



Conflict-based negotiation strategy for human-agent negotiation

Mehmet Onur Keskin¹ · Berk Buzcu¹ · Reyhan Aydoğan^{1,2} 

Accepted: 6 September 2023 / Published online: 3 November 2023
© The Author(s) 2023

Abstract

Day by day, human-agent negotiation becomes more and more vital to reach a socially beneficial agreement when stakeholders need to make a joint decision together. Developing agents who understand not only human preferences but also attitudes is a significant prerequisite for this kind of interaction. Studies on opponent modeling are predominantly based on automated negotiation and may yield good predictions after exchanging hundreds of offers. However, this is not the case in human-agent negotiation in which the total number of rounds does not usually exceed tens. For this reason, an opponent model technique is needed to extract the maximum information gained with limited interaction. This study presents a conflict-based opponent modeling technique and compares its prediction performance with the well-known approaches in human-agent and automated negotiation experimental settings. According to the results of human-agent studies, the proposed model outperforms them despite the diversity of participants' negotiation behaviors. Besides, the conflict-based opponent model estimates the entire bid space much more successfully than its competitors in automated negotiation sessions when a small portion of the outcome space was explored. This study may contribute to developing agents that can perceive their human counterparts' preferences and behaviors more accurately, acting cooperatively and reaching an admissible settlement for joint interests.

Keywords Opponent modelling · Preference modelling · Human-agent negotiation · Automated negotiation

1 Introduction

Negotiation is an interaction among self-interested parties that have a conflict of interests and aim to achieve a joint agreement. It can occur daily basis when parties need to make decisions collectively on any matters such as personal activities (e.g., arranging holiday plans), professional procedures (e.g., job interviews, task or resource allocations), or societal matters (e.g., effective energy distribution). Depending on the complexity of the decisions, this process can be time-consuming and cumbersome for human stakeholders. Therefore, researchers in the field of Artificial Intelligence have put their effort into automating this process over the

last decades [1, 8, 14]. Recently, there has been a high interest in human-agent negotiations in which intelligent agents negotiate with their human counterparts [4, 28]. Creating large-scale social impact by such intelligent systems requires understanding how human decisions are made and their preferences and interests [26]. That is, agents should be capable of understanding why their opponent made such offers and what is acceptable to their opponent so that it can adapt its bidding strategy accordingly to increase the chance of reaching mutually beneficial agreements. That shows the importance of the opponent modeling during the negotiation.

There are a variety of opponent modeling techniques proposed in automated negotiation literature [6]. As far as the existing opponent models to predict the opponent's preferences are concerned, it is observed that they attempt to learn a model from bid exchanges and mostly have some particular assumptions about both opponent's bidding behavior and preference model (e.g., having an additive utility function and employing time-based concession strategy). Even simple heuristic models such as the frequentist approach [19, 31] perform well in negotiation. Although there are relatively much fewer offer exchanges in human-agent negotiation in contrast to automated one (i.e., the number of

✉ Reyhan Aydoğan
R.Aydogan@tudelft.nl
Mehmet Onur Keskin
onur.keskin@ozu.edu.tr
Berk Buzcu
berk.buzcu@ozu.edu.tr

¹ Özyeğin University, Istanbul, Turkey

² Computer Science, Interactive Intelligence, TU Delft, Delft, Netherlands

offers typically does not exceed 20-30 offers in human-agent negotiation [22]), some studies adopt their variants in human-agent negotiation [25].

This study pursues an alternative way of modeling human opponents' preferences by searching for cause-effect relationships in human negotiators' bidding patterns. Here, the main challenge is to learn meaningful preference relations that enable our agents to generate offers that are more likely to be acceptable by their opponents despite the small number of offer exchanges. Accordingly, this study proposes a novel conflict search-based opponent modeling strategy mainly designed to learn human opponents' preferences in multi-issue negotiations to generate well-targeted offers leading to mutually beneficial agreements (i.e., high social welfare). The proposed opponent modeling approach has been evaluated experimentally concerning performance metrics, such as the model's accuracy and the model's effect on the negotiation and negotiation outcome. To show the performance of the proposed approach, we conducted two human-agent negotiation experiments involving 70 participants in total and compared our agent performance with those of the aforementioned well-known frequentist approaches [19, 31]. Our results showed that the proposed conflict-based opponent model outperformed them dramatically in terms of their prediction accuracy. Furthermore, we studied the effect of the model on the negotiation outcome in automated negotiation by involving 15 state-of-the-art negotiation agents from the International Automated Negotiating Agents Competition (ANAC) [15] on six different negotiation scenarios. Our results showed that our agent gained the highest average individual utility and social welfare (i.e., both product of utilities and the sum of utilities) on average.

The rest of the paper is organized as follows. Section 2 reviews the related work on opponent modeling. The proposed opponent model is explained in Section 3, and the negotiation strategy of the agent utilizing the model is defined in Section 4. Section 5 presents our experimental setup and analysis of the results. Finally, we conclude our work with a discussion involving future work directions in Section 6.

2 Related work

Automated negotiation has been widely studied for several decades, and a variety of negotiation frameworks have been proposed so far [2, 8]. By their nature, automated agents try to find the most beneficial agreement for both parties by making many consecutive offers up to a particular deadline (time or round). As Hindriks, Jonker, and Tykhonov point out that agents can benefit from learning about their opponent during negotiation [12], a variety of opponent modeling approaches have been proposed in the negotiation community, such as opponent's preferences (e.g., [12, 19, 21, 25,

31, 32]), the acceptability of an offer (e.g., [20, 26, 29]) and negotiation strategy/attitude (e.g., [16, 23, 27]). The main opponent strategy is identifying the opponent's preferences by analyzing offer exchanges between parties. Afterward, the agent examines the opponent's negotiation offers with estimated opponent preferences to get an idea about its strategy/attitude. Various modeling techniques have been used in these strategies, such as kernel density, Bayesian learning, and frequentist models. While building up their model, those opponent modeling approaches rely on some assumptions such as having a predetermined deadline, capturing their preferences in the form of an additive utility function, and following a turn-taking negotiation protocols such as (Stacked) Alternating Offers Protocol [2] and conceding over time (e.g., time-based concession strategies). In the following part, we mention the most relevant works. A more detailed explanation about opponent modeling can be found in the survey [6].

Another common preference model technique in the literature is based on Bayesian learning [12, 32]. Hindriks et al. use Bayesian learning to predict the shape of the opponent's utility function, the corresponding rank of issue values, and issue weights [12]. As an extension of Hindrik's work, Yu et al. incorporate regression analysis into Bayesian learning by comparing the predicted future bids and actual incoming bids. Accordingly, they update the Bayesian belief model by considering both current and expected coming bids.

Recent studies' most common preference modeling strategies are variations of the frequentist models. The winner of the Second Automated Negotiating Agents Competition [15] called Hardheaded agent [19] uses a simple counting mechanism for each issue value and analyzes the contents of the opponent's consequent offers. The main heuristic is that the opponent would concede less on the essential issues while using the preferred values in its offers. Therefore, while analyzing the opponent's current and previous offer, if the value of an issue is changed, that issue's weight is decreased by a certain amount. While such a simple approach is initially intuitive, information loss seems inevitable. Due to the nature of the negotiation, the opponent may need to concede even on important issues. Those moves may mislead the model. In addition, the concession amount may vary during the negotiation, which the frequentist approach needs to capture.

Moreover, suppose the opponent repeats the same bid multiple times. In that case, the model may overvalue those repeated issue values while underestimating the unobserved values (e.g., converging a zero utility since it is not seen). Tunali, Aydoğan, and Sanchez [31] aims to resolve those problems by comparing the windows of offers instead of consecutive pairs of offers and offering a more robust estimation of the opponent's behavior. It adopts a decayed weight update to avoid incorrect updates when opponents concede on the most critical issues. Furthermore, it smoothly increases the importance of issue values to avoid unbalanced issue

value distributions when the opponent offers the same offer repeatedly. Although their approach outperforms the classical frequency approach, it still suffers from only counting the issue value appearance because it ignores the varying utility patterns of the opponent's offers.

Apart from the opponent modeling in automated negotiation, we review the opponent modeling approaches particularly designed for human-agent negotiations. Lin et al. introduce the QOAgent using kernel density estimation (KDE) for modeling opponent's preferences [21]. According to their results, the QOAgent can reach more agreements. In most cases, it achieves better agreements than the human counterpart playing the same role in individual utility. As an extension of QOAgent, Oshrat et al. present an agent unlike most other negotiating agents in the automated negotiation, the KBAgent attempts to utilize previous negotiations with other human opponents having the same preferences to learn the current human opponent [26]. This approach requires an essential assumption that human participants will behave similarly to each other. Thus, KBAgent builds a broad knowledge base from its previous opponents and accordingly offers based on a probabilistic model constructed from the knowledge base utilizing kernel density estimation. In their experimental comparison, the KBAgent outperformed the QOAgent. In these studies, the authors focus on the overall negotiation performance rather than the performance of the proposed opponent modeling approach. However, we examine the accuracy of the opponent modeling and the performance of the whole negotiation strategy.

