

An Adaptive Emotion-Aware Strategy for Human-Agent Negotiation: Insights from Real-World Human-Robot Experiments

M. Onur Keskin
onur.keskin@ozu.edu.tr
Department of Computer Science,
Özyeğin University
Istanbul, Turkey

Umut Çakan
umut.cakan@ozu.edu.tr
Department of Computer Science,
Özyeğin University
Istanbul, Turkey

Reyhan Aydoğan
reyhan.aydogan@ozyegin.edu.tr
Artificial Intelligence & Data
Engineering, Özyeğin University
Istanbul, Turkey
Interactive Intelligence, Delft
University of Technology
Delft, The Netherlands

Abstract

Negotiation is pivotal for conflict resolution in human-agent interactions, where emotional and behavioral dynamics can significantly shape the outcomes. However, many existing strategies prioritize time- or behavior-based tactics and overlook the dynamic role of emotional awareness. This paper presents the Solver Agent, which integrates real-time facial expression recognition into a hybrid strategy incorporating time- and behavior-based approaches. It is deployed on a humanoid robot with multimodal interaction capabilities (speech, gestures, facial expression analysis) to dynamically refine its bidding and concession strategies based on an opponent's emotional cues and negotiation patterns. In user studies with 28 participants, the Solver Agent achieved higher agent scores, improved social welfare, and faster agreements than a baseline hybrid strategy without compromising participant satisfaction. Participants also viewed the Solver Agent as more attuned to their preferences and goals. These findings highlight that embodied emotion-aware negotiation can foster equitable and efficient collaboration, pointing to new opportunities in human-agent interaction research.

CCS Concepts

• **Computing methodologies** → *Intelligent agents*.

Keywords

Human-Agent Negotiation, Emotion, Negotiation Strategy, Human-Robot Interaction

ACM Reference Format:

M. Onur Keskin, Umut Çakan, and Reyhan Aydoğan. 2025. An Adaptive Emotion-Aware Strategy for Human-Agent Negotiation: Insights from Real-World Human-Robot Experiments. In *Proceedings of ACM International Conference on Intelligent Virtual Agents (IVA '25)*. ACM, New York, NY, USA, 9 pages. <https://doi.org/XXXXXXX.XXXXXXX>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

IVA '25, Berlin, Germany

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-XXXX-X/2018/06
<https://doi.org/XXXXXXX.XXXXXXX>

1 Introduction

Negotiation is a complex process in which various parties with different preferences seek to reach a consensus [13]. Various approaches [3, 9, 12, 40] have been proposed to automate this process. In automated negotiations, agents can exchange thousands of offers and reach an agreement in seconds. Designing an effective strategy in human-agent negotiations requires addressing time constraints and opponent model uncertainty and considering *human factors* such as bounded rationality, reciprocity, and emotional awareness [26]. This paper examines how *emotionally aware* negotiation strategies can improve outcomes in *human-robot* settings through user experiments. On average, the number of offers in human-agent negotiation does not exceed 20 [15, 26, 33], and most negotiations end in fewer than 20 rounds [30]. In addition to exchanging offers, human negotiators exchange arguments and *emotional signals* [14]. While designing a negotiating agent for such settings, it is crucial to identify the best ways to utilize *nonverbal expressions* and other social signals. Moreover, human negotiators expect *reciprocal behavior*: If they make a cooperative move, they want the opponent to cooperate similarly; otherwise, attitudes may shift drastically [26]. Hence, awareness of the other side's attitude is pivotal in human negotiations.

Emotions can play a crucial role in shaping cognitive appraisals and concession behavior. Several works investigate the effect of emotions in negotiation [35, 41, 46], such as the finding that people *tend to concede more* to an angry counterpart [10] or that dominant emotional expressions lead to higher scores [45]. Although many of these works concentrate on expressing emotions, fewer have rigorously examined how an agent perceives its human partner's emotional state and adapts its strategy accordingly. The IAGO framework [28] enables emoticons or textual emotion sharing in human-agent negotiation; however, the opponent's emotional state is not deeply integrated. Moreover, in the annual Human-Agent Negotiation Competition [30], emotional states remain largely unexplored. In contrast, our work proposes a novel strategy that continuously perceives the opponent's emotional state and adjusts to the opponent's *bid exchanges* and *remaining negotiation time* for a more *contextual* approach.

The physical embodiment also influences human interaction [7]. Negotiating with a humanoid robot can intensify emotional expressions compared to virtual agents, creating a richer context for

adaptive strategies. Previous human-agent negotiation works [6, 42] do not provide *fully autonomous* negotiation strategies as we do. Furthermore, the rise of socially interactive agents motivates a deeper investigation of how agents can interpret and respond to emotional signals, especially under real-time negotiation demands. Hence, we aim to design an *emotionally aware* negotiation agent that perceives the emotional state of a human partner and adapts its offers accordingly.

Our main contribution is an experimental evaluation of how integrating *emotional awareness* into a human-agent negotiation framework affects the negotiation outcomes. In user studies, we find that agents leveraging the opponent's emotional state achieve *significantly higher agent scores* and *better social welfare* in less time than agents dismissing such signals. We further examine participant attitudes and discuss how emotional adaptation shapes the perceived fairness or empathy of the agent. Specifically, this paper investigates:

- **RQ1:** Does the emotionally aware agent affect the negotiation outcomes (e.g., Individual Utilities, Social Welfare, and agreement time/round)?
- **RQ2:** Does the emotionally aware agent influence the participants' attitudes toward the agent itself?
- **RQ3:** Do participants' *pre-negotiation priorities* (e.g., self-interest vs. cooperative stances) affect outcomes if the agent does not leverage emotional inputs?

The rest of the paper is organized as follows. Section 2 presents related work. Section 3 details our human-agent negotiation framework and bidding strategies, while Section 4 describes the experimental design and results. Finally, Section 5 concludes the paper with future directions.

