68 (2004)

13. Dirkzwager, A., Hendrikx, M.: An adaptive negotiation strategy for real-time bilateral negotiations. 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. SCI, vol. 535, pp. 163–170. Springer, Tokyo (2014). doi:10.1007/978-4-431-54758-7_10
14. Faratin, P., Sierra, C., Jennings, N.R.: Using similarity criteria to make issue trade-offs in automated negotiations. Artif. Intell. **142**(2), 205–237 (2002)
15. Faratin, P., Sierra, C., Jennings, N.R.: Negotiation decision functions for autonomous agents. Robot. Auton. Syst. **24**(3–4), 159–182 (1998)
16. Fatima, S., Kraus, S., Wooldridge, M.: Principles of Automated Negotiation. Cambridge University Press, Cambridge (2014)
17. Fatima, S.S., Wooldridge, M., Jennings, N.R.: An agenda-based framework for multi-issue negotiation. Artif. Intell. **152**(1), 1–45 (2004)
18. van Galen, L.N.: Agent smith: opponent model estimation in bilateral multi-issue negotiation. In: Ito, T., Zhang, M., Robu, V., Fatima, S., Matsuo, T. (eds.) New Trends in Agent-Based Complex Automated Negotiations. SCI, vol. 383, pp. 167–174. Springer, Heidelberg (2012). doi:10.1007/978-3-642-24696-8_12
19. Gerla, M., Lee, E.K., Pau, G., Lee, U.: Internet of vehicles: from intelligent grid to autonomous cars and vehicular clouds. In: IEEE World Forum on Internet of Things, pp. 241–246 (2014)
20. Hindriks, K., Jonker, C.M., Kraus, S., Lin, R., Tykhonov, D.: Genius: negotiation environment for heterogeneous agents. In: Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems. pp. 1397–1398 (2009)
21. Hindriks, K., Tykhonov, D.: Opponent modelling in automated multi-issue negotiation using bayesian learning. In: 7th International Joint Conference on Autonomous Agents and Multiagent Systems, pp. 331–338 (2008)
22. Ikarashi, M., Fujita, K.: Compromising strategy using weighted counting in multi-times negotiations. In: Proceedings of the 3rd International Conference on Advanced Applied Informatics, pp. 453–458 (2014)
23. Jennings, N.R., Faratin, P., Lomuscio, A.R., Parsons, S., Wooldridge, M.J., Sierra, C.: Automated negotiation: prospects, methods and challenges. Group Decis. Negot. **10**, 199–215 (2001)
24. Kawaguchi, S., Fujita, K., Ito, T.: AgentK: Compromising strategy based on estimated maximum utility for automated negotiating agents. In: Ito, T., Zhang, M., Robu, V., Fatima, S., Matsuo, T. (eds.) New Trends in Agent-Based Complex Automated Negotiations. SCI, vol. 383, pp. 137–144. Springer, Heidelberg (2012). doi:10.1007/978-3-642-24696-8_8
25. van Krimpen, T., Looije, D., Hajizadeh, S.: HardHeaded. In: Ito, T., Zhang, M., Robu, V., Matsuo, T. (eds.) Complex Automated Negotiations: Theories, Models, and Software Competitions. SCI, vol. 435, pp. 223–227. Springer, Heidelberg (2013). doi:10.1007/978-3-642-30737-9_17
26. Luo, X., Jennings, N.R., Shadbolt, N., Leung, H.F., Lee, J.H.M.: A fuzzy constraint based model for bilateral, multi-issue negotiations in semi-competitive environments. Artif. Intell. **148**(1), 53–102 (2003)
27. Ossowski, S., Sierra, C., Botti, V.: Agreement technologies: a computing perspective. In: Ossowski, S. (ed.) Agreement Technologies. LGTS, vol. 8, pp. 3–16. Springer, Dordrecht (2013). doi:10.1007/978-94-007-5583-3_1
28. Sanchez-Anguix, V., Aydogan, R., Julian, V., Jonker, C.: Unanimously acceptable agreements for negotiation teams in unpredictable domains. Electron. Commer. Res. Appl. **13**(4), 243–265 (2014)

29. Sanchez-Anguix, V., Julian, V., Botti, V., García-Fornes, A.: Tasks for agent-based negotiation teams: analysis, review, and challenges. Eng. Appl. Artif. Intel. **26**(10), 2480–2494 (2013)
30. Williams, C.R., Robu, V., Gerding, E.H., Jennings, N.R.: IAMhaggler2011: a gaussian process regression based negotiation agent. In: Ito, T., Zhang, M., Robu, V., Matsuo, T. (eds.) Complex Automated Negotiations: Theories, Models, and Software Competitions. SCI, vol. 435, pp. 209–212. Springer, Heidelberg (2013). doi:10.1007/978-3-642-30737-9_14