Furthermore, Nazari, Lucas, and Gratch follow a similar intuition with frequentist models for human-agent negotiation [25]. However, they take into account only the preference ranking of the negotiation issues instead of estimating the overall utility of each outcome. For issue values, they consider a predefined ordering. However, those assumptions may not hold in negotiations where a human participant may have a different evaluation of issue values. In their negotiation, their agent considers the importance of the issues and the expected ordering of the issue values while generating their offers. A similar heuristic with the frequentist approach holds here. That is, an issue is more important if the opponent consistently asks for more on that issue. It leads to the same intrinsic problem of the frequentist approach.

Instead of learning an explicit preference model, some studies focus on understanding what offers would be acceptable for their opponent. Sanchez et al. use Bayesian classifiers to learn the acceptability of partial offers for each team member in a negotiation team [29]. They present a model for negotiation teams that guarantees unanimous decisions consisting of predictable, compatible, and unforeseen issues. The model maximizes the probability of being accepted by both sides. While their model relies on predictable issues such as price, our model is designed to handle unpredictable discrete

issues. Lastly, it is good to mention the reinforcement learning approach proposed for human-agent negotiation [20]. Lewis et al. collect a large dataset consisting of offers represented in natural language from 5808 sessions on Amazon's Mechanical Turk. They present a reinforcement learning model to maximize the agent's reward against human opponents. Accordingly, they aim to estimate the negotiation states acceptable for their human counterparts.

More recently, researchers have been trying to incorporate deep learning models in opponent modeling. For instance, Sengupta et al. has implemented a reinforcement learning-based agent that can adapt to unknown agents per experiences with other agents. In order to model the opponent, they have applied the Recurrent Neural Networks model, specifically LSTM, since they use time series data from the negotiation steps. However, they switched their implementation to a 1D-CNN classifier instead due to data limitations. They observe an opponent agent's bidding strategy according to the agent's self-utility and try to cast it into a class of known behaviors. According to this classification, the agent swaps negotiation strategy within the runtime [30].

Meanwhile, Hosokawa and Fujita expand upon the classical frequentist approach through the addition of the ratio of offers within specified slices of the negotiation timeline, and they implement a weighting function to stabilize the ratios as time passes to capture the change of an opponent's concession toward the end of negotiation [13].

3 Proposed conflict-based opponent modelling (CBOM)

Our opponent modeling called *Conflict-Based Opponent Modeling* (CBOM) aims to estimate the opponent's preferences represented utilizing an additive utility function shown in (1) where w_i represents the importance of the negotiation issue I_i (i.e., issue weight), 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 $\sum_{i \in n} w_i = 1$ and the domain of V_i is $(0, 1)$ for any i . The higher the V_i is, the more preferred an issue value is.

$$\mathcal{U}(o) = \sum_{i=1}^n w_i \times V_i(o_i) \quad (1)$$

Regarding the issue valuation/weight functions (i.e., preferences on issue values), targeting to learn these functions directly from the opponent's offer history may not be a reliable approach since the opponent's negotiation strategy may mislead us. Although the contents of the opponent's offers give insight into which values are more preferred over others, depending on the employed strategy, we may end up with

a different model estimation. For instance, the frequency of the issue value appearance might be a good indicator for understanding the ranking of issue values. However, it is not sufficient to deduce to what extent each issue value is preferred. Fluctuations in the opponent's offers or repeating the exact offers often mislead the agent into accurately estimating the additive utility function. Therefore, we aim first to detect the preference ordering pattern rather than quantifying an evaluation function directly and then interpolate it.

As most of the existing opponent models in the literature do, our model assumes the opponent concedes over time. Initially, the agent does not know anything about its opponent's preferences; therefore, it creates a template estimation model according to the given domain configuration. In other words, the agent starts with an initial belief in ranking the issue values (V_i) and issues (W_i). The agent may assume that the opponent's value function is the opposite of its value function. For instance, If the agent prefers $V_1 > V_2$, it may consider that its opponent prefers $V_2 > V_1$. Alternatively, it may consider an arbitrary ordering for the opponent. Consider that we have n issues and for each issue i there are possible issue values denoted by $D_i = \{v_1^i, \dots, v_m^i\}$. Assuming an arbitrary preference ordering for the opponent, (2) and (3) shows how the initial valuation values and issue weight are initiated. The agent keeps the issues and issue values in order in line with the estimated opponent preferences. Meaning that v_k^i is preferred over v_j^i by the opponent where where $k > j$. Accordingly, (2) assigns compatible evaluation values via max normalization. Similarly, the issue weights are initialized by sum normalization, where each issue weight is in the $[0, 1]$ range, and their sum is equal to 1. Equation (3) ensures that their sum equals one.

$$v_j^i = \frac{j}{|D_i|} \quad (2)$$

$$W_i = \begin{cases} W_{i-1} + \frac{W_{i-1}}{n} & i > \frac{n}{2} \\ W_{i+1} - \frac{W_{i+1}}{n} & i < \frac{n}{2} \\ \frac{1}{n} & i = \lceil \frac{n}{2} \rceil \end{cases} \quad (3)$$

As the agent receives the opponent's offers during the negotiation, it updates its belief incrementally based on the inconsistency between the current model and the opponent's offers. To achieve this, it stores all bids made by the opponent so far, and when a new offer arrives, the current offer is compared with the previous bids with respect to any conflicting ordering. In particular, common and different values in the offer contents are detected. For different values, the system checks whether there is any conflicting situation with the current model. Recall that the current offer is expected to have the less preferred values since the model assumes that the opponent concedes over time. However, according to the

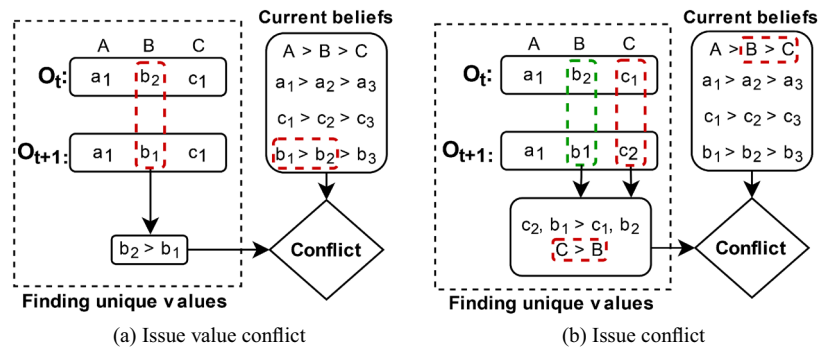
learned model, the ordering may not match the expectation. In such a case, the model is updated.

Two types of conflict in the estimated model could be detected: *issue value conflict* and *issue conflict*. To illustrate those conflicts, let us examine some examples where agents negotiate over three issues (i.e., A , B , and C) in Fig. 1. As current belief indicates $b_1 > b_2 > b_3$ where b_i denotes a possible issue value for the issue B , b_1 is preferred b_2 . In the given negotiation dialogue in Fig. 1a, it can be seen that the opponent's previous offer and current offers are $O_t = \langle a_1, b_2, c_1 \rangle$ and $O_{t+1} = \langle a_1, b_1, c_1 \rangle$, respectively. Agents can examine the contents of the offers and find unique value changes to make some inferences on the preferences. For our case, the only difference in the offers is the value of the issue B . Relying on the assumption that the human negotiator leans towards concession over time, the agent could infer $b_2 > b_1$. Recall that the most preferred values would appear early. As seen clearly, this preference ordering conflicts with that of the agent's belief. We call this type of conflict "issue value conflict" in our study.

The latter conflict type is about the importance of the issues. In the given an example in Fig. 1b, the consecutive offers involve more than one issue value difference, particularly on issues B and C . Then, the agent can deduce $(c_1, b_2) > (c_2, b_1)$ by relying on the concession assumption mentioned above. Individually, ordering in issue C is consistent with the belief (i.e., $c_1 > c_2$); however, the ordering on issue B is conflicting (i.e., $b_2 > b_1$). Therefore, in order to have $(c_1, b_2) > (c_2, b_1)$, the importance of the issue C should be higher than that of B (i.e., $C > B$). This inference conflicts with the current belief of the agent, which says B is more important than C .

Algorithm 1 shows how the ranking of the issue values is extracted. When the opponent makes an offer (O_c), the agent compares the content of the current offer with that of each offer in the opponent's offer history to find the unique values and consequently extract some preferential comparisons (Lines 1–4). Afterward, the agent keeps all those comparisons in a dictionary called CM (Line 3). By reasoning on each comparison in this set by considering the current belief set, each conflict is extracted and stored in AC (Line 5–7). It is worth noting that the method can find issue value conflicts consisting of multiple issues. After keeping track of all possible conflicts, the agent must determine how to update its beliefs. Counting the number of conflicts on each issue value pair, it considers the issue value orderings having the least conflicts and updates its belief accordingly (Lines 8–16). The agent detects issue value pairs in the conflict set for each issue and compares their occurrences to determine which one to stick on. For instance, if the agent observes conflicting information, the more frequent ordering becomes more dominant, and the agent adapts its beliefs accordingly. After updating the beliefs about issue value orderings, it does the same kind

Fig. 1 Preference conflict extraction example



of updates for the issue ordering (Lines 18-25). After finalizing the updates on the rankings, it estimates the utility space of the opponent by utilizing the update operations in (2) & (3).