2 Related Work

Even though human-agent negotiation has gained attention, researchers generally adapt existing *automated negotiation* strategies without considering human factors. Some of these works are slight modifications of existing strategies (e.g., time-based or behavior-based) [17, 19]. In contrast, others introduce new aspects such as *arguments* or *emotional expression* [24, 29]. Although specific agents consider opponent behavior to some degree [17], they rarely account for *opponent awareness* of changing offers and do not *continuously measure the user's affect*.

Vahidov *et al.* introduce a variant of a time-based concession function for human-agent negotiation [44]. Aydoğan *et al.* propose a stochastic time-based concession strategy picking random offers within a Boulware-Conceder utility range [12, 44]. Jonker *et al.* present *Deniz Agent*, which adapts its moves (e.g., concession, selfish, silent) based on the opponent's actions [17], while the amount of concession is determined by an optimal bidding strategy [4]. Lin *et al.* introduce 'QOAgent' to negotiate with boundedly rational agents under incomplete information [27]. KBAgent extends QOAgent by exploiting history to avoid offers that previously annoyed humans [31]. These approaches highlight *time-based* and *behavior-based* concessions for human-agent negotiation but do not incorporate an opponent's *emotional signals* or their awareness of changing offers. Our work addresses this gap by adapting the agent's strategy based on the user's real-time *facial expressions*.

The ANAC organizers have encouraged research on human-agent negotiation through dedicated leagues [30]. Examining these league participants reveals diverse tactics: "LyingAgent" misleads opponents about preferences for higher gains, "Elphaba" seeks mutually beneficial offers, "Murphy" uses jokes to build rapport, and some agents (e.g., Agent Cena, Boulware) rely on utility thresholds. Pinocchio postpones revealing all issues to propose beneficial offers for both parties. However, none of these agents *continuously* evaluate the user's *emotional state* to adapt their bidding strategy. Moreover, some agents try to *express* manipulate human counterparts by provoking anger or friendliness [29, 30]. However, expressing emotion differs from *perceiving* and adapts to the user's affect in real time. Our approach leverages *facial-expression-based recognition* to fine-tune concessions, thus extending these emotional strategies from mere *expression* to *reciprocal adaptation*. A range of works confirm that human participants can concede more to an agent *expressing* anger [10, 46], that *dominant* movements can yield higher agent scores [45], or that warmth influences the willingness to renegotiate [36]. These studies focus primarily on how the agent's emotional *expression* affects a human's behavior. In contrast, our work *perceives* the user's emotional state and adapts accordingly.

Most existing human-agent negotiation frameworks rely on text-based interfaces or 2D avatars, such as IAGO [28] and NegoChat [38], or focus on speech-based virtual agents [11]. Lewis *et al.* [24] learn chat-based negotiations from transcripts of human-human talks. In contrast, a physically embodied *humanoid* robot can amplify social presence, leading to stronger emotional displays and more *immersive* interaction. Although some studies explored *human-robot* negotiation [6, 42], they mainly address interaction characteristics (e.g., handshake feedback, disagreement style) rather than implementing a *real-time emotion-aware* strategy. While robotic persuasion frameworks [32] examine how dominance-based concessions can shape outcomes, we focus on real-time emotional awareness and adaptation rather than explicitly persuasive moves. In our work, a humanoid robot autonomously negotiates through speech and gestures, utilizing the power of machine learning to adapt its offers based on *facial expression recognition*. Thus, we integrate *embodied negotiation* with *emotional perception*, examining how these elements jointly enhance negotiation outcomes. Overall, our approach fills a gap in the literature by moving beyond time- or behavior-based concessions to *emotion-driven adaptation*. This utilizes an *embodied* platform capable of *perceiving* and *responding* to user emotions in real time.

3 Human-Agent Negotiation Framework

In our study, we adopt a human-agent negotiation (HAN) framework [21] in which a *physically embodied* Nao robot interacts with a human negotiator through speech, camera, and microphone. Nao uses a *pre-trained convolutional neural network* [25] to detect the user's facial expressions in real-time and applies *text-to-speech* technology, as well as predefined gestures and verbal statements to communicate offers, basic emotional states (e.g., offended, pleasant) similar to the human-agent negotiation framework proposed in [2]. Additionally, our system utilizes Google's automated speech recognition and speech-to-text APIs. This multimodal approach

aims to replicate *human-like* negotiation signals more closely than text-only/speech-only systems [11, 18, 28, 38].

We adopt *Alternating Offers Protocol* [1] as illustrated in Figure 1. The *human* initiates each negotiation with an offer in our experimental setup. Nao *accepts* or *counteroffers* until a termination condition (deadline or agreement) is reached. Human participants specify desired resource allocations in natural language (e.g., ‘I want three apples’). Nao parses each sentence using speech-to-text technology and domain-specific grammar. Mispronunciations are reduced by matching user utterances against a negotiation corpus, ensuring robust extraction of structured offers. For reproducibility purposes, all source code and related contents are available in the GitHub repository ¹.




Figure 1: Negotiation Protocol

To facilitate clarity, we rely on *regular expressions* to interpret each user’s acceptance or rejection of the offer. For example, if the participant says “deal” or “agree”, Nao logs an accept action and ends the negotiation if it is consistent with the user’s final allocation.

3.1 Hybrid Agent: Time & Behavior Based Agent

Since agents must deal with limited time, the remaining time should be taken into consideration during negotiation. To balance *time pressure* and *opponent behavior*, we adopt a hybrid strategy [22], combining a time-based concession function [44] and a behavior-based approach inspired by [12, 37]. The principal intuition is that when the deadline is distant, the agent pays more attention to its opponent’s behavior while deciding the agent’s target utility for its next offer. As the deadline approaches, it tends to find an agreement urgently; therefore, it prioritizes the remaining time.

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

In equation 1, TU_{Hybrid} denotes the agent’s target utility of the hybrid function where $t \in [0, 1]$ is the scaled time. TU_{Time} is derived from Vahidov’s time-dependent function [44] (Eq. 2), while $TU_{Behavior}$ captures the windowed changes in the user’s offers (Eq. 3–5). This ensures that at the start of the negotiation, *opponent behavior* dominates, while the *time-based* concession becomes dominant as the deadline approaches. The coefficients P_0 , P_1 , and

P_2 respectively define the maximum, the curve of the concession strategy, and minimum agent utility for each stage of negotiation (e.g., 0.9, 0.7, 0.4).

$$TU_{Time} = (1 - t)^2 \cdot P_0 + 2(1 - t)t \cdot P_1 + t^2 \cdot P_2 \quad (2)$$

$$TU_{Behavior} = U(O_j^{t-1}) - \mu \times \Delta U \quad (3)$$

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

$$\mu = P_3 + t \times P_3 \quad (5)$$

For behavior-based updates, we adopt an extension of *Tit-For-Tat* [12], in which the agent tracks up to the last n offers of the opponent to capture short-term fluctuations without being confused by older and inconsistent offers. While avoiding missing the opponent’s general bidding pattern, our tactic estimates the utility changes of the opponent’s offers within this window by giving more priority (e.g., $W_1 > W_2$) to the changes on the most recent ones.

In this work, the agent considers the opponent’s last five ($n = 5$) offers and estimates the weighted utility difference, as human-agent negotiation sessions typically last 20 rounds on average. To mimic the opponent’s behavior, the agent scales the overall utility change by a time-dependent empathy parameter, μ , to estimate the agent’s target utility as seen in Equation 3. $U(O_j^{t-1})$ denotes the utility of the agent’s previous offer. ΔU measures how much the opponent’s utility has changed over its last n bids (we set $n=5$). If the opponent has *conceded* (i.e., increasing utility of the agent, positive ΔU), our agent reduces its utility proportionally; if the opponent raises demands (negative ΔU), our agent’s target utility also increases. In addition to that, ΔU is controlled by the *empathy* coefficient μ (Eq. 5), where P_3 is the initial empathy parameter. In our study, P_3 is set to 0.5. As the negotiation time progresses, μ grows so that *mimicking* the opponent’s moves becomes more impactful, yet near the end, the *time-based* part of Eq. 1 dominates, preventing ending negotiation without an agreement.

3.2 Solver Agent: Emotion-Aware Hybrid Agent

In this work, we introduce the *Solver Agent* that augments the hybrid approach by integrating two new parameters into the behavior-based negotiation strategy:

- P_A : An *awareness* coefficient captures how closely the opponent’s *behavior changes* with our agent’s bidding behavior.
- P_E : An *emotion* coefficient that captures the real-time facial expression feedback from the opponent.

Thus, Equation 6 refines $TU_{Behavior}$ to incorporate emotional cues and opponent awareness:

$$TU_{Behavior} = U(O_j^{t-1}) + (P_A^2 \times P_E) - [(1 - P_A^2) \times (\mu \times \Delta U)] \quad (6)$$

Here, P_A and P_E denote the awareness and the emotion coefficient, respectively. P_A^2 modulates how much emotional input (P_E) shifts agent’s target utility, while $(1 - P_A^2)$ balances the *mimicking* factor, $\mu \times \Delta U$. If the opponent’s behavior strongly tracks ours (meaning that P_A is high), we rely more on their *expressed effect*; otherwise, we pivot on observed concessions or demands.

3.2.1 Estimating the Emotion Coefficient (P_E): CNN-based model gives a single dominant emotion prediction for each frame [25].

¹<https://github.com/anonimipanda/human-agent-negotiation-framework>

Instead of selecting one *dominant* emotion, our approach aggregates a *vector* of certainty values (sad, happy, angry, neutral, surprised) using a CNN-based facial expression model. Our early analysis revealed that pilot sessions rarely recognized *disgust* and *fear* emotions, so we exclude these two categories to improve model robustness. Figure 3 shows that the agent collects instant images of human negotiators during the negotiation. The model outputs the certainty of the prediction for each emotion (e.g., sad: 0.7, happy: 0.1). Let m be the number of frames collected from the starting time of Nao's offer until the time of the opponent's response. Equation 7 shows how P_E is calculated where F_i^k denotes the certainty value of i^{th} emotion in k^{th} frame and V_i indicate the weights of the i^{th} emotion as seen in Figure 2. Here, V_i is associated with negative weights if the facial expression is labeled with a negative category, such as sadness and anger. Otherwise, they are associated with a positive value. That is, P_E is the weighted average of certain values of each emotion. Studies show that facial expressions (e.g., a smile) do not necessarily mean the person is happy, and there are different reasons for the same facial feedback [5, 39]. The motivation behind using the vector of emotions instead of a dominant emotion is that the dominant emotion might be misleading since it depends on the context.

$$P_E = \sum_{i=1}^5 \left(V_i \times \left(\sum_{k=1}^m F_i^k \right) / m \right) \quad (7)$$




Figure 2: Weights of Categorical Facial Expressions




Figure 3: Example Facial Expression Feedback Vector

3.2.2 Opponent-Awareness Coefficient (P_A): In addition to the emotion vector, the agent should consider to what extent the opponent's facial feedback aligns with the agent's offer pattern changes, since the opponent may try to deceive by showing negative facial expressions while being pleased with the agent's offer. To estimate the opponent awareness coefficient P_A – the degree of the opponent's response to the agent's behavior changes, both agents' subsequent moves [16] (e.g., silent, nice, concession, unfortunate, fortunate, selfish) are analyzed.

$$P_A = C_H / C_A \quad (8)$$

First, the agent calculates the number of times the opponent changes its behavior from one type to another when the agent changes its behavior type (C_H). It corresponds to the degree of the opponent's response to the agent's behavior changes. This number is normalized (i.e., divided by the total number of the agent's behavior changes, C_A). This percentage equals P_A and allows the agent to understand the correlation between emotional changes received with the camera and the opponent's offer. In this way, the emotion-aware bidding strategy can deal with human manipulation and camera errors.