To illustrate this, we trace the negotiation in Fig. 2, where we can observe how this opponent model works. Following the same domain, we first arbitrarily set our initial beliefs about preference ordering (e.g., $a_3 > a_2 > a_1$ for the values of issue A). The agent keeps track of offers made by the

opponent so far. In our example, you can see the offer history at time $t + 2$. Following, the agent compares all previous offers with each other (i.e., pairwise comparison) and tries to extract an ordering relation. Here, O_t and O_{t+1} denote the first and second offer made by the opponent. Since the agent believes that the opponent’s earlier offers are more preferred over the later offers, it extract that $a_1, b_2 > a_2, b_1$. This knowledge does not give any novel insight to update our beliefs, but we store this ordering for future analysis in the following rounds. When the opponent makes the offer O_{t+2} , the model compares it with all the previous offers pairwise as well as the previously extracted information (e.g., the $a_1, b_2 > a_2, b_1$ relation). Starting from the first offer in the offer history (O_t - O_{t+2}), the model acquires the information of $a_1 > a_2$, since there is only one issue with a different value. When it compares O_{t+1} with O_{t+2} , it extracts $(b_1 > b_2)$ and updates its beliefs accordingly. Similarly, the extracted information could be utilized to reason about the ordering of the negotiation issues (e.g., $A > B$) based on the contradiction between $(a_1, b_2 > a_2, b_1)$ and $(b_1 > b_2)$. Consequently, the agent deduces that the importance of issue A is more than issue B considering the assumption that one-issue comparisons are more reliable than multi-issue comparisons, which conflicts with the current belief and updates its belief accordingly.

Algorithm 1 Conflict-based Opponent Model (CBOM).

```

O: Offer history, Oc: Current Offer, Bc: Current Belief, I:
Issues, Vi: Values;
UV: A list of unique values for a given offer pair;
CM: Comparison map of unique values per each offer pair;
AC: All conflicts extracted from comparison map;
VC: A list of conflicted value pairs for an issue;
IC: A list of conflicted issue pairs;
Uopp: The estimated opponent utility space;
1 for each Oi ∈ O do
2   UV ← findUniqueValues(Oi, Oc);
3   CM.append(UV);
4 end
5 for each c ∈ CM do
6   AC.append(extractConflictsFrom(c, Bc));
7 end
8 for each i ∈ I do
9   VC ← getValueConflictbyIssue(AC, i);
10  for each (p, r) ∈ VC do // p, r ∈ Di
11    if ||VC(r, p)|| > ||VC(p, r)|| then
12      Bc(i) ← r > p;
13    else
14      Bc(i) ← p > r;
15    end
16  end
17 end
18 IC ← getIssueConflict(AC);
19 for each (A, B) ∈ IC do // A, B ∈ I
20   if ||IC(A, B)|| > ||IC(B, A)|| then
21     Bc(I) ← A > B;
22   else
23     Bc(I) ← B > A;
24   end
25 end
26 Uopp ← estimateOppUtilitySpace(Bc);

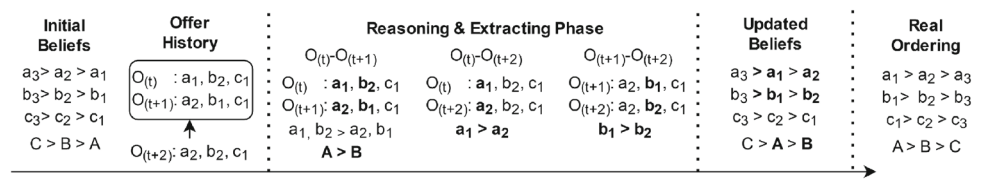
```

4 Proposed conflict-based negotiation strategy

This section presents our negotiation strategy employing the opponent model mentioned above. This strategy incorporates the estimated opponent modeling into the Hybrid strategy [16], which estimates the target utility of the current offer based on time and behavior-based concession strategies.

The Algorithm 2 elaborates how the agent makes its decisions during the negotiation. In each round, it calculates a target utility by employing the hybrid bidding strategy. Consequently, it generates candidate offers that were not offered by the agent (i.e., CBOM Agent). Its utility is in the range of lower and upper target utility (i.e., $TU_{cbom} - \epsilon$ and $TU_{cbom} + \epsilon$) (Lines 1–9). If there is no such an offer, the boundary is

Fig. 2 Example process of the conflict-based opponent model (CBOM)



enlarged with a dynamically generated small number according to the domain size (Line 8). We define a round count n where we believe the agent has enough offers from its opponent to estimate their preferences. The value of n may vary depending on whether the agent negotiates with a human or agent negotiator. If the number of received offers from the opponent is less than n , the agent picks the offer maximizing its utility among potential offers (Lines 10–12). Shortly, the system does not engage the opponent model until there are enough offers accumulated in the history of opponent offers. Otherwise, the agent selects the offer whose estimated utility product is the maximum (Line 14). If the opponent made an offer with a utility higher than our lowest utility bid and the utility of the current candidate’s offer (Line 16), the agent accepted its opponent’s offer instead of making the offer. This acceptance condition is slightly more cooperative than the AC_{next} acceptance strategy. Otherwise, it makes the chosen offer (Line 19).

$$TU_{Hybrid} = (t^2) \times TU_{Times} + (1 - t^2) \times TU_{Behavior} \quad (4)$$

$$TU_{Times} = (1 - t)^2 \times P_0 + [2 \times (1 - t) \times t \times P_1 + t^2 \times P_2] \quad (5)$$

$$TU_{Behavior} = U(O_j^{t-1}) - \mu \times \Delta U \quad (6)$$

$$\Delta U = \sum_{i=1}^4 [W_i \times (U(O_h^{t-i}) - U(O_h^{t-i-1}))] \quad (7)$$

$$\mu = P_3 + t \times P_3 \quad (8)$$

Equation (4) outlines how the agent computes the target utility for its upcoming offer, according to [16]. The concession function (TU_{Times}), represented by (5), incorporates t , the scaled time ($t \in [0, 1]$), and P_0, P_1, P_2 , which correspond to the curve’s maximum value, curvature, and minimum value, respectively. Note that the values of P_0, P_1 , and P_2 in our experiment are 0.9, 0.7, and 0.4, respectively. The behavioral aspect of the "Hybrid" strategy involves scaling the overall utility change by a time-varying parameter, μ , to estimate the target utility, as demonstrated in (4). $U(O_a^{t-1})$ signifies the agent’s utility for its preceding offer. Positive changes imply that the opponent has made concessions; hence, the agent should also make concessions.

In (6), $U(O_a^{t-1})$ again represents the agent’s utility for its prior offer. Positive changes indicate that the opponent has conceded, prompting the agent to make concessions. Considering the opponent’s previous n bids, where W_i represents the weights of each utility difference, the behavior-based approach determines overall utility changes, as demonstrated in (7). Equation (8) reveals that the value of the coefficient μ is determined by the current time and P_3 , which controls the degree of mimicry. Initially, the agent decreases or increases the target utility less than its opponent; subsequently, the degree of mimicry rises over time. Therefore, the “HybridAgent” strategy can smoothly conform with domains of varying sizes and harmonize with distinctive opponents utilizing behavior-based components of the HybridAgent. As an extension of the bidding strategy, CBOM agent also generates a target utility value by combining different p-values for various domain sizes. It cares about the social welfare score for both parties, choosing the most agreeable offer

Algorithm 2 Conflict-based negotiation strategy.

```

O_space: Offer space, n: Minimum number of round required for CBOM;
O_opp, O_cbom: Opponent’s & CBOM agent’s offer history, respectively ;
U_cbom(o): The utility of an offer o for the CBOM agent ;
U_opp: The estimated utility space of the opponent ;
ε: Parameter controlling the bid utility window ;
TU_Hybrid: Calculated target utility with Hybrid Bidding Strategy ;
1 O_potential ← {} ;
2 while ||O_potential|| > 0 do
3   for each o ∈ O_space do
4     if (U_cbom(o) ∈ [TU_Hybrid - ε, TU_Hybrid + ε]) & (o ∉ O_cbom) then
5       O_potential ← O_potential + o ;
6     end
7   end
8   ε ← ε + 0.01 ;
9 end
10 if ||O_opp|| < n then
11   O_cbom^t ← arg max_o U_cbom(O_potential) ;
12 else
13   U_opp ← CBOM(O_opp);
14   O_cbom^t ← arg max_o U_cbom(O_potential) * U_opp(O_potential) ;
15 end
16 if (min(U(O_cbom ∪ O_cbom^t)) ≤ U(O_opp^t)) then
17   Accept O_opp^t ;
18 else
19   Return O_cbom^t
20 end
    
```

from the list of offers that the opponent model creates. Thus, it is expected that CBOM agent can achieve higher utility while maximizing social welfare and finding quicker mutual agreement.