The accurate calculation of P_A depends on how well the agent estimates its opponent's utility function. Conflict-based opponent modeling [20] is utilized for this work. According to our experimental results, this opponent modeling approach outperforms the well-known frequentist opponent modeling approaches in automated negotiation [43, 47]. The calculated RMSE and Spearman correlation values are given in Table 1. The conflict-based opponent model estimated the opponent's utility function more accurately. Therefore, we adopted this opponent model for our study.

Table 1: Accuracy Comparison of the Opponent Models

Opponent Models	RMSE	SPEARMAN
Conflict-Based [20]	0.179 ± 0.05	0.820 ± 0.10
Scientist [43]	0.258 ± 0.04	0.560 ± 0.16
Frequency [47]	0.267 ± 0.04	0.562 ± 0.15

Adjusting Concession Parameters via Clustering

Finally, the Solver Agent adjusts related concession parameters in Equation 1 after a certain number of rounds – the average number of rounds to complete human-agent negotiation. The agent's move plays a crucial role in the received utility at the end of the negotiation when it approaches the deadline. Therefore, the agent acts more carefully and adjusts its concession parameters strategically. To achieve this, we classify the human negotiators' behavior in terms of the percentage of each move type (%concession, %fortune, %nice, %selfish, %unfortunate, %silent) they made in another human-agent negotiation dataset comprising 116 negotiations [2]. A clustering algorithm, K-means, categorizes the human players according to their move percentages with elbow analysis. This categorization is named S_T according to opponents' dominant negotiation moves defined in [16]. By analyzing centroids and deviations of the clusters, we found out that there are five dominant categories: *fortunate*, *neutral* (i.e., no dominant moves), *silent*, *selfish*, and *concession* based on their dominant moves. The Solver Agent calculates the move types of its opponent after reaching a certain number of rounds n by checking which category the opponent fits and accordingly updates the parameters of its strategy as specified in Table 2. Note that the initial parameters are set according to the neutral category. Those parameters are updated according to the opponent's dominant moves. If no dominant move is detected, the current values of those parameters are not updated.

Algorithm 1 outlines how the Solver Agent generates its offers. Initially, it uses the *time-based* strategy alone (lines 10–11), then

Table 2: Dominant Move Type and Actions

Dominant Move Type	Action
Fortunate	Decrease Concession Rate (P_1)
Silent	Decrease the Empathy Score (μ)
Selfish	Increase the Empathy Score (μ)
Concession	Increase Concession Rate (P_1)
	Increase Time-Based Target Utility (P_2)

calculates P_E (line 13). Before n rounds elapse, it updates the offer using the hybrid tactic (line 15). After n rounds, it updates P_A and refines the tactic parameters (lines 18–22). Finally, it compares the new Nash offer (i.e., the offer maximizing the product of utilities) and picks the more beneficial one (line 23). This cyclical approach ensures real-time adaptation to both the *emotional* and *behavioral* signals.

Algorithm 1 Solver Agent’s Offer Strategy

```

1:  $t_{cur}$ : current time,  $O_j^{t_{cur}}$ : Nao’s current offer
2:  $O_h^{t_{cur}}$ : human opponent’s current offer
3:  $nash_{offer}$ : generated Nash offer
4:  $U(nash_{offer})$ : utility of the Nash offer for Nao
5:  $U(O_h^t)$ : utility of the human opponent’s offer for Nao
6:  $H_o$ : human opponent’s bid history,  $A_o$ : Nao’s bid history
7:  $EstOpp_{pref}$ : estimated opponent’s preference profile
8:  $tactic_{time}$ : Nao’s time-based bidding tactic
9:  $tactic_{solver}$ : Nao’s time+behavior-based bidding tactic
10: if  $|H_o| < 2$  then
11:    $O_j^{t_{cur}} \leftarrow \text{generateOffer}(tactic_{time})$ 
12: else
13:    $P_E \leftarrow \text{updateEmotionEffect}(opponent_{emotions})$ 
14:   if  $|H_o| < n$  then
15:      $O_j^{t_{cur}} \leftarrow \text{generateOffer}(tactic_{solver})$ 
16:   else
17:      $EstOpp_{pref} \leftarrow \text{updateOpponentProfile}(H_o)$ 
18:      $P_A \leftarrow \text{updateAwareness}(A_o, H_o, EstOpp_{pref})$ 
19:      $S_T \leftarrow \text{updateSensitivityClass}(H_o, EstOpp_{pref})$ 
20:      $tactic_{solver} \leftarrow \text{updateTacticParams}(S_T, P_A, P_E)$ 
21:      $O_j^{t_{cur}} \leftarrow \text{generateOffer}(tactic_{solver})$ 
22:      $nash_{offer} \leftarrow \text{generateNashOffer}(EstOpp_{pref})$ 
23:     if  $U(nash_{offer}) > U(O_j^{t_{cur}})$  then
24:        $O_j^{t_{cur}} \leftarrow nash_{offer}$ 
25:     end if
26:   end if
27: end if

```

The Solver Agent *extends* a standard time-and-behavior approach by *continuously* factoring in user emotion (P_E) and awareness (P_A). This design aims to (1) respond quickly to genuine frustration or contentment and (2) reduce overreaction to deceptive signals. By combining speech-based negotiation, facial expression input, and opponent modeling, the Solver Agent provides an *adaptive* and *contextually responsive* negotiation platform.

4 Evaluation

We conducted in-person user experiments to assess our proposed negotiation agents, where a humanoid robot (Nao) negotiated with participants over resource allocation. Each participant interacted with Nao in *two* separate sessions: (1) one session employing the *Hybrid* strategy and (2) a second session deploying the *Solver Agent*. We randomized the order of these sessions (counterbalancing) to minimize learning or fatigue effects: half of the participants encountered the Solver Agent first, and the other half faced its second session. Both agents use an *AC-Next* acceptance condition [13], accepting an offer if its utility for the agent is at least as high as that of the agent’s *next* offer’s target utility. We obtained Institutional Review Board (IRB) approval from XYZ University; all data were masked (e.g., facial expression logs, offer transcripts) to address ethical concerns, and no participant reported physical or emotional harm.