In accordance, Fig. 3 illustrates an example of the selected offer of the CBOM Agent according to its boundaries. The figure is structured with the y-axis representing the agent's utility and the x-axis representing the opponent's utility for the potential outcomes in the domain. Each outcome is depicted by blue dots on the graph. The red dot represents the target utility offer at a given time, indicating the preferred outcome for the agent. The red circle is drawn by adding the target utility to an epsilon value. Within this boundary, the agent selects the offer closest to the Nash offer (d_N), depicted in green. When the domain size is limited, only a few bids may remain within a specific utility range. Without enlarging the epsilon, the agent might end up repeating certain offers. The number of offers within the offer window should be increased in such situations. Expanding this boundary allows the agent to explore more offers, which helps avoid sending repetitive final offers to the opponent while still adhering to the target window. Examining other offers within the same window allows the agent to identify a more appropriate choice while upholding its target utility.

5 Experimental analysis

We first examine the performance of the proposed Conflict-based Opponent Model (CBOM) by conducting two different human experiments (Section 5.1) and extend this evaluation by considering the performance of the proposed strategy using this opponent model through agent-based negotiation simulations (Section 5.2).

5.1 Evaluation of opponent modeling via human-agent experiments

To show how well the proposed opponent modeling approach predicts the human opponent's preferences, we conducted experiments where participants negotiated with our agent on a given scenario to find a consensus within limited rounds by following the Alternating Offers Protocol (Section 3).

We consider the performance metrics to assess the quality of the predictions: Spearman's correlation and root-mean-square error (RMSE). The former metric indicates the accuracy of the predicted order of the outcomes according to the learned utility function, whereas the latter measures how accurate estimated utilities are. For correlation estimation, possible outcomes are sorted concerning the learned oppo-

nent model, and this ranking is compared with the actual ordering. Consequently, the Spearman correlation is calculated between the actual outcome ranking and the estimated one. The correlation would be high when both orderings are similar to each other. The correlation coefficient r ranges between -1 and 1, where the sign of the coefficient shows the direction, and the magnitude is the strength of the relationship. For RMSE, the utility of each outcome is estimated according to the learned model, and the error in the prediction is calculated (See (9)). When the estimated utility values are close to the actual utility values, the RSME values would be low. In summary, low RSME and high correlation values are desired in our case.

$$RMSE = \sqrt{\frac{1}{n} \sum_{i=1}^n (U(o_i) - \hat{U}(o_i))^2} \quad (9)$$

Baarslag et al. compare the performance of the existing opponent models in automated negotiation [5]. Their results show that frequentist-based opponent modeling approaches are the most effective among the existing ones despite the approach's simplicity. Therefore, we use a benchmark involving two different state-of-the-art frequentist opponent modeling approaches widely used in automated negotiation employed in HardHeaded [19] and Scientist [31] agents to evaluate the performance of the proposed opponent model. Frequentist opponent modeling techniques mostly rely on heuristics, assuming that the opponent would concede less on the essential issues and the preferred values appear more often than less preferred ones. Consequently, they check the frequency of each issue value's appearance in the offers. Furthermore, they compare the content of the consecutive offers and find out the issues with changed values. In other words, if the value of an issue is changed in the opponent's consecutive offers, the weights of those issues are decreased by a certain amount (i.e., becoming less critical). In the Scientist Agent, Tunali et al. aims to resolve some update problems and enhance the model by comparing a group of offer exchanges instead of only consecutive pairs of offers and adopting a decayed weight update mechanism. Each opponent model is fed and updated in each round by simulating the negotiation data obtained from human-agent negotiation experiments. At the end of each negotiation, the estimated models are evaluated according to the RMSE and Spearman correlation metrics explained above.

5.1.1 Study 1: human-agent negotiation in deserted Island scenario

We analyzed and utilized the negotiation log data collected during the human-agent negotiation experiments in

Fig. 3 Offer selection example of the CBOM Agent according to ϵ boundary

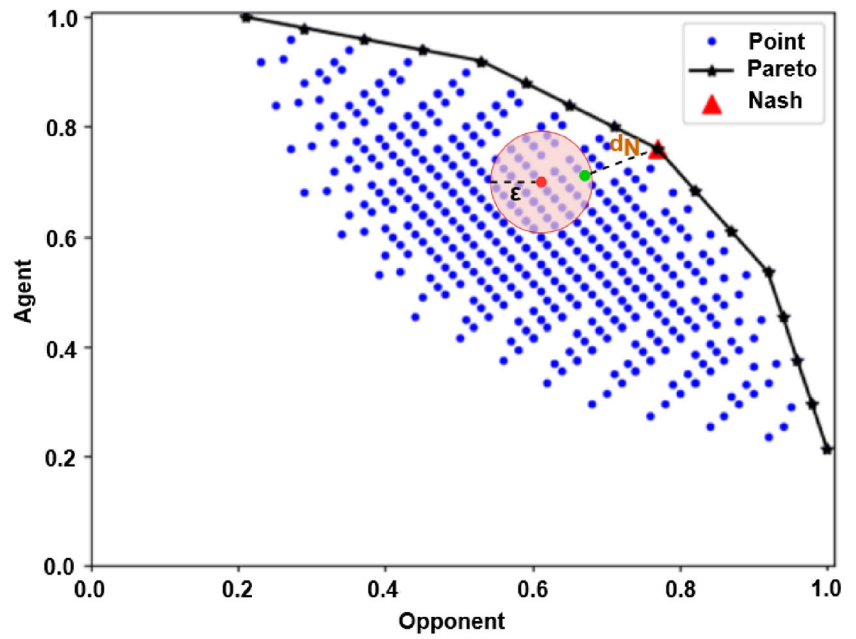


Table 1 Agent's and participants' preference profiles in deserted Island scenario

Items	First Negotiation		Second Negotiation	
	Agent	Participant	Agent	Participant
Compass	13	5	6	13
Container	22	20	13	5
Food	17	7	20	22
Hammer	6	13	5	10
Knife	5	10	10	17
Match	20	22	7	6
Medicine	7	6	17	7
Rope	10	17	22	20

Fig. 4 RMSE & Spearman Correlations for the experiment in Island Scenario (★ represents $p < 0.001$)

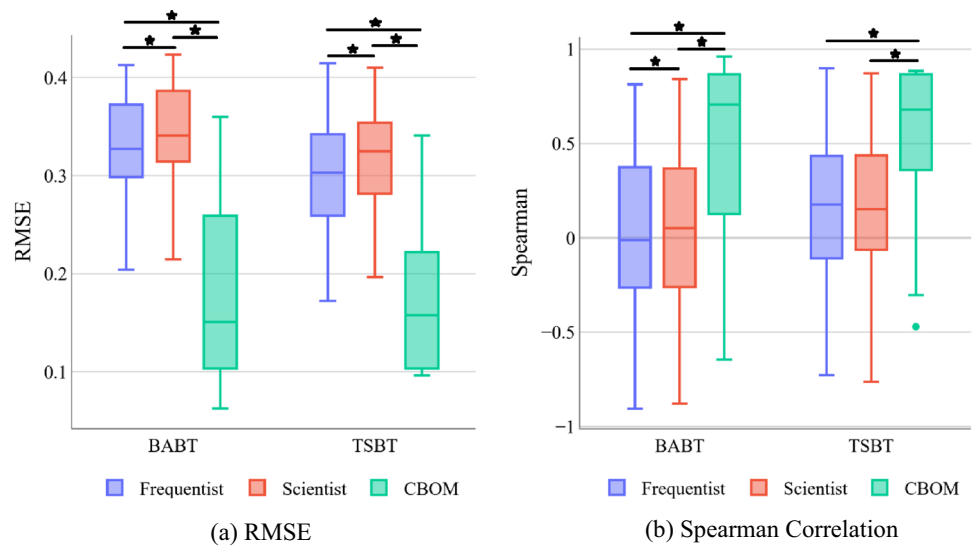


Table 2 Preference profiles for grocery negotiation sessions

Items	First Negotiation		Second Negotiation	
	Agent	Participant	Agent	Participant
Watermelon	4	12	12	4
Banana	1	8	8	1
Orange	12	4	4	12
Apple	8	1	1	8

[4], where the participants negotiated on a particular scenario called “Deserted Island”. They negotiated resource allocation based on the division of eight survival products by two partners who fell on the deserted island. Each participant attended two negotiation sessions where the utility distributions of the issues were the same, but the orderings differed. Table 1 shows the preference profiles for both sessions. During the experiments, participants only know their preferences, and so does the agent. In this study, 42 participants (21 men, 21 women, median age: 23) were included and asked to negotiate with our agent on a face-to-face basis, and the agent made counteroffers. Offer exchanges in both sessions were recorded separately for each session. At the end of this data collection process, 46 sessions using the time-based stochastic bidding tactic (TSBT) and 38 sessions using the behavior-based adaptive bidding tactic (BABT) were obtained. The average negotiation rounds to reach an agreement was 14.84, with a standard deviation of 5.2.

Figure 4 shows box plots for each opponent modeling technique’s RMSE and Spearman correlation values. As far as the correlation values are concerned, it can be said that CBOM’s ranking predictions are better than Scientist and Frequentist (See Fig. 4b). To apply the appropriate statistical significance test, we first check the normality of the data distribution via the Kolmogorov-Smirnov normality test and then the homogeneity of variance via Levene’s Test. We

applied the dependent sample t-test or the Wilcoxon Signed Rank test, depending on the results. If the data distribution passes these tests, the paired t-test is applied; otherwise, a non-parametric statistics test, namely the Wilcoxon-Signed Rank test. All statistical test results are given at the 99% confidence interval (i.e., $\alpha = 0.01$). When we apply the statistical tests, it is seen that CBOM’s ranking performance is statistically significantly better than others ($p < 0.01$). Similarly, the errors on the estimated utilities via CBOM are lower than the errors via other approaches (see Fig. 4a). Furthermore, it is seen that Scientist statistically significantly performed better than Frequentist for both metrics except when the agent employs the TSBT strategy.

5.1.2 Study 2: human-agent negotiation in grocery scenario

In this part, we analyzed and utilized the negotiation log data collected during another human-agent negotiation experiment in [16] where the participants negotiated on a particular scenario called “Grocery”. Different from the first study, the negotiation domain does not consist of binary resource items (i.e., *allocate* or *not allocate*). Instead, the negotiation parties negotiate on the number of items to be allocated (i.e., how many items will be allocated). In this scenario, there are four types of fruits, where each participant can have up to four

Fig. 5 RMSE & Spearman Correlations for the experiment in Grocery Scenario (★ represents $p < 0.001$)

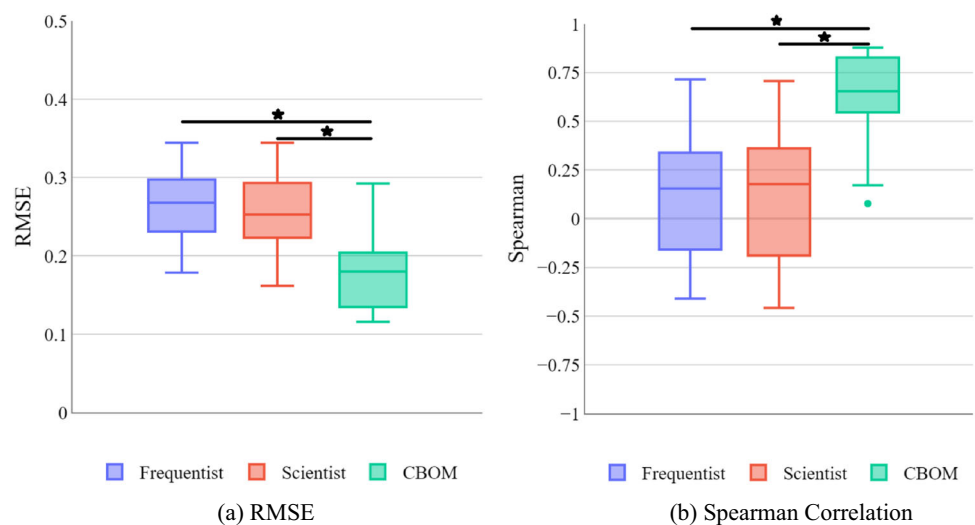


Table 3 Negotiation scenarios in automated negotiation settings

Domain (ANAC Year)	# of Values for Each I	Total Bids	Opposition
Car (2015)	3, 4, 4, 5	240	0.209
Smart Energy Grid (2016)	5, 5, 5, 5	625	0.362
Grocery (2011)	4, 4, 4, 5, 5	1,600	0.354
Party (2012)	3, 4, 4, 4, 4, 4	3,072	0.191
Politics (2015)	2, 3, 3, 4, 4, 4, 4, 5	23,040	0.221
Supermarket (2012)	4, 6, 6, 7, 7, 7	49,392	0.155

of each, and the opponent gets the rest. The participants aim to find an adequate division of the fruits. Table 2 shows the agent and participant's preference profiles for both sessions. It is worth noting that each party only knows its scores. In this experiment, the participants negotiated against an agent employing the hybrid strategy where TSBT and BABT strategies are used together for bidding. 28 participants attended two negotiation sessions where all negotiation sessions ended with an agreement, thus, totaling up to 56 negotiation sessions against the agent. The average negotiation rounds to reach an agreement was 19.39, with a standard deviation of 11.82.

Similar to the previous study, we update each opponent modeling by using the offer exchanges by the human participants and calculate the error and correlation values with the final model at the end of each negotiation. Figure 5 shows box plots for RMSE and Spearman correlation values per each opponent modeling technique in this scenario. We can conclude that *CBOM* statistically significantly outperformed others, whereas *Frequentist* performs better than *Scientist* when we analyze the statistical test results. Those results are in line with the first study and strongly show the success of the proposed opponent modeling in human-agent negotiations. It is worth noting that the prediction error in the grocery scenario is lower than in the island scenario. Although the number of possible outcomes in these scenarios is the same (256), the number of issues in grocery scenarios is lower than in the island scenario. Therefore, one can intuitively think it is easier to predict the evaluation values in the grocery scenario compared to island scenarios. In addition, this study's average number of rounds is higher (19.39 versus 14.84). When

we receive more offers, the model's accuracy may increase depending on the model.

5.2 Evaluation of the CBOM agent via automated negotiation experiments

In this section, we evaluate the performance of our agent employing the proposed *CBOM* opponent modeling by comparing its performance with that of the state-of-the-art negotiating agents available in automated negotiation literature. We built a rich benchmark of 15 successful negotiating agents who competed in the International Automated Negotiating Agents Competition ANAC [15] between 2011 and 2017. We ran negotiation tournaments in Genius, where each agent bilaterally negotiated on various negotiation scenarios. Six negotiation scenarios were used during the tournament, and the details of those scenarios are given in Table 3. As can be seen, the size and opposition degree of preference profiles in the given scenario is different. The size of the scenarios determines the search space. The larger the search space is, the more difficult it might be to estimate an accurate model based on the opponent's offers exchanges. Next, the opposition is valuable information regarding understanding the domain's capacity to satisfy both parties [5]. That is, it indicates how difficult it is to find a consensus. Taking the opposition of the preference profiles into account while analyzing the negotiation results may help us get an insight into how well the proposed negotiation strategy is in terms of social welfare with varying difficulties in finding an agreement.

Table 4 Average Spearman and RMSE results for six domains

Domains	SPEARMAN			RMSE		
	CBOM	Scientist	Hardheaded	CBOM	Scientist	Hardheaded
Car	0.73 ± 0.1	0.29 ± 0.0	0.33 ± 0.0	0.15 ± 0.0	0.31 ± 0.0	0.27 ± 0.0
Energy Grid	0.68 ± 0.1	0.25 ± 0.0	0.25 ± 0.0	0.16 ± 0.0	0.22 ± 0.0	0.30 ± 0.0
Grocery	0.81 ± 0.0	0.85 ± 0.0	0.82 ± 0.0	0.12 ± 0.0	0.17 ± 0.0	0.20 ± 0.0
Party	0.90 ± 0.0	0.82 ± 0.0	0.50 ± 0.0	0.10 ± 0.0	0.13 ± 0.1	0.26 ± 0.0
Politics	0.82 ± 0.0	0.86 ± 0.0	0.79 ± 0.0	0.20 ± 0.0	0.22 ± 0.0	0.16 ± 0.0
Supermarket	0.66 ± 0.0	0.58 ± 0.0	0.56 ± 0.0	0.21 ± 0.0	0.18 ± 0.0	0.30 ± 0.0

Table 5 Average agent individual utility

Agent Name	Car	Energy Grid	Grocery	Party	Politics	Supermarket	Average \pm Std
CBOMAgent	0.833	0.777	0.846	0.838	0.594	0.868	0.793 \pm 0.09
Atlas3	0.875	0.758	0.828	0.809	0.578	0.848	0.783 \pm 0.10
AgentKN	0.832	0.728	0.793	0.820	0.609	0.756	0.756 \pm 0.07
NiceTitForTat	0.804	0.680	0.767	0.848	0.569	0.853	0.753 \pm 0.10
HardHeaded	0.882	0.670	0.795	0.832	0.533	0.741	0.742 \pm 0.12
ParsCat	0.800	0.726	0.792	0.833	0.501	0.767	0.736 \pm 0.11
CUHKAgent	0.750	0.684	0.763	0.800	0.543	0.765	0.717 \pm 0.09
IAMcrazyHaggler	0.836	0.546	0.819	0.771	0.416	0.872	0.710 \pm 0.17
IAMhaggler2012	0.782	0.496	0.796	0.841	0.491	0.792	0.700 \pm 0.15
PonPokoAgent	0.773	0.606	0.785	0.783	0.427	0.802	0.696 \pm 0.14
Caduceus	0.731	0.546	0.801	0.703	0.489	0.711	0.663 \pm 0.11
ParsAgent2	0.779	0.540	0.713	0.698	0.424	0.805	0.660 \pm 0.14
Boulware	0.712	0.611	0.671	0.752	0.460	0.642	0.641 \pm 0.09
AgentX	0.672	0.582	0.655	0.801	0.347	0.642	0.616 \pm 0.14
YXAgent	0.770	0.553	0.721	0.625	0.397	0.553	0.603 \pm 0.12
Conceder	0.598	0.505	0.391	0.649	0.306	0.335	0.464 \pm 0.13

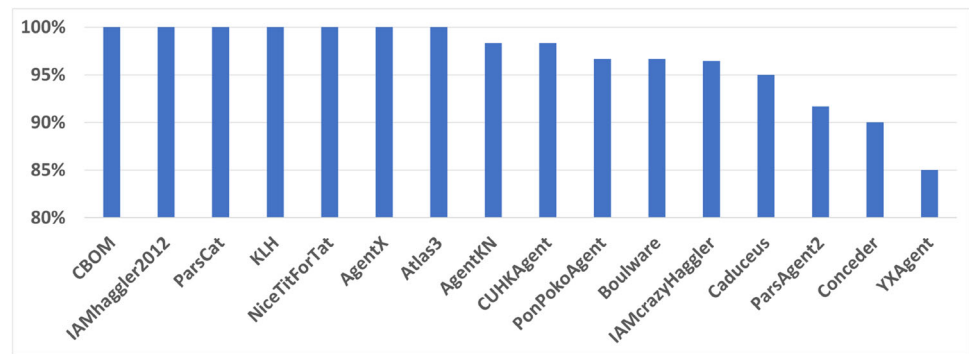
We formed a pool of agents involving our Conflict-based agent and the ANAC finalists in different categories. We ran a tournament in Genius where each agent bilaterally negotiated with each other on scenarios described in Table 3. The ANAC agents used in this evaluation are listed as follows:

– **Boulware** and **Conceder** are baseline agents available in Genius framework.