4.1 Experimental Scenario

We frame the negotiation as a *supermarket scenario* where Nao and the human participant must *split* four types of fruits (each with four units) to achieve a minimum target score, which is 40. If the participant total is less than 40, they earn a score of zero. A 10-minute deadline ensures time pressure. All they need is to find an agreement on how they will share the fruits between them. As shown in Table 3, each fruit has a different preference score for Nao and the participant, and these preferences *change* between the first and second sessions. Participants are told that Nao does *not* know their preferences.

Table 3: Preference Profiles for Negotiation Sessions

Items	First Negotiation		Second Negotiation	
	Nao’s Preferences	Participant’s Preferences	Nao’s Preferences	Participant’s Preferences
Watermelon	4	12	12	4
Banana	1	8	8	1
Orange	12	4	4	12
Apple	8	1	1	8

Before the first session, a demo video explains the interaction protocol to the participants. They also perform a simpler ‘Demo negotiation’ to familiarize themselves with the protocol by employing Hybrid Agent. The demo negotiation session has 5 5-minute deadline. After the training session, participants receive their preference profile for their negotiation session. After studying the given profile, they are asked to negotiate with Nao accordingly. After completing their negotiation, participants fill out a questionnaire form regarding their first negotiation. After a 10-minute break, they start a second negotiation with a new preference profile. They are told that their preferences are utterly different to prevent the learning effect. However, only the order of the scores is changed, and the value distribution of the scores remains the same for a fair evaluation. During their negotiation, participants can see their preference profile, offer information, remaining time, and whose turn it is (Figure 4).



Figure 4: Experiment Setup from Participant's Perspective

4.2 Participants and Priorities

We recruited 28 participants (23 men and 5 women), with a mean age of 23.7 years, and all were university students (B.Sc., M.Sc., Ph.D.). Before negotiating, each participant ranked four priorities from highest (4) to lowest (1): (i) *decreasing the opponent's utility*, (ii) *decreasing agreement time*, (iii) *increasing their utility*, and (iv) *finding the best deal for both sides*. Figure 5 illustrates the distribution of these self-reported priorities, showing various negotiation perspectives, where the y-axis represents the number of participants that prioritized that question with that order (e.g., "Finding the best offer" is ranked 1st, denoted by blue, by 7 participants).




Figure 5: Ranking of Participants' Priorities

We also applied K-means clustering to each participant's rank pattern, identifying four groups (*competitive*, *individualist*, *selfish*, *prosocial*) as seen in Figure 6. To give a meaningful category name

for each cluster, we examined the centroid points of each cluster. Note that the inner and outermost tiles correspond to the least and most essential priority. Here, the yellow category consists of participants caring about decreasing their opponent's score, but the importance of their score is the least important. Therefore, we called this group as "competitive". The grey category consists of participants who aim to maximize their utility at most; therefore, we called them "individualist". The blue group cares about decreasing the opponent's score at most while not considering the best deal for all ("Selfish"). We called the participants who care most about both sides' scores as "prosocial". Those categories will be used in the detailed analysis of the negotiation outcomes.




Figure 6: Participant Clusters According to Priorities

4.3 Experimental Results

Figure 7 shows the average *agent* and *user* scores across the two conditions, plus the normalized product of scores. On average, the agent score is *significantly higher* with the Solver Agent than with the Hybrid strategy (0.73 vs. 0.68). A two-tailed paired sample t-test reveals $t = -2.134$ at $p = 0.042$, indicating a medium-large effect size (Cohen's $d = 0.5056$). In contrast, users' scores remain roughly the same (0.77 vs. 0.79, $p = 0.302$), suggesting that *emotional awareness* increases the agent's score without decreasing the user's score.

We also compute a *normalized score product* by dividing the product of user and agent scores by the Nash product (0.64 in our scenario). The Solver Agent yields a higher average product than the Hybrid (0.87 vs. 0.82), $t = -1.720$ at $p = .096$. This indicates a slightly better *joint utility* when the agent uses emotional signals, but not significantly.

Cluster Analysis. Table 4 reports agent vs. user scores by cluster. Given our small sample, we observe that "Selfish" participants yield higher agent scores overall, but do not find significant differences across clusters. Future work may reexamine these patterns with a larger population.

Time and Number of Bids. Figure 8 shows that *average normalized agreement time* is 0.41 for the Solver Agent vs. 0.51 for the




Figure 7: Individual & Normalized Product Scores

Table 4: Avg. Agent & User Score According to Priority Types

Priority Types	Agent Score $\pm$ STD		User Score $\pm$ STD	
	Solver Agent	Hybrid Agent	Solver Agent	Hybrid Agent
Selfish (9)	0.74 $\pm$ 0.07	0.72 $\pm$ 0.12	0.73 $\pm$ 0.08	0.74 $\pm$ 0.10
Prosocial (8)	0.72 $\pm$ 0.07	0.64 $\pm$ 0.10	0.80 $\pm$ 0.05	0.81 $\pm$ 0.06
Individualist (6)	0.74 $\pm$ 0.07	0.71 $\pm$ 0.12	0.78 $\pm$ 0.09	0.80 $\pm$ 0.05
Competitive (5)	0.70 $\pm$ 0.07	0.62 $\pm$ 0.15	0.77 $\pm$ 0.14	0.80 $\pm$ 0.13

Hybrid. A paired t-test shows $t = 1.795$ and $p = 0.093$, which is *not* significant at $\alpha = 0.05$. However, the total number of bids is significantly lower (14.96 vs. 19.39, $p = 0.048$), suggesting that while both strategies often reach agreement within 10 minutes, the Solver Agent requires *fewer* offers, indicating more *efficient* negotiation.




Figure 8: Average Total Offers and Agreement Time per Session

We also examined the responses to the 9-point questionnaire in Table 5, where 1 = “strongly disagree” and 9 = “strongly agree.” In most of the questions, participants gave similar ratings for Solver vs. Hybrid, except for Q5 (“Nao cared about my preferences”, which is notably higher for Solver (7.71 vs. 6.82). Statistical tests (t-test or Wilcoxon when normality was violated as Q1, Q5, Q6, and Q7) confirm a medium-large effect (Cohen’s $d = 0.54$, $p = 0.047$). This aligns with our hypothesis that emotional adaptation leads participants to perceive Nao as *more considerate*, although some participants did not consciously notice emotional signals (see Q3).

Interestingly, although the Solver Agent utilizes real-time facial-expression feedback, participants did not strongly perceive that Nao was adapting its offers based on their emotional states (see Q3 in Table 5). One potential explanation is that Solver Agent operates

behind the scenes: it modifies its utility target and concession rate based on users’ facial-expression cues (Eq. 6), but it does *without* announcing that it is reacting to negative or positive affect. As a result, the user’s focus on reaching a better agreement may overshadow any awareness of the agent’s internally adaptive mechanisms. Despite the lack of perceived emotional adaptation reported in Q3, participants rated the Solver Agent significantly higher in Q5, indicating that even if users did not *explicitly* detect emotion-driven changes, they found the Solver Agent *more empathetic* overall.

Table 5: Questionnaire Results

Questions	Points $\pm$ STD	
	Solver Agent	Hybrid Agent
1) Nao negotiated fairly.	7.39 $\pm$ 1.73	6.75 $\pm$ 2.03
2) Nao negotiated with me like a human.	7.39 $\pm$ 1.23	7.21 $\pm$ 1.13
3) Nao determined her next offer according to my emotional state.	6.61 $\pm$ 1.68	6.29 $\pm$ 2.08
4) Nao tried to find the best deal for us.	7.25 $\pm$ 1.73	6.50 $\pm$ 2.04
5) Nao cared about my preferences.	7.71 $\pm$ 1.27	6.82 $\pm$ 1.94
6) Nao considers my behavior.	6.89 $\pm$ 1.81	6.79 $\pm$ 1.57
7) I’m satisfied with my performance.	7.54 $\pm$ 1.31	6.82 $\pm$ 1.78
8) Nao often made very unfair offers.	3.89 $\pm$ 2.39	4.39 $\pm$ 2.23

5 Conclusion and Future Work

This paper presented a *Solver Agent* that integrates an opponent’s *emotional state* and *awareness* of the agent’s own changing behavior into a negotiation strategy. By combining time- and behavior-based concessions with real-time *facial expression* inputs, we aimed to effectively adapt to user frustration or satisfaction, encouraging faster and more beneficial agreements. We evaluated our approach in *human-robot* negotiation experiments with 28 participants, countering the Solver Agent with a Hybrid baseline for finding answers to three research questions:

- **RQ1:** Utility of the agent reached significantly higher when human participants negotiate with an emotionally aware agent (0.73 versus 0.68) while the total number of bids is significantly lower (14.96 vs. 19.39).
- **RQ2:** Participants thought that Nao cared about their preferences while negotiating an emotionally aware agent significantly (7.71 vs. 6.82).
- **RQ3:** Due to the lack of participants ($N=28$), the RQ3 is still an open question since the participant clusters cannot exceed at least 10 or differentiate other clusters according to the pre-negotiation participant priorities survey.

As a summary of the contribution, the Solver Agent outperforms the Hybrid approach in terms of (1) *agent score*, (2) *fewer overall bids*, and (3) improved user perception of “caring.” Meanwhile, user scores remain statistically unaffected, suggesting that integrating emotional signals can be beneficial without sacrificing user utility. Subjective feedback also indicates that the participants viewed the emotion-aware agent as *more considerate* of their preferences. These results highlight the potential of *emotion-driven* adaptation in embodied negotiation, although some participants did not consciously recognize the agent’s affective cues. Larger-scale studies with participants from varied backgrounds would further validate these findings and determine whether emotion-awareness scales

beyond a student population. This study shows the feasibility of emotion-aware negotiation with a physically embodied agent.

Future work may explore how *transparency* or user-acknowledged affect tracking could shape perceptions of fairness and trust, and whether more diverse participant pools or multi-round negotiations lead to different outcomes. We plan to investigate *dimensional* as shown in Geneva Emotion Wheel [8] and *appraisal-based* emotion models [23] to handle richer social conflicts and more precise emotional signals. Embedding *contextual information* within the argumentation modules may further enhance the agent's negotiation logic. Since human arguments can contain *domain-specific preferences* and *sentiment information*, taking advantage of these signals in a *multimodal interaction* framework could yield deeper insights. Finally, testing with agents of varying anthropomorphic characteristics (e.g., ABOT The database [34]) would clarify how the physical embodiment modulates user trust and compliance.

Acknowledgments

This work has been supported by a grant of The Scientific and Research Council of Turkey (TÜBİTAK) with grant number 118E197. This article's contents reflect the authors' ideas and positions and do not necessarily reflect the opinions or positions of TÜBİTAK.