- **Hardheaded** [19] was the winner of individual utility category in ANAC 2011.
- **NiceTitForTat** [7] was the finalist of individual utility category in ANAC 2011.
- **CUHKAgent** [11] was the winner of individual utility category in ANAC 2012.
- **IAMHaggler2012** [15] was the winner of the Nash category in ANAC 2012.

Table 6 Average Nash Distances

Agent name	Car	Energy Grid	Grocery	Party	Politics	Supermarket	Average \pm Std
CBOMAgent	0.077	0.104	0.068	0.120	0.311	0.058	0.123 \pm 0.09
Atlas3	0.183	0.216	0.105	0.230	0.253	0.198	0.197 \pm 0.05
AgentKN	0.139	0.194	0.166	0.213	0.250	0.229	0.198 \pm 0.04
ParsCat	0.133	0.173	0.130	0.156	0.438	0.178	0.201 \pm 0.11
NiceTitForTat	0.154	0.209	0.178	0.230	0.187	0.250	0.201 \pm 0.03
IAMhaggler2012	0.115	0.500	0.095	0.152	0.363	0.145	0.228 \pm 0.15
HardHeaded	0.142	0.357	0.183	0.198	0.487	0.279	0.274 \pm 0.12
AgentX	0.210	0.273	0.230	0.267	0.418	0.262	0.277 \pm 0.07
Boulware	0.217	0.271	0.254	0.278	0.326	0.334	0.280 \pm 0.04
IAMcrazyHaggler	0.156	0.495	0.176	0.156	0.686	0.064	0.289 \pm 0.22
Caduceus	0.190	0.462	0.178	0.238	0.431	0.254	0.292 \pm 0.11
PonPokoAgent	0.201	0.392	0.177	0.205	0.640	0.154	0.295 \pm 0.17
CUHKAgent	0.361	0.428	0.280	0.205	0.457	0.261	0.332 \pm 0.09
ParsAgent2	0.197	0.474	0.274	0.306	0.576	0.210	0.340 \pm 0.14
YXAgent	0.268	0.526	0.350	0.376	0.718	0.450	0.448 \pm 0.15
Conceder	0.321	0.360	0.534	0.459	0.458	0.652	0.464 \pm 0.11

Fig. 6 Agreement rates of agents

- **Atlas3** [24] was the winner of individual utility category in ANAC2015, .
- **ParsAgent2** [17] was the winner of the Nash category in ANAC 2015.
- **AgentX** [9] was fourth of the Nash category in ANAC 2015.
- **Caudeceus** [10] was the winner of individual utility category in ANAC 2016.
- **YXAgent** [3] was the second of individual utility category in ANAC 2016.
- **PonPoko Agent** [3] was winner of individual utility category in ANAC 2017.
- **AgentKN** [3] was the second of the Nash category in ANAC 2017.

In order to study how well our opponent model performs when it negotiates with automated negotiating agents, we compare the performance of opponent models used in

Conflict-based (CBOM), Scientist, and HardHeaded by integrating those opponent models into our negotiation strategy. We calculated the Spearman correlation between the actual and estimated ranks of the outcomes per each scenario and reported their averages. Note that the higher correlation is, the better the prediction is. Table 4 shows those Spearman correlations and RMSE in the utility calculations where the best scores are boldfaced. It is seen that CBOM is more successful than others in terms of Spearman correlation, except for the results obtained in the grocery and politics domains. Furthermore, RMSE results show that the CBOM is more successful in all domains.

Next, we analyze the performance of the proposed negotiation strategy relying on the CBOM opponent modeling against the ANAC finalists. The most widely used performance metric in negotiation is the final received utility, which is intuitive and in line with Kiruthika’s approach to Multi-Agent Negotiation systems [18]. There are other metrics,

Table 7 Average agreement rounds

Agent name	Car	Energy Grid	Grocery	Party	Politics	Supermarket	Average \pm Std
AgentX	205	816	440	477	1,868	162	661 \pm 580
Conceder	1,222	1,826	2,626	2,017	2,559	3,095	2,224 \pm 611
IAMhaggler2012	2,765	4,393	2,096	2,829	4,146	2,378	3,101 \pm 864
ParsCat	2,899	3,988	2,741	3,034	4,254	2,567	3,247 \pm 639
CBOMAgent	3,506	3,859	2,978	3,452	4,407	2,403	3,434 \pm 633
IAMcrazyHaggler	3,526	4,202	2,984	3,104	4,489	2,553	3,476 \pm 682
PonPokoAgent	3,673	4,321	3,330	3,564	4,489	3,202	3,763 \pm 481
Boulware	3,604	4,005	3,650	3,705	4,317	3,560	3,807 \pm 270
NiceTitForTat	3,367	4,002	4,061	3,654	4,142	4,053	3,880 \pm 277
AgentKN	3,612	4,054	3,809	3,898	4,362	3,712	3,908 \pm 246
Atlas3	4,252	4,473	2,298	4,355	4,369	3,973	3,953 \pm 756
YXAgent	3,907	4,171	3,597	4,053	4,565	3,754	4,008 \pm 312
ParsAgent2	3,658	4,249	3,870	4,024	4,527	3,736	4,011 \pm 301
Caduceus	4,070	4,501	3,628	4,148	4,473	4,012	4,139 \pm 296
HardHeaded	4,099	4,399	3,968	4,196	4,646	4,243	4,258 \pm 217
CUHKAgent	4,301	4,458	4,512	4,272	4,659	4,236	4,406 \pm 150

such as nearness to Pareto optimal solutions/Kalai point/the Nash point, the sum of both agents' agreement utility (i.e., social welfare), and the product of those agreement utilities. Accordingly, we evaluate the performances regarding average individual received utility, Nash distance, and social welfare.

First, we analyze the average individual utilities received by each agent. Table 5 shows those utilities per each agent in each negotiation scenario where the highest scores are boldfaced. The last column shows the average scores of each agent in all domains. Our agent took in the first top three agents. We noticed that the worst performance of our agent was in the supermarket domain, where the outcome space is too large to search. Moreover, our agent performed well in the smart energy grid and grocery scenarios, whose opposition levels are high.

Table 6 shows the average Nash distance for each agent in all scenarios separately, and the final column indicates the average of all scenarios. Here, the lower the Nash distance is, the fairer the agent's outcomes are. Our conflict-based agent outperformed the ANAC finalist agents except for Nice TitForTat, which is known for maximizing social welfare in the Politics scenario (See Table 8 in Appendix). Similar results were obtained when we analyzed the social welfare in terms of summation of both agents' agreement utilities (See Appendix). Overall results support the success of our agent, and the reason may stem from the fact that our agent aims to learn its opponent's preferences over time and aims to find win-win solutions for both sides.

When we investigate the overall agreement rate, it can be seen that most of the agents have a high acceptance rate, and the leading ones, like ours and Atlas3, found agreements in all negotiations, as seen in Fig. 6. The final metric that we investigated is the average rounds to reach an agreement. In our experiments, the deadline is set to 5000 rounds per negotiation scenario. Table 7 shows the average rounds that the agent reached their agreement. It can be observed that the size of the outcome space and the opposition level may influence the agreement round. In large and competitive scenarios, agents needed more rounds to reach an agreement. Among all agents, Agent X tended to reach a consensus sooner than all other agents. Furthermore, IAMHaggler2012 and ParsCat agents tend to explore the offered space as much as they can in the given time. Therefore, these are the agents least affected by the size of the outcome space and its competitiveness. Our conflict-based agent could reach an agreement sooner than more than half of the agents but it is worth noting that it took more time in terms of seconds due to its computational complexity similar to AgentKN (See Table 9 in Appendix).

As a result of all the automated negotiations, we determined the six most successful agents in both the individual and fairness category. Figure 7 shows clearly that our agent gains the highest individual gain while having a fairer win-win solution (i.e., minimum distance to Nash solution). It is worth noting that while having high utility, our agent lets its opponent gains relatively high utility in contrast to other top agents.

6 Conclusion and future work

In conclusion, this work presents a conflict-based opponent modeling approach and a bidding strategy employing this model for bilateral negotiations. Apart from evaluating the performance of the proposed opponent model in two different human-negotiation experiment settings, the proposed strategy was also tested against the finalist of the ANAC agents considering various performance metrics such as individual utility and distance to the Nash solution. Our results show that the proposed approach outperformed the state of the negotiating agents, and the proposed opponent model performed better than other frequency-based models. The contribution of this study is twofold: (1) introducing a novel opponent modeling approach to learn human negotiators' preferences from limited bid exchanges and (2) presenting a suitable bidding strategy relying on the proposed opponent model for both collaborative and competitive negotiation settings.

Due to the algorithm's complexity, the agent's performance decreases when the outcome space becomes more extensive or the number of generated offers made by the opponent increases in automated negotiation. We are planning to reduce the computational complexity of the opponent modeling by adopting dynamic programming properties and local search. The upcoming study will focus on opponent model strategies that decrease the human-agent negotiation duration with the optimal number of rounds. It would be interesting to create stereotype profiles by mining the previous negotiation history and matching the current opponent's profile based on their recent offer exchanges.

Understanding and discovering the opponent's preferences over negotiation may play a key role in adopting strategic bidding strategies to find mutually beneficial agreements. However, as stated before, it is challenging to create a mental model for the opponent's preferences based on a few bid exchanges. In contrast to automated negotiation, the number of exchanged bids is limited in human-agent negotiation. That requires a bidding strategy smartly exploring the potential bids and building upon an opponent model, capturing the

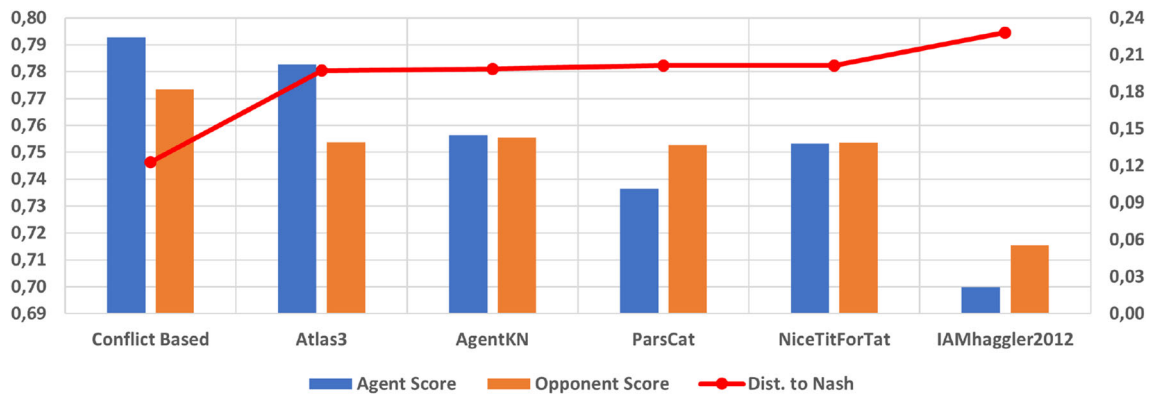


Fig. 7 Best six agents in all automated negotiation results

critical components of the opponent's preferences. While creating such modeling is not trivial with limited bid exchanges, the agent can exploit its previous negotiation experiences and take advantage of repeated patterns. As future work, it would be interesting to create different mental models from previous negotiation experiences by applying our model and trying to detect which mental model fits better for the current human negotiators. Consequently, instead of starting to learn from

scratch, our model can enhance the chosen model by analyzing the current bid exchanges. Furthermore, the agent can exploit different types of inputs, such as the opponent's arguments and facial expressions, to enhance opponent modeling.

Appendix

Table 8 Average social welfare score

Agent Name	Car	Energy Grid	Grocery	Party	Politics	Supermarket	Average \pm Std
CBOMAgent	165,655	157,227	150,312	167,216	106,945	172,876	153,372 \pm 21,995
AgentX	162,310	149,437	157,913	157,453	124,845	159,043	151,834 \pm 12,678
Atlas3	159,891	151,089	153,578	153,295	126,653	157,466	150,329 \pm 10,976
AgentKN	166,027	150,985	133,303	156,921	126,087	150,364	147,281 \pm 13,615
ParsCat	166,225	152,253	145,902	164,745	85,448	161,411	145,997 \pm 27,999
NiceTitForTat	165,437	144,536	129,573	152,589	127,955	145,374	144,244 \pm 12,914
IAMhaggler2012	161,990	93,551	156,376	165,018	93,827	164,709	139,245 \pm 32,337
HardHeaded	166,258	120,624	131,202	161,969	85,146	142,000	134,533 \pm 27,260
Boulware	152,574	134,493	118,818	141,937	116,171	127,207	131,867 \pm 12,750
IAMcrazyHaggler	155,615	95,329	136,363	157,082	58,579	172,290	129,210 \pm 39,828
PonPokoAgent	149,828	108,912	132,648	154,698	62,888	158,976	127,992 \pm 33,601
Caduceus	151,600	100,644	130,134	144,040	93,133	146,992	127,757 \pm 22,888
CUHKAgent	130,850	120,066	114,833	157,693	92,766	145,382	126,932 \pm 21,067
ParsAgent2	150,379	100,984	115,091	135,542	69,915	154,718	121,105 \pm 29,599
Conceder	136,634	128,370	81,617	112,968	114,208	68,503	107,050 \pm 24,320
YXAgent	139,677	93,292	111,212	121,324	54,423	107,581	104,585 \pm 26,484

Table 9 Average agreement run times

Agent Name	Car	Energy Grid	Grocery	Party	Politics	Supermarket	Average \pm Std
AgentX	9,831	18,626	3,413	7,771	81,721	23,662	24,171 \pm 26,610
Conceder	17,312	29,410	19,058	27,922	114,666	213,279	70,275 \pm 72,241
Boulware	40,855	50,588	32,838	43,204	142,281	189,926	83,282 \pm 60,379
ParsCat	36,405	66,436	30,499	41,597	172,065	184,158	88,527 \pm 64,418
PonPokoAgent	47,276	66,870	40,068	47,167	169,917	177,437	91,456 \pm 58,747
HardHeaded	46,853	57,890	34,569	53,148	165,997	198,421	92,813 \pm 64,299
Atlas3	52,283	72,105	18,138	66,174	161,367	199,162	94,872 \pm 63,694
CUHKAgent	50,303	55,542	38,336	53,566	168,014	214,496	96,709 \pm 68,404
YXAgent	55,307	75,352	46,004	59,929	194,091	162,102	98,797 \pm 57,485
IAMcrazyHaggler	57,375	80,281	42,139	44,002	199,154	175,726	99,780 \pm 63,581
IAMhaggler2012	78,696	124,397	63,920	88,375	226,875	231,055	135,553 \pm 68,526
ParsAgent2	154,667	164,868	46,388	117,370	227,899	254,998	161,032 \pm 68,806
Caduceus	48,768	70,825	62,566	101,401	521,733	508,281	218,929 \pm 209,987
NiceTitForTat	48,707	90,214	118,381	163,795	494,087	748,441	277,271 \pm 256,374
CBOMAgent	149,535	248,868	116,199	183,300	626,341	937,576	376,970 \pm 302,796
AgentKN	240,108	322,627	104,224	149,689	776,897	863,015	409,427 \pm 299,331

Funding This work has been supported by a grant of The Scientific and Research Council of Turkey (TÜBİTAK) with grant number 118E197 and partially supported by the CHIST- ERA grant CHIST-ERA-19-XAI-005, the Scientific and Research Council of Turkey (TÜBİTAK, G.A. 120N680).