References

- [1] Reyhan Aydoğan, David Festen, Koen V. Hindriks, and Catholijn M. Jonker. 2017. Alternating Offers Protocols for Multilateral Negotiation. In *Modern Approaches to Agent-based Complex Automated Negotiation*. Springer, 153–167.
- [2] Reyhan Aydoğan, Onur Keskin, and Umut Çakan. 2022. Would You Imagine Yourself Negotiating With a Robot, Jennifer? Why Not? *IEEE Transactions on Human-Machine Systems* 52, 1 (2022), 41–51.
- [3] Tim Baarslag, Reyhan Aydoğan, Koen V. Hindriks, Katsuhide Fujita, Takayuki Ito, and Catholijn M. Jonker. 2015. The Automated Negotiating Agents Competition, 2010–2015. *AI Magazine* 36, 4 (2015), 115–118.
- [4] Tim Baarslag, Enrico H. Gerding, Reyhan Aydoğan, and Monica M. C. Schraefel. 2015. Optimal Negotiation Decision Functions in Time-Sensitive Domains. In *WI-IAT 2015, Volume II*. IEEE Computer Society, 190–197.
- [5] Lisa Feldman Barrett, Ralph Adolphs, Stacy Marsella, Aleix M. Martinez, and Seth D. Pollak. 2019. Emotional Expressions Reconsidered: Challenges to Inferring Emotion From Human Facial Movements. *Psychological Science in the Public Interest* 20, 1 (2019), 1–68. doi:10.1177/1529100619832930 arXiv:https://doi.org/10.1177/1529100619832930 PMID: 31313636.
- [6] Chris Bevan and Danaë Stanton Fraser. 2015. Shaking Hands and Cooperation in Tele-present Human-Robot Negotiation. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction, HRI 2015, Portland, OR, USA*. ACM, 247–254.
- [7] Umut Çakan, M. Onur Keskin, and Reyhan Aydoğan. 2023. Effects of Agent's Embodiment in Human-Agent Negotiations. In *Proceedings of the 23rd ACM International Conference on Intelligent Virtual Agents (IVA '23)*. New York, NY, USA, Article 11, 8 pages. doi:10.1145/3570945.3607362
- [8] Adam K. Coyne, Andrew Murtagh, and Conor McGinn. 2020. Using the Geneva Emotion Wheel to Measure Perceived Affect in Human-Robot Interaction. In *2020 15th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 491–498.
- [9] Dave de Jonge and Carles Sierra. 2015. NB³: a multilateral negotiation algorithm for large, non-linear agreement spaces with limited time. *Auton. Agents Multi Agent Syst.* 29, 5 (2015), 896–942.
- [10] Celso M. de Melo, Peter J. Carnevale, and Jonathan Gratch. 2011. The effect of expression of anger and happiness in computer agents on negotiations with humans. In *AAMAS 2011*. 937–944.
- [11] Rahul R. Divekar, Hui Su, Jeffrey O. Kephart, Maira Gatti de Bayser, Melina Albeiro Guerra, Xiangyang Mou, Matthew Peveler, and Lisha Chen. 2020. HUMAINE: Human Multi-Agent Immersive Negotiation Competition. In *Conference on Human Factors in Computing Systems CHI*. ACM, 1–10.
- [12] Peyman Faratin, Carles Sierra, and Nick R Jennings. 1998. Negotiation decision functions for autonomous agents. *Robotics and Autonomous Systems* 24, 3–4 (1998), 159–182.
- [13] Shaheen Fatima, Sarit Kraus, and Michael Wooldridge. 2014. *Principles of automated negotiation*. Cambridge University Press.
- [14] Onat Güngör, Umut Çakan, Reyhan Aydoğan, and Pinar Öztürk. 2020. Effect of Awareness of Other Side's Gain on Negotiation Outcome, Emotion, Argument and Bidding Behavior. In *Recent Advances in Agent-based Negotiation*. Springer.
- [15] Galit Haim, Ya'akov (Kobi) Gal, Michele Gelfand, and Sarit Kraus. 2012. A cultural sensitive agent for human-computer negotiation. In *AAMAS 2012, Valencia, Spain, June 4–8, 2012 (3 Volumes)*. IFAAMAS, 451–458.
- [16] Koen V. Hindriks, Catholijn M. Jonker, and Dmitry Tykhonov. 2011. Let's dans! An analytic framework of negotiation dynamics and strategies. *Web Intelligence and Agent Systems* 9, 4 (2011), 319–335.
- [17] Catholijn M. Jonker and Reyhan Aydoğan. 2021. Deniz: A Robust Bidding Strategy for Negotiation Support Systems. In *Advances in Automated Negotiations*. Springer Singapore, Singapore, 29–44.
- [18] C. M. Jonker, R. Aydoğan, T. Baarslag, J. Broekens, C. A. Detweiler, K. V. Hindriks, A. Hultgren, and W. Pasman. 2016. An Introduction to the Pocket Negotiator: A General Purpose Negotiation Support System. In *EUMAS 2016*. Springer, 13–27.
- [19] Gregory Kersten and Gordon Lo. 2003. Aspire: An integrated negotiation support system and software agents for e-business negotiation. *Int. Journal of Internet and Enterprise Management* 1 (2003), 293–315.
- [20] Mehmet Onur Keskin, Berk Buzcu, and Reyhan Aydoğan. 2023. Conflict-based negotiation strategy for human-agent negotiation. *Applied Intelligence* 53 (11 2023), 1–17. doi:10.1007/s10489-023-05001-9
- [21] Mehmet Onur Keskin, Berk Buzcu, Berkecan Koçyiğit, Umut Çakan, Anil Doğru, and Reyhan Aydoğan. 2024. NEGOTIATOR: A comprehensive framework for human-agent negotiation integrating preferences, interaction, and emotion. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence (Jeju, Korea) (IJCAI '24)*. Article 1012, 4 pages. doi:10.24963/ijcai.2024/1012
- [22] Mehmet Onur Keskin, Umut Çakan, and Reyhan Aydoğan. 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*. 1557–1559.
- [23] Su Lei and Jonathan Gratch. 2023. Emotional Expressivity is a Reliable Signal of Surprise. *IEEE Transactions on Affective Computing* 14, 4 (2023), 2913–2924. doi:10.1109/TAFFC.2023.3234015
- [24] Mike Lewis, Denis Yarats, Yann Dauphin, Devi Parikh, and Dhruv Batra. 2017. Deal or No Deal? End-to-End Learning of Negotiation Dialogues. In *EMNLP*. 2443–2453.
- [25] Shan Li and Weihong Deng. 2019. Reliable Crowdsourcing and Deep Locality-Preserving Learning for Unconstrained Facial Expression Recognition. *IEEE Transactions on Image Processing* 28, 1 (2019), 356–370.
- [26] Raz Lin and Sarit Kraus. 2010. Can automated agents proficiently negotiate with humans? *Commun. ACM* 53, 1 (2010), 78–88.
- [27] Raz Lin, Sarit Kraus, Jonathan Wilkenfeld, and James Barry. 2008. Negotiating with bounded rational agents in environments with incomplete information using an automated agent. *Artif. Intell.* 172, 6–7 (2008), 823–851.
- [28] Johnathan Mell and Jonathan Gratch. 2016. IAGO: Interactive Arbitration Guide Online (Demonstration). In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems, Singapore, May 9–13, 2016*. ACM, 1510–1512.
- [29] Johnathan Mell and Jonathan Gratch. 2017. Grumpy & Pinocchio: Answering Human-Agent Negotiation Questions through Realistic Agent Design. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, AAMAS*. ACM, 401–409.
- [30] Johnathan Mell, Jonathan Gratch, Tim Baarslag, Reyhan Aydoğan, and Catholijn M. Jonker. 2018. Results of the First Annual Human-Agent League of the Automated Negotiating Agents Competition. In *18th Int. Conference on Intelligent Virtual Agents, IVA*. ACM, 23–28.
- [31] Yinon Oshrat, Raz Lin, and Sarit Kraus. 2009. Facing the challenge of human-agent negotiations via effective general opponent modeling. In *AAMAS (1)*. IFAAMAS, 377–384.
- [32] Lydia Ould Ouali, Nicolas Sabouret, and Charles Rich. 2017. A Computational Model of Power in Collaborative Negotiation Dialogues. In *Intelligent Virtual Agents*, Jonas Beskow, Christopher Peters, Ginevra Castellano, Carol O'Sullivan, Iolanda Leite, and Stefan Kopp (Eds.). Springer, Cham, 259–272.
- [33] Noam Peled, Ya'akov (Kobi) Gal, and Sarit Kraus. 2015. A study of computational and human strategies in revelation games. *Autonomous Agents and Multi-Agent Systems* 29, 1 (2015), 73–97.
- [34] Elizabeth Phillips, Xuan Zhao, Daniel Ullman, and Bertram F. Malle. 2018. What is Human-like?: Decomposing Robots' Human-like Appearance Using the Anthropomorphic roBOT (ABOT) Database. In *2018 13th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 105–113.
- [35] Davide Pietroni, Gerben A Van Kleef, Carsten KW De Dreu, and Stefano Pagliaro. 2008. Emotions as strategic information: Effects of other's emotional expressions on fixed-pie perception, demands, and integrative behavior in negotiation. *Journal of Experimental Social Psychology* 44, 6 (2008), 1444–1454.
- [36] Pooja Prajod, Mohammed Al Owayyed, Tim Rietveld, Jaap-Jan van der Steeg, and Joost Broekens. 2019. The Effect of Virtual Agent Warmth on Human-Agent Negotiation. In *18th International Conference on Autonomous Agents and MultiAgent Systems (Montreal QC, Canada) (AAMAS '19)*. 71–76.