Declarations

Conflicts of interest The authors declare that they have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Amini M, Fathian M, Ghazanfari M (2020) A boa-based adaptive strategy with multi-party perspective for automated multilateral negotiations. *Appl Intell* 50. <https://doi.org/10.1007/s10489-020-01646-y>
- Aydoğan R, Festen D., Hindriks K, Jonker C (2017) Alternating Offers Protocols for Multilateral Negotiation, pp 153–167. *Studies Comput Intell Springer Sci.* https://doi.org/10.1007/978-3-319-51563-2_10
- Aydoğan R, Fujita K, Baarslag T, Jonker CM, Ito T (2021) ANAC 2017: Repeated multilateral negotiation league. *Advances in automated negotiations.* Springer Singapore, Singapore, pp 101–115
- Aydoğan R, Keskin MO, Çakan U (2022) Would you imagine yourself negotiating with a robot, jennifer? why not? *IEEE Trans Hum Mach Syst* 52(1):41–51. <https://doi.org/10.1109/THMS.2021.3121664>
- Baarslag T, Hendriks M, Hindriks K, Jonker C (2013) Predicting the performance of opponent models in automated negotiation. In: 2013 IEEE/WIC/ACM International joint conferences on Web Intelligence (WI) and Intelligent Agent Technologies (IAT), vol 2, pp 59–66. <https://doi.org/10.1109/WI-IAT.2013.91>
- Baarslag T, Hendriks MJC, Hindriks K, Jonker C (2016) A survey of opponent modeling techniques in automated negotiation. In: *AAMAS*
- Baarslag T, Hindriks K, Jonker C (2013) A tit for tat negotiation strategy for real-time bilateral negotiations. In: *Complex automated negotiations: theories, models, and software competitions.* Springer Berlin Heidelberg, Berlin, Heidelberg, pp 229–233. https://doi.org/10.1007/978-3-642-30737-9_18
- Fatima S, Kraus S, Wooldridge M (2014) *Principles of automated negotiation.* Cambridge University Press. <https://doi.org/10.1017/CBO9780511751691>
- Fujita K, Aydoğan R, Baarslag T, Hindriks K, Ito T, Jonker C (2017) The Sixth automated negotiating agents competition (ANAC 2015). Springer International Publishing, Cham, pp 139–151 (2017). https://doi.org/10.1007/978-3-319-51563-2_9
- Güneş TD, Arditi E, Aydoğan R (2017) Collective voice of experts in multilateral negotiation. *PRIMA 2017: Principles and practice of multi-agent systems.* Springer International Publishing, Cham, pp 450–458
- Hao J, Leung Hf (2014) CUHKAgent: An Adaptive negotiation strategy for bilateral negotiations over multiple items. Springer Japan, Tokyo, pp 171–179. https://doi.org/10.1007/978-4-431-54758-7_11
- Hindriks K, Jonker C, Tykhonov D (2009) The benefits of opponent models in negotiation. In: 2009 IEEE/WIC/ACM International joint conference on Web Intelligence and Intelligent Agent Tech-

- nology (WI-IAT). pp 439–444. <https://doi.org/10.1109/WI-IAT.2009.192>
13. Hosokawa Y, Fujita K (2020) Opponent's preference estimation considering their offer transition in multi-issue closed negotiations. *IEICE Trans Inf Syst E103.D*, 2531–2539 (2020). <https://doi.org/10.1587/transinf.2020SAP0001>
 14. Jonge Dd, Sierra C (2015) NB3: a multilateral negotiation algorithm for large, non-linear agreement spaces with limited time. *Auton Agents Multi-Agent Syst* 29(5):896–942. <https://doi.org/10.1007/s10458-014-9271-3>
 15. Jonker C, Aydoğan R, Baarslag T, Fujita K, Ito T, Hindriks K (2017) Automated negotiating agents competition (ANAC). *Proceedings of the AAAI Conference on Artificial Intelligence* 31(1). <https://ojs.aaai.org/index.php/AAAI/article/view/10637>
 16. Keskin MO, Çakan U, Aydoğan R (2021) Solver agent: Towards emotional and opponent-aware agent for human-robot negotiation. In: *Proceedings of the 20th international conference on Autonomous Agents and MultiAgent Systems, AAMAS 21*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, pp 1557–1559
 17. Khosravimehr Z, Nassiri-Mofakham F (2017) Pars agent: hybrid time-dependent, random and frequency-based bidding and acceptance strategies in multilateral negotiations. *Springer International Publishing, Cham*, pp 175–183. https://doi.org/10.1007/978-3-319-51563-2_12
 18. Kiruthika U, Somasundaram TS, Subramanian KS (2020) Lifecycle model of a negotiation agent: A survey of automated negotiation techniques. *Group Decision and Negotiation* 29:1–24. <https://doi.org/10.1007/s10726-020-09704-z>
 19. van Krimpen T, Looije D, Hajizadeh S (2013) *HardHeaded*. Springer Berlin Heidelberg, Berlin, Heidelberg, pp 223–227. https://doi.org/10.1007/978-3-642-30737-9_17
 20. Lewis M, Yarats D, Dauphin Y, Parikh D, Batra D (2017) Deal or no deal? end-to-end learning of negotiation dialogues. In: *Proceedings of the 2017 conference on empirical methods in natural language processing*. Assoc Comput Ling, Copenhagen, Denmark, pp 2443–2453. <https://doi.org/10.18653/v1/D17-1259>
 21. Lin R, Kraus S, Wilkenfeld J, Barry J (2008) Negotiating with bounded rational agents in environments with incomplete information using an automated agent. *Artif Intell* 172(6–7):823–851
 22. Mell J, Gratch J, Aydoğan R, Baarslag T, Jonker CM (2019) The likeability-success tradeoff: Results of the 2nd annual human-agent automated negotiating agents competition. In: *2019 8th international conference on affective computing and intelligent interaction (ACII)*, pp 1–7. <https://doi.org/10.1109/ACII.2019.8925437>
 23. Mell J, Lucas GM, Mozgai S, Gratch J (2020) The effects of experience on deception in human-agent negotiation. *J Artif Intell Res* 68:633–660
 24. Mori A, Ito T (2017) *Atlas3: A negotiating agent based on expecting lower limit of concession function*. Springer International Publishing, Cham, pp 169–173. https://doi.org/10.1007/978-3-319-51563-2_11
 25. Nazari Z, Lucas G, Gratch J (2015) Opponent modeling for virtual human negotiators. In: *International conference on intelligent virtual agents*. pp 39–49. https://doi.org/10.1007/978-3-319-21996-7_4
 26. Oshrat Y, Lin R, Kraus S (2009) Facing the challenge of human-agent negotiations via effective general opponent modeling. *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS 2*:377–384. <https://doi.org/10.1145/1558013.1558064>
 27. Peled N, Gal YK, Kraus S (2015) A study of computational and human strategies in revelation games. *Auto Agents Multi-Agent Syst* 29(1):73–97
 28. Rosenfeld A, Zuckerman I, Segal E, Drein O, Kraus S (2014) *Negochat: a chat-based negotiation agent*. In: *AAMAS*. pp 525–532
 29. Sanchez-Anguix V, Aydoğan R, Julian V, Jonker C (2014) Unanimously acceptable agreements for negotiation teams in unpredictable domains. *Elec Comm Res Appl* 13(4):243–265
 30. Sengupta A, Mohammad Y, Nakadai S (2021) An autonomous negotiating agent framework with reinforcement learning based strategies and adaptive strategy switching mechanism. In: *Proceedings of the 20th international conference on autonomous agents and multiAgent systems, AAMAS '21*. pp 1163–1172
 31. Tunali O, Aydoğan R, Sanchez-Anguix V: Rethinking frequency opponent modeling in automated negotiation. In: *International*. pp 263–279 (2017). https://doi.org/10.1007/978-3-319-69131-2_16
 32. Yu C, Ren F, Zhang M (2013) *An adaptive bilateral negotiation model based on bayesian learning*. Springer, vol 435, pp 75–93. https://doi.org/10.1007/978-3-642-30737-9_5

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Mehmet Onur Keskin obtained his B.Sc. in Computer Science and MSc. in Artificial Intelligence from Özyeğin University, Istanbul. As a research assistant at the OZU Interactive Intelligence Systems Laboratory since 2017, Keskin has been focused on advancing human-agent/robot interaction frameworks for complex decision-making problems. His academic journey was enriched with research internships at well-known institutions such as the University of Cambridge, Brown 2 University, TU Delft, and the Oregon State University Collaborative Robotics Institute. He won IBM's Human multi-agent immersive negotiation competition at IJCAI 2020 and achieved second place in two categories at the Supply Chain Management League at IJCAI 2021. His research effort was recognized with the Richard E. Merwin Student Scholarship by the IEEE Computer Society and the Best Research Assistant Award by Özyeğin University in 2022. Keskin was one of the main organizers of the International Automated Negotiating Agents Competition 2023 (ANAC'23) at AAMAS.



Berk Buzcu earned his Bachelor of Science in Computer Science from Özyeğin University. Subsequently, he pursued further academic education as a Master's Student at the same institution. During this time, he actively participated in the CHIST-ERA (locally Tübitak) funded project, "EXPECTATION: Personalized Explainable AI for decentralized agents with heterogeneous knowledge." His focus within the scope of the project was the application of explainability techniques

in recommender systems which aligned with his research interests. The main goal was to enhance the success of recommender systems in achieving user goals via sociability. In addition to his thesis research, Berk developed a human-robot negotiation environment while touching on other multiagent-related topics involving Blockchain and drone cooperation.



Reyhan Aydoğan is an assistant professor in Computer Science at Özyeğin University and a research fellow in the Interactive Intelligence Group at Delft University of Technology. She received her M.Sc. and Ph.D. degrees in Computer Engineering from Boğaziçi University. Her primary research focuses on modeling, developing, and analyzing agent technology that integrates different aspects of intelligence, such as reasoning, decision-making, and learning. She applies AI techniques such as

machine learning and semantic reasoning in designing and developing agent-based decision support systems, e.g., negotiation support systems and automated negotiation tools. She is particularly interested in supporting human decision-makers in complex and dynamic environments, which also requires the design of effective human-computer interaction (e.g., preference elicitation). She serves as a program committee member in conferences such as AAMAS, IJCAI, and ECAI and is a reviewer for journals such as JAAMAS and JAIR. Dr. Aydoğan is one of the main organizers of the International Automated Negotiating Agents Competition (ANAC).