- [37] Raquel Ros and Carles Sierra. 2006. A Negotiation Meta Strategy Combining Trade-off and Concession Moves. *Autonomous Agents and Multi-Agent Systems* 12 (03 2006), 163–181.
- [38] Avi Rosenfeld, Inon Zuckerman, Erel Segal-Halevi, Osnat Drein, and Sarit Kraus. 2014. NegoChat: a chat-based negotiation agent. In *AAMAS. IFAAMAS/ACM*, 525–532.
- [39] Magdalena Rychlowska, Rachael E. Jack, Oliver G. B. Garrod, Philippe G. Schyns, Jared D. Martin, and Paula M. Niedenthal. 2017. Functional Smiles: Tools for Love, Sympathy, and War. *Psychological Science* 28, 9 (2017), 1259–1270. doi:10.1177/0956797617706082 arXiv:https://doi.org/10.1177/0956797617706082 PMID: 28741981.
- [40] Victor Sanchez-Anguix, Reyhan Aydoğan, Vicente Julian, and Catholijn Jonker. 2014. Unanimously acceptable agreements for negotiation teams in unpredictable domains. *Electronic Commerce Research and Applications* 13, 4 (2014), 243–265.
- [41] Marwan Sinaceur and Larissa Z Tiedens. 2006. Get mad and get more than even: When and why anger expression is effective in negotiations. *Journal of Experimental Social Psychology* 42, 3 (2006), 314–322.
- [42] Leila Takayama, Victoria Groom, and Clifford Nass. 2009. I’m Sorry, Dave: I’m Afraid i Won’t Do That: Social Aspects of Human-Agent Conflict. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09)*, 2099–2108.
- [43] Okan Tunali, Reyhan Aydoğan, and Víctor Sánchez-Anguix. 2017. Rethinking Frequency Opponent Modeling in Automated Negotiation. In *20th International Conference Principles and Practice of Multi-Agent Systems*, Vol. 10621. Springer, 263–279.
- [44] Rustam M. Vahidov, Gregory E. Kersten, and Bo Yu. 2017. Human-Agent Negotiations: The Impact Agents’ Concession Schedule and Task Complexity on Agreements. In *50th Hawaii Int. Conference on System Sciences, HICSS, Hawaii, USA*, 1–9.
- [45] Wim F. J. van der Ham, Joost Broekens, and Peter H. M. P. Roelofsma. 2014. The Effect of Dominance Manipulation on the Perception and Believability of an Emotional Expression. In *Emotion Modeling - Towards Pragmatic Computational Models of Affective Processes*. Lecture Notes in Computer Science, Vol. 8750. Springer, 101–114.
- [46] Gerben A Van Kleef, Carsten KW De Dreu, and Antony SR Manstead. 2004. The interpersonal effects of anger and happiness in negotiations. *Journal of personality and social psychology* 86, 1 (2004), 57.
- [47] Thijs van Krimpen, Daphne Looije, and Siamak Hajizadeh. 2013. HardHeaded. In *Complex Automated Negotiations: Theories, Models, and Software Competitions*. Studies in Computational Intelligence, Vol. 435. Springer, 223–